

CROSSCUT REPORT









EXASCALE REQUIREMENTS REVIEWS

March 9–10, 2017 – Tysons Corner, Virginia

An Office of Science review sponsored by:
Advanced Scientific Computing Research
Basic Energy Sciences
Biological and Environmental Research
Fusion Energy Sciences
High Energy Physics
Nuclear Physics



U.S. DEPARTMENT OF
ENERGY



On the cover

Left column, top to bottom:

1. *Appears in the **Fusion Energy Sciences** Exascale Requirements Review:* Electromagnetic field simulation using the TORIC solver showing incoming long-, intermediate-, and short-wavelength waves in the Alcator C-Mod tokamak (Source: Bonoli, Parker, and Wukitch et al. 2007).
2. *Appears in the **Basic Energy Sciences** Exascale Requirements Review:* Direct numerical simulation of a turbulent di-methyl ether lifted jet flame with multistage ignition in the negative temperature coefficient regime at 5 atm (Source: Minamoto and Chen 2016).
3. *Appears in the **Biological and Environmental Research** Exascale Requirements Review:* Model of the cellulase synthase enzyme, CesaA, derived by integrating neutron scattering and high-performance computing (Image credit: Splettstoesser, scistyle.com).
4. *Appears in the **High Energy Physics** Exascale Requirements Review:* One of seven tape libraries at the Feynman Computing Center at Fermilab; each library can hold 10,000 tapes. More than 150 PB of experimental and observational data are stored (total storage capacity exceeds 1000 PB).
5. *Appears in the **Nuclear Physics** Exascale Requirements Review:* Quantum fluctuations of gluons captured in a field configuration. The solid regions show enhanced action density, and the vectors show the space-time orientation of one of the gluon fields (Image reproduced with permission from D. Leinweber).

Right column, top to bottom:

1. *Appears in the **Basic Energy Sciences** Exascale Requirements Review:* A 16,611-atom quantum molecular dynamics simulation of H_2 production from water using a LiAl-alloy particle. Produced H_2 molecules are represented by green ellipsoids; water molecules are not shown for clarity (Source: Shimamura et al. 2014).
2. *Appears in the **Biological and Environmental Research** Exascale Requirements Review:* Simulation analysis of pretreated lignocellulosic biomass. The extension of molecular dynamics simulations to large macromolecular assemblies will allow long length-scale cooperative behavior to be understood (Image credit: Splettstoesser, scistyle.com).
3. *Appears in the **High Energy Physics** Exascale Requirements Review:* The ATLAS detector under construction at CERN in 2007.
4. *Appears in the **Fusion Energy Sciences** Exascale Requirements Review:* Turbulence spreading in from the edge to the core. Instabilities in the steep gradient edge region drive turbulence to propagate radially inward (Source: Chowdhury et al. 2014).

Disclaimer

This report was prepared as an account of a meeting sponsored by the U.S. Department of Energy. Neither the United States Government nor any agency thereof, nor any of their employees or officers, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of document authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof. Copyrights to portions of this report (including graphics) are reserved by original copyright holders or their assignees, and are used by the Government's license and by permission. Requests to use any images must be made to the provider identified in the image credits.

CROSSCUT REPORT

EXASCALE REQUIREMENTS REVIEWS

Crosscut Organizing Committee

James J. Hack, Oak Ridge National Laboratory/Oak Ridge Leadership Computing Facility
Katherine Riley, Argonne National Laboratory/Argonne Leadership Computing Facility
Richard Gerber, Lawrence Berkeley National Laboratory/National Energy
Research Scientific Computing Center

Report Section Leads

Katie Antypas, National Energy Research Scientific Computing Center
Richard Coffey, Argonne Leadership Computing Facility
Eli Dart, ESnet
Richard Gerber, National Energy Research Scientific Computing Center
James J. Hack, Oak Ridge National Laboratory/Oak Ridge Leadership Computing Facility
Katherine Riley, Argonne Leadership Computing Facility
Tjerk Straatsma, Oak Ridge Leadership Computing Facility
Jack C. Wells, Oak Ridge Leadership Computing Facility

Sponsors and Representatives

Barbara J. Helland, DOE/Office of Advanced Scientific Computing Research
Carolyn B. Lauzon, DOE/Office of Advanced Scientific Computing Research

Additional Contributors

Argonne National Laboratory Document Design and Development Team:
Mary Fitzpatrick, Andrea Manning, and Vicki Skonicki, Business & Information Systems
Michele Nelson, Communications, Education, and Public Affairs
Beth Cerny and Richard Coffey, Argonne Leadership Computing Facility

This page is intentionally left blank.

TABLE OF CONTENTS

Executive Summary	1
ES.1 Introduction and Background.....	1
ES.2 Findings.....	2
ES.2.1 Computing	2
ES.2.2 Data	2
ES.2.3 Software and Application Development	3
ES.2.4 Training and Workforce Development	3
ES.2.5 Opportunities for Collaboration.....	4
1 Overview of Exascale Requirements Reviews	5
2 Priority Scientific Research Directions	7
2.1 Priority Research Directions in Basic Energy Sciences (BES)	7
2.2 Priority Research Directions in Biological and Environmental Research (BER)	8
2.2.1 Biological Systems Science	8
2.2.2 Climate and Environmental Sciences	9
2.3 Priority Research Directions in Fusion Energy Sciences (FES).....	10
2.3.1 Fusion Energy Science.....	10
2.3.2 Plasma Surface Interactions and Structural Materials	11
2.3.3 Discovery Plasma Science.....	11
2.3.4 Verification, Validation, and Uncertainty Quantification	11
2.4 Priority Research Directions in High Energy Physics (HEP)	12
2.5 Priority Research Directions in Nuclear Physics (NP)	13
2.6 Priority Research Directions for Advanced Scientific Computing Research (ASCR).....	14
3 Findings.....	15
3.1 Computing	15
3.1.1 Hardware	15
3.1.2 Allocations, Access, and Policies	16
3.2 Data	17
3.2.1 Large-Scale Data Storage and Advanced Analytics	18
3.2.2 Experimental and Simulation Workflows	18
3.2.2.1 Steering of Experiments and Co-scheduling.....	19
3.2.2.2 Direct Streaming of Large Datasets from Experiments	19
3.2.2.3 Integration of Simulation and Experimental Data for Analysis.....	19
3.2.2.4 Facilitating Analysis Workflows across Multiple, Distributed Facilities.....	19

3.2.3	Data Management, Archiving, and Curation.....	19
3.2.4	I/O Performance	20
3.2.5	Remote Access, Sharing, and Data Transfer.....	20
3.3	Software and Application Development Environment	21
3.3.1	Models, Methods, and Algorithms	21
3.3.2	Software Portability, Performance, and Sustainability.....	22
3.3.3	Common Environment/Sensible Standards	23
3.4	Training and Workforce Development	24
3.4.1	Workforce Development	24
3.4.2	Training	25
3.4.3	Existing Programs	25
4	Path Forward	27
4.1	Areas for Collaboration	27
4.1.1	Co-evolving ASCR Facilities and DOE Experimental Science	27
4.1.2	Addressing Data-related Needs and Developing New Methods of Analysis and Data Management	27
4.1.3	Pursuing Formal, Close Collaborations between Application Scientists and ASCR Experts	28
4.1.4	Creating an Easy-to-use HPC Environment That Is Integrated across ASCR Centers and Other Work Environments	28
4.1.5	Sharing the Scale-of-Data “Burden” among Facilities	28
4.1.6	Developing Closer Collaborations among ASCR Research and Facilities to Support DOE Science.....	28
4.1.7	Recruiting and Providing Training to DOE Researchers Who Are Late Adopters.....	29
4.1.8	Supporting Training and Workforce Development.....	29
4.1.9	Continuing to Collect Requirements, Identify Needs, and Deliver HPC Resources in Support of Science	30
4.2	Ongoing Facility Efforts and Programs Already Under Way to Reach an Exascale Ecosystem	30
4.2.1	ESnet.....	30
4.2.2	LCF Programs.....	31
4.2.3	NERSC.....	31
4.2.4	Exascale Computing Project (ECP)	31
5	References.....	33
6	Acronyms and Abbreviations	35

Appendices.....	37
Appendix A Executive Summary from the Basic Energy Sciences Exascale Requirements Review Report.....	A-1
Appendix B Executive Summary from the Biological and Environmental Research Exascale Requirements Review Report.....	B-1
Appendix C Executive Summary from the Fusion Energy Sciences Exascale Requirements Review Report.....	C-1
Appendix D Executive Summary from the High Energy Physics Exascale Requirements Review Report.....	D-1
Appendix E Executive Summary from the Nuclear Physics Exascale Requirements Review Report.....	E-1
Appendix F Executive Summary from the Advanced Scientific Computing Research Exascale Requirements Review Report.....	F-1
Appendix G Sections Cited in This Report from the Exascale Requirements Review for Basic Energy Sciences (BES).....	G-1
Appendix H Sections Cited in This Report from the Exascale Requirements Review for Biological and Environmental Research (BER).....	H-1
Appendix I Sections Cited in This Report from the Exascale Requirements Review for Fusion Energy Sciences (FES).....	I-1
Appendix J Sections Cited in This Report from the Exascale Requirements Review for High Energy Physics (HEP).....	J-1
Appendix K Sections Cited in This Report from the Exascale Requirements Review for Nuclear Physics (NP).....	K-1
Appendix L Sections Cited in This Report from the Exascale Requirements Review for Advanced Scientific Computing Research (ASCR)	L-1

This page is intentionally left blank.

EXECUTIVE SUMMARY

ES.1 Introduction and Background

The mission of the U.S. Department of Energy Office of Science (DOE SC) is the delivery of scientific discoveries and major scientific tools to transform our understanding of nature and to advance the energy, economic, and national security missions of the United States. To achieve these goals in today's world requires investments in not only the traditional scientific endeavors of theory and experiment, but also in computational science and the facilities that support large-scale simulation and data analysis.

The Advanced Scientific Computing Research (ASCR) program addresses these challenges in the Office of Science. ASCR's mission is to discover, develop, and deploy computational and networking capabilities to analyze, model, simulate, and predict complex phenomena important to DOE. ASCR supports research in computational science, three high-performance computing (HPC) facilities — the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory and Leadership Computing Facilities at Argonne (ALCF) and Oak Ridge (OLCF) National Laboratories — and the Energy Sciences Network (ESnet) at Berkeley Lab.

ASCR is guided by science needs as it develops research programs, computers, and networks at the leading edge of technologies. As we approach the era of exascale computing, technology changes are creating challenges for science programs in SC for those who need to use high performance computing and data systems effectively. Numerous significant modifications to today's tools and techniques will be needed to realize the full potential of emerging computing systems and other novel computing architectures.

To assess these needs and challenges, ASCR held a series of Exascale Requirements Reviews in 2015–2017, one with each of the six SC program offices,¹ and a subsequent Crosscut Review that sought to integrate the findings from each. Participants at the reviews were drawn from the communities of leading domain scientists, experts in computer science and applied mathematics, ASCR facility staff, and DOE program managers in ASCR and the respective program offices.

The purpose of these reviews was to identify mission-critical scientific problems within the DOE Office of Science (including experimental facilities) and determine the requirements for the exascale ecosystem that would be needed to address those challenges. The exascale ecosystem includes exascale computing systems, high-end data capabilities, efficient software at scale, libraries, tools, and other capabilities. This effort will contribute to the development of a strategic roadmap for ASCR compute and data facility investments and will help the ASCR Facility Division establish partnerships with Office of Science stakeholders. It will also inform the Office of Science research needs and agenda.

The results of the six reviews have been published in reports available on the web at <http://exascale.org/>. This report presents a summary of the individual reports and of common and crosscutting findings, and it identifies opportunities for productive collaborations among the DOE SC program offices.

¹ Advanced Scientific Computing Research, Biological and Environmental Research, Basic Energy Sciences, Fusion Energy Sciences, High Energy Physics, and Nuclear Physics.

ES.2 Findings

The science communities for each SC office drove the agendas and selection of topics in each review. Each review independently arrived at findings and needs specific to its domain, but across all six offices there were strikingly more commonalities than differences. Here we present an integrated view of requirements and findings across the reviews. The major areas of concern and need were categorized into four high-level areas: Computing, Software, Data, and Workforce Development.

ES.2.1 Computing

- The Office of Science has substantial needs for both extreme-scale computing and computing at a somewhat smaller scale that still exceeds the capability of local resources. Large numbers of medium-scale runs are needed to support uncertainty quantification (UQ) and verification and validation of models.
- Scientists' computational tasks vary widely in computational complexity and have a growing range of workflow needs that must be supported in an exascale ecosystem.
- Usable and appropriate systems greatly enhance scientists' productivity.
- Researchers need wide and early access to testbed and prototype systems to develop codes and workflows prior to deployment on leading edge production systems at scale.
- There will need to be a substantial increase in compute capacity and/or capability for each project over what is available today. For example, computational needs are expected to increase by a factor of 1,000 by 2025 in some of the High Energy Physics frontiers.
- Scheduling and allocation policies are needed to support workflow needs, especially data analysis requirements at experimental facilities: real-time, pseudo-real time, co-scheduling, variable job requirements, and allocations based on other resources like disk and memory. Multiyear commitments of HPC resources are needed to assure the success of multiyear scientific campaigns and missions.

ES.2.2 Data

- A close integration of HPC simulation and data analysis will greatly aid in interpreting the results of experiments. Such an integration will minimize data movement and facilitate interdependent workflows.
- Scientists need long-term storage space and support for analysis tools, many of which differ significantly from traditional simulation software (e.g., heavy use of Python). The actions of performing analyses of big datasets and drawing inferences based on these data are revolutionizing many fields. New approaches are needed for analyzing large datasets including advanced statistics and machine learning.
- Workflows in both simulation and analysis are becoming more complex and require the scale of computing available at ASCR HPC facilities, where they need to be accommodated. This complexity is often related to needs for data movement, including over the wide area network. Scientists at experimental facilities want to use HPC to help guide experiments in real time, which requires co-scheduling between ASCR facilities and facilities from other DOE offices.
- The complexity, volume, and rapidity of data from experiments and simulation requires data management, archiving, and curation well beyond what is in common practice today. Scientists are looking for increased collaboration with ASCR in this area.
- As science increasingly becomes a community effort, the need to share, transfer, and access data at remote sites becomes more important. Large scientific projects no longer work in isolation.
- The input/output (I/O) capabilities of large HPC systems need to scale with their computational capability, and sometimes grow faster. Simulations cannot spend excessive time blocking

on I/O, and data read/write rates can be the primary factor that limits performance of data analysis pipelines.

ES.2.3 Software and Application Development

- Mathematical, software, and algorithm development is required in all Office of Science programs to take advantage of exascale computing architectures and to meet data analysis, management, and workflow needs.
- Improved mathematical, algorithmic, and workflow methods for verification, validation, and UQ are needed for a wide array of Office of Science missions.
- Scalable data processing, data analysis, machine learning, discrete algorithms, and multiscale/multiphysics simulations are crucial for reducing and understanding the large-scale data that will be produced by exascale systems.
- All SC program communities urgently require sufficiently performance-portable and expressive programming models.
- The requirements reviews highlighted the need for improved developer productivity and the ability to package the artifacts of DOE research for broad reuse. There is a recognized need for community planning, policies, processes to improve software quality, interoperability, testing, and deployment within the DOE exascale ecosystem.
- Broad consensus exists among the reports that the sustainability of DOE software infrastructure and the productivity of DOE researchers can be advanced by establishing a common computing and data analysis environment across platforms and facilities and through the promotion of sensible, community-based standards.
- Standard mathematical libraries and frameworks must transition to new architectures in a predictable and timely fashion. Urgency is being expressed through these reports to define which common tools and libraries will be supported and sustained, especially for tools and libraries for which DOE has been the sponsor of the software development.
- Workshop participants identified a requirement to work with ASCR Facilities in the development and deployment of capabilities to continuously test and integrate application and system software under development by DOE researchers on production systems.

ES.2.4 Training and Workforce Development

- The need for significant investments in workforce development and training was a common theme in all six requirements workshops and created significant discussion in the Crosscut workshop, as well.
- The reviews also highlighted the desire for deeper partnerships with ASCR for workforce development from multiple offices.
- There is a pressing need for a workforce trained in computational science: individuals who are able to develop methods and algorithms optimized for exascale systems; effectively and efficiently use and configure those systems; and write, provide, and maintain applications software (libraries, workflows, etc.).
- Respected career paths to support an HPC workforce are needed.
- Training is critical for developing better workflows and transitioning codes to next-generation platforms.

Participants in the ASCR Research review identified additional training and workforce development-related needs, including for access to prototype systems and early access to large pre-production systems at ASCR facilities. There were calls for closer collaborations between ASCR Research and ASCR Facilities.

ES.2.5 Opportunities for Collaboration

An overarching theme in all of the reviews was the belief that effective scientific utilization of high-end capability computing requires dynamic partnerships among application scientists, applied mathematicians, computer scientists, and facility support staff. Therefore, close coordination both within and across ASCR subprograms and with partner SC offices is key to the success of the ASCR program and the Office of Science mission as a whole.

Throughout the reviews and at the Crosscut Meeting, participants identified a number of areas where collaborations among stakeholders would enable significant new capabilities.

- As the data-related needs of experimental science at DOE user facilities and other DOE-supported missions/projects continue to grow quickly, a co-evolution of capabilities, services, and practices between experimental scientists and ASCR Facilities will be required.
- Analysis and management of data both from simulation and experiment are needed; and there is an opportunity for application teams and ASCR to collaborate on creating new, scalable algorithms; build better memory management techniques; explore data-related methods; and develop end-to-end analysis pipelines to solve workflow, visualization, and optimization concerns.
- An expansion of programs like SciDAC would enable close collaborations between application scientists and ASCR experts, which leads to the development and ongoing maintenance of methods, models, algorithms, workflows, and software libraries and tools that perform well on exascale platforms.
- Creating an easy-to-use HPC environment would help integrate ASCR centers and other work environments.
- Developing closer collaborations among ASCR Research and Facilities could play a major role in preparing application teams and their codes for exascale architectures.
- Providing training could enable the broad science community to use exascale-class computing and data systems efficiently.
- Implementing formal organizational structure within DOE would help encourage and support cross-cutting collaborations among the SC program offices.

These high-level and common findings were discussed and derived at the crosscutting meeting by the scientific leaders of each review, facility staff, and DOE program managers.

1 OVERVIEW OF EXASCALE REQUIREMENTS REVIEWS

During fiscal years (FYs) 2015 and 2016, the Exascale Requirements Reviews brought together key computational domain scientists, U.S. Department of Energy (DOE) planners and administrators, and experts in computer science, applied mathematics, and scientific computation to determine the requirements for an exascale ecosystem that will support scientific research of interest to the DOE Office of Science (SC). The ecosystem includes high-performance computing (HPC) and data systems, software, services, and other programmatic or technological elements that may be needed to support forefront scientific research.

Meetings were held for each of the DOE's six SC program offices, as follows:

- The High-Energy Physics (HEP) review was held in June 2015.
- The Basic Energy Sciences (BES) review was held in November 2015.
- The Fusion Energy Sciences (FES) review was held in January 2016.
- The Biological and Environmental Research (BER) review was held in March 2016.
- The Nuclear Physics (NP) review was held in June 2016.
- The Advanced Scientific Computing Research (ASCR) was held in September 2016.

A cross-cutting review with the leads from each of the six reviews was held in Tysons Corner, Virginia, in March 2017 to discuss the material contained in this report.

Each Exascale Requirements Review resulted in a report that contained programmatic goals and HPC requirements for the given SC program office.

The goals of the reviews were to:

- Identify forefront scientific challenges that will require exascale-class computing and data systems and the opportunities for scientific discovery that could be enabled by such systems.
- Identify the key elements that must be present in an exascale ecosystem to enable scientists to address the challenges and opportunities described above.
- Describe gaps in the current HPC ecosystem that need to be addressed.
- Promote the exchange of ideas among domain scientists, computer scientists, applied mathematicians, HPC facility staff, and DOE managers to maximize the potential for use of exascale computing to advance scientific discovery.

DOE will use the Exascale Requirements Review reports to guide investments and strategic planning, specifically including efforts to:

- Articulate future HPC needs to DOE and SC management, the Office of Management and Budget, and Congress.
- Identify emerging technology, software, and services needs for SC research.
- Develop a strategic roadmap for the facilities based on scientific needs and objectives.

This page is intentionally left blank.

2 PRIORITY SCIENTIFIC RESEARCH DIRECTIONS

The drive for exascale systems originates with the scientific challenges and priority research directions (PRDs) within the community of DOE SC researchers. These were articulated in each of the six Exascale Requirements Review reports.² Here we present a summary of priority research areas for which HPC is necessary in each program office; details are available in the individual reports for each program office (exascaleage.org).

2.1 Priority Research Directions in Basic Energy Sciences (BES)

BES supports fundamental research to understand, predict, and ultimately control matter and energy across scales from atoms to molecules to materials. This research probes the electronic and magnetic properties of matter and energy — an effort that is foundational to discovery of future energy solutions for the nation. Theory, computation, and the user facilities have played key roles in the BES portfolio to provide these solutions and have been recognized as essential for the future of the program.

Priority research directions and challenges are identified within the BES Exascale Requirements Review report in eight broad areas.

■ *Novel Quantum Materials and Chemicals.*

Future energy technologies will rely on specific combinations of elements, materials, and phases whose behaviors are “emergent” or not predictable by studying their components in isolation. Developing new, predictive theories and efficient and adaptive software and exploiting the full capabilities of new and future computing architectures at BES/ASCR facilities are critical to designing chemicals and materials with desired emergent properties.

■ *Catalysis, Photosynthesis and Light Harvesting, and Combustion.*

Understanding and controlling chemical transformations and energy conversion are core to the BES mission. Current computational limitations severely restrict the size and complexity of systems that can be studied with sufficient fidelity. Emerging computing ecosystems, along with advances in theoretical methods and algorithms, will enable study of realistic heterogeneous environments at long time and length scales, significantly enhancing prediction reliability.

■ *Materials and Chemical Discovery.*

The biggest challenges to realizing the vision of materials and chemical discovery “by design” are predicting novel materials and chemicals with targeted properties and creating corresponding synthesis and degradation pathways. Progress will require advances in theory, predictive modeling capabilities, hardware resources, and experimental techniques to migrate to an adaptive, multiscale modeling paradigm.

■ *Soft Matter.*

The complexity of soft matter — polymers, liquids, gels, or amorphous solids, etc. — presents scientific as well as computational challenges in designing functional matter for applications in energy storage and production, chemical separations, enhanced oil recovery, food packaging, chip manufacturing, and health care. Researchers expect to access many interesting phenomena in soft matter with exascale resources, enabling modeling that will have a significant and broader impact on critical national needs.

² In addition, for ease of reference, Appendices A through F reprint the complete text of the Executive Summary from each Exascale Requirements Review.

■ *Advances in Algorithms for Quantum Systems.*

While advances in quantum and multiscale methods affect computational chemistry, biology, physics, and materials science, the associated high computational cost of scaling poses significant challenges. Facilitating use of these methods will involve reducing the scaling and developing multilevel parallel algorithms for them, requiring collaboration among application scientists and developers, applied mathematicians, computer scientists, and software engineers.

■ *Computing and Data Challenges at BES Facilities.*

The BES program operates major scientific user facilities where more than 240 different types of instruments require complex data acquisition and analysis methods. Future facilities will support expanded instrumentation and techniques, leading to an exponential growth in data. Managing and extracting useful scientific information from these data, combined with real-time modeling and simulations during the experiments, will require exascale computational resources.

■ *Mathematics and Computer Science Transforming BES Science.*

Bridging the gap between BES scientific goals and ASCR computing capabilities will depend on the collective abilities of the science domains and facilities to deliver breakthroughs in mathematics and computer science. These needs include improvements in speed and accuracy in predictive modeling; algorithms and software environments for fast, multimodal analysis of multisource data; and tools that make the efficient programming of tomorrow's machines as straightforward as programming today's laptops.

■ *Next-Generation Workforce.*

Fielding a sufficiently skilled workforce is a major challenge in realizing exascale computational science. Significant investments must be made in training a new generation of scientists who are not only well grounded in their technical disciplines, but who are also knowledgeable about relevant computer science and applied mathematics issues.

2.2 Priority Research Directions in Biological and Environmental Research (BER)

BER's programs are organized within two divisions: the Biological Systems Science Division (BSSD) and the Climate and Environmental Sciences Division (CESD).

2.2.1 Biological Systems Science

The core ambitions of Biological Systems Science — to predict, control, and design the function, environmental health, and productivity of plants, microbes, and the biomes they support — require discovery and characterization of the causal linkages in biomolecular networks within and across organisms and abiotic interfaces. Four PRDs address these core ambitions:

■ *Multiscale Biophysical Simulation from Molecules to Cells.*

Understanding biological function requires understanding the complex heterogeneous physics that drives living processes, which, in turn, requires a deep mechanistic description of how proteins, ribonucleic acids (RNAs), and other biomolecules perform their functions through the power of deoxyribonucleic acid (DNA) sequencing. Current modeling approaches need to be revisited to exploit future computational architectures, and new algorithms for incorporating data derived from innovations in genomics, molecular imaging, structural biology, and spectroscopy need to be developed.

■ *Mapping Sequence to Models.*

Disruptive experimental technologies and sophisticated molecular functional assays are driving the need for extraordinary computational innovation and scaling to understand the functional encoding in organisms' genomes. Current algorithms for transforming raw data into genes and genomes, protein structures, and chemical activities are computationally and storage

intensive. There are critical needs for developing new algorithms that can exploit future exascale ecosystems and for developing and maintaining such key resources as constantly updated taxonomic and gene family phylogenetic trees, open-access publications, and data in large-scale functional genomic resources.

■ ***Microbes to the Environment.***

CESD and BSSD come together to incorporate predictions and data into integrated models of biomes and their environmental functions and to propagate their effects — as contributors to the carbon, nitrogen, sulfur, and phosphorous cycles — up to global scales. These processes depend on microbial mechanisms that scale to pores in soil particles and through reactive transport in flowing watersheds. The output from these models can then help drive Earth system models from CESD. Detailed modeling of microbe-microbe, microbe-plant, and plant-plant interactions becomes critical.

■ ***Biological Big-Data Challenges.***

The size, quality, and structure of data present unique challenges to scientists in making effective predictions about biological identity, function, and behavior. Even a single bacterial cell has 10 billion molecules drawn from 10,000 chemical species; and complexity only grows with groups of cells in tissues or organisms in communities. Improvements in the algorithms used to cluster, reduce dimensionality, compute graphs of probabilistic dependencies, and generally find models of the data are essential, as are knowledge systems that also incorporate ontological systems for classifying data.

2.2.2 Climate and Environmental Sciences

Increasingly, environmental systems simulation includes physical, chemical, and biological processes spanning ever-broader ranges of spatial and temporal scales. The interactions among scales produce complex feedbacks that determine the system behaviors and have emerged as a central challenge in simulation, data collection, data management, and scientific analysis. Straddling these challenges, CESD pursues these seven PRDs:

■ ***Atmospheric Simulation and Data Assimilation within the Earth System.***

Advances in computational technology are enabling researchers to substantially improve the representation of clouds in global modeling frameworks and their impacts on energy and water cycles. Through increases in resolution, along with the appropriate parameterizations, Earth System Models (ESMs) will be able to more realistically capture heavy precipitation events; droughts; floods; and other low-probability, high-impact events. Future research also needs to develop data assimilation techniques at the appropriate spatial scales and with the appropriate targeted observations.

■ ***Terrestrial and Subsurface Research.***

Many land processes of importance to the integrated functioning of the Earth system operate, and will continue to operate, on spatial scales much finer than those represented in the host ESMs and thus will remain unresolvable in global frameworks. These phenomena interact directly with the physical and chemical environments with significant impacts at all scales. Comprehensive and robust theories must be developed that allow process knowledge to migrate effectively up in scale from process resolving to process parameterized to improve the modeling of terrestrial and subsurface processes.

■ ***Oceans and Cryospheric Research.***

The ocean-cryosphere system comprises the global ocean, including the main deep basins, marginal seas, coastal ocean, and estuaries, along with all sea-ice and land-ice systems. It is estimated that over the twentieth century, the ocean system has absorbed approximately 90% of the heat trapped by greenhouse gases, whereas the cryosphere is undergoing the most rapid

recent changes within the entire Earth system, most clearly in the form of systematic transition to summertime ice-free conditions in the Arctic. Enhancing computational capabilities may dramatically improve the fidelity of ocean-cryosphere system simulations and their coupling with other Earth system components.

■ ***Earth System Models.***

Fully coupled ESMs integrate the physical and biogeochemical components of the Earth's system to capture the many nonlinear interactions and feedbacks. In some cases, high-resolution models remove the need for approximate parameterizations in favor of directly resolving processes.

■ ***Integrated Assessment and Impacts-Adaptation-Vulnerability Modeling.***

Integrated assessment models (IAMs) have historically focused on understanding the implications of human activity on the Earth system at the global scale. For example, in coastal regions where many cities and important infrastructure assets are located, understanding human–natural process interactions is essential for evaluating vulnerability to floods and measures to mitigate flood risk. IAMs are increasingly coupled to other models involving multiple scales, processes, sectors, disciplines, institutions, and sets of heterogeneous data. These model couplings thus present a variety of theoretical, operational, and computational challenges.

■ ***Transforming Science through Exascale Capabilities: Model-Data Fusion and Testbeds.***

Use of model development testbeds will be critical to addressing model development challenges by helping scientists identify systematic errors in simulations that point to specific structural and parametric uncertainties and by providing insights into dominant processes and process interactions. However, testbeds generate large quantities of simulation output; performing efficient post-processing for model output and applying proper metrics and diagnostics packages are other challenges. Developing innovative data storage and analytics algorithms that can exploit exascale ecosystems represents a major scientific opportunity for accelerating progress in ESM development.

■ ***Transforming Science through Exascale Capabilities: Algorithms and Computational Science.***

The algorithms that will enable the use of advanced architectures need to be as broad and diverse as the problems these models seek to address and must span model testing, integration with multiple scales, coupling, analysis, and understanding. In addition, while both Fortran and C/C++ are interoperable on almost all computing platforms available today, there is a strong desire for other productivity-oriented scripting languages in a distributed environment (e.g., Python/pyMPI) and to explore alternatives to the flat MPI model. Thus, pragma-based hybrid programming models are needed; attention must also be directed to compiler, performance, and scientific portability challenges when dealing with multiple architectures and to determining how parallelism is exploited.

2.3 Priority Research Directions in Fusion Energy Sciences (FES)

The FES program mission is to expand the fundamental understanding of matter at very high temperatures and densities and to build the scientific foundation needed to develop a fusion energy source (FES 2016). This mission is accomplished through the study of plasma, the fourth state of matter, and how it interacts with its surroundings. For the purposes of this review, the FES was organized in four topical research areas: Fusion Energy Science; Plasma Surface Interactions and Structural Materials; Discovery Plasma Science; and Verification, Validation, and Uncertainty Quantification.

2.3.1 Fusion Energy Science

■ ***Turbulence and Transport in a Fusion Reactor.***

Plasma turbulence and transport determine the viability of a fusion reactor: if plasma energy is lost too quickly, fusion burn cannot occur or be sustained, and core confinement will not be

realized if plasma is not confined at the edge. Running well-resolved, full-torus gyrokinetic simulations for ITER and other future fusion reactors will require exascale-class supercomputers to obtain reliable prediction of radial temperature and density profiles.

■ ***Energetic Particles and Magnetohydrodynamic Instabilities in a Fusion Reactor.***

Energetic particles (EPs) in burning plasma experiments can readily excite mesoscale or macroscopic magnetohydrodynamic (MHD) instabilities, which can lead to disruptions or threaten a machine's integrity. Predictive capability requires exascale-level, integrated first-principles simulation of nonlinear interactions of multiple kinetic-MHD processes.

■ ***RF Heating in a Fusion Reactor.***

The robust and efficient application of high-power radio frequency (RF) systems is critical in the ion cyclotron, electron cyclotron, and lower hybrid ranges of frequencies, necessitating fidelity in modeling how RF waves interact with the tenuous edge plasma of a fusion device. Emerging exascale architectures would enable models that fully account for the multiscale and multiphysics nature of RF heating and current drive.

■ ***Whole-Device Fusion Modeling.***

In the hot fusion plasma in a toroidal geometry, several multiphysics processes are working together; most are scale inseparable and interacting nonlinearly in a self-organized manner. A whole-device modeling approach is needed to understand and predict the fusion reactor plasma. This effort cannot be realized without exascale (or beyond) computational capability at high fidelity.

2.3.2 Plasma Surface Interactions and Structural Materials

- Developing first-wall material and component solutions for future fusion devices (ITER, DEMO) is prohibitively costly, necessitating a multiscale modeling approach. An integrated and first-principles-based suite of advanced codes will be needed to model the boundary plasma and material surface.

2.3.3 Discovery Plasma Science

■ ***General Plasma Science.***

Modeling plasma turbulence remains beyond the capabilities of today's most advanced computers and algorithms. Exascale computing will enable direct numerical simulations of the high-dimensional, nonlinear turbulent dynamics.

■ ***High-Energy-Density Laboratory Plasmas.***

High-energy-density laboratory plasmas (HEDLPs) are extreme states of matter having pressures in excess of 1 Megabar. The physics of laser-plasma interactions and HEDLPs is multiscale, highly nonlinear, and must often be described by a kinetic modeling approach. Extreme HPC resources will enable increasing the problem size and grid resolution, running ensembles, and reducing turnover time.

■ ***Low-Temperature Plasmas.***

Low-temperature plasmas (LTPs) are partially ionized gases involved in the manufacture of electronics components and are typically in a strongly nonequilibrium state. Providing high-confidence models will transform the applied use of LTPs in industry.

2.3.4 Verification, Validation, and Uncertainty Quantification

- Confidence in "validated predictive models" must be earned through systematic confrontation with experimental data and a sharp focus on careful and quantitative estimates of errors and uncertainties. New methodologies and algorithms must address mathematical obstacles in multiphysics integration and use of computationally expensive codes in multiscale integration.

2.4 Priority Research Directions in High Energy Physics (HEP)

The U.S. high energy physics community has set forth a number of high-level opportunities for discovery over the coming years, including the following: to use the Higgs Boson as a new tool for discovery; pursue the physics associated with neutrino mass, identify the new physics of dark matter, understand cosmic acceleration (dark energy and inflation), and explore the unknown: new particles, interactions, and physical principles. In each of the above, HPC and data analysis will play crucial roles. In particular, very large datasets can be generated by the HPC applications — potentially much larger than those from experiments. Analyzing these datasets can prove to be as difficult a problem as running the original simulation. Thus, data-intensive tasks currently go hand in hand with many HEP HPC applications that were originally viewed as being only compute-intensive.

HEP divides its investigation into these topics into the energy, intensity, and cosmic frontiers and augmented by research into accelerator design and optimization to support research in the frontiers:

■ *Energy Frontier.*

The Energy Frontier is most visibly represented by two large HEP experiments, ATLAS and CMS (for the Compact Muon Solenoid detector), at the Large Hadron Collider (LHC), where the Higgs Boson was discovered in 2012. Both of those experiments continue to investigate the nature of the Higgs and search for additional physics. Research into Energy Frontier science makes heavy use of HPC for simulation to understand and verify the fundamental physics; use of HPC for detector simulations and analysis of large amounts of data is in an early growth phase.

■ *Intensity Frontier.*

The Intensity Frontier comprises the set of experiments that require intense particle beams and/or highly sensitive detectors to study rare processes with ever-greater precision and sensitivity. The study of neutrino physics represents much of the emphasis in the Intensity Frontier. Intensity Frontier experiments, while not as large as those at the LHC, are more numerous, and many aspects of neutrino physics are puzzling so the experimental picture is incomplete. While historically these experimental data streams have not been at the forefront of driving computational requirements, they require computationally demanding lattice quantum chromodynamics (QCD) calculations for the underlying theory.

■ *Cosmic Frontier.*

The primary science thrusts within the Cosmic Frontier are (1) understanding the nature of cosmic acceleration (investigating dark energy); (2) discovering the origin and physics of dark matter, the dominant matter component in the universe; and (3) investigating the nature of primordial fluctuations, which is also a test of the theory of inflation. A number of sky surveys in multiple wavebands are now scanning the sky to shed light on these questions. Near-future observations will generate extremely large datasets — hundreds of petabytes (PB) in size — and will be accompanied by large-scale cosmological simulation campaigns that will require significant computational resources.

■ *Accelerator Design and Optimization.*

An important field for all of HEP is accelerator design and optimization. Simulations are used to study electromagnetics and beam dynamics for current and near-future technology machines, as well as for designing future accelerators. The accelerator modeling community — which has been and will continue to be an important part of the field of high-energy physics — has demanding computational requirements.

2.5 Priority Research Directions in Nuclear Physics (NP)

Nuclear physics pursues a vital research program of exploring and quantifying the structure and dynamics of matter in our universe, especially the strongly-interacting matter composed of quarks and gluons described by the theory of quantum chromodynamics (QCD). From the fundamental properties of these building blocks of nature, through to cataclysmic astrophysical events, computational nuclear physics research establishes the essential link between the very smallest-scale and the very largest-scale events in our universe. Large-scale simulations of astrophysical objects, such as supernova and mergers, and studies of nuclei and their reactions address the central questions of how the chemical elements emerged, starting at the earliest moments of our universe. Large-scale numerical calculations of the structure and interactions of the neutron and proton from QCD, their behavior under extreme and non-equilibrium conditions, and how they respond to electroweak forces provide essential components of this research program. An extensive and vibrant experimental nuclear physics program, utilizing a range of accelerators and detectors, motivates and confirms our understanding of the nature of matter and its interactions.

Computational challenges in NP can be divided into five broad categories:

■ *Nuclear Astrophysics.*

Cutting-edge efforts in computational nuclear astrophysics are fundamentally multiphysics in nature, and large-scale numerical simulations are at the heart of this science. Key objectives in the Long-Range Plan (DOE and NSF 2015) include exploring the origin of the elements (e.g., ascertaining the site[s] where the heavy elements including uranium and gold are synthesized); the physics of ultra-dense, neutron-rich nuclear matter in neutron stars; and the nature of neutrinos and their interactions with nuclei and in dense nuclear matter.

■ *Experiment and Data.*

Nuclear physics experiments access a multidimensional and multichannel problem space, requiring beam intensity, polarization, and careful treatment of backgrounds and systematics. DOE NP experimental efforts at major facilities include nuclear studies at the Argonne Tandem Linac Accelerator System (ATLAS) at Argonne National Laboratory (Argonne); electron scattering from nuclear matter at Thomas Jefferson National Accelerator Facility (JLab); and proton and heavy-ion collisions at the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory (BNL), the Facility for Rare Isotope Beams (under construction), and the planned Electron Ion Collider (EIC). The ever-increasing data and analysis needs of NP experiments can be addressed, in part, through collaboration with ASCR.

■ *Nuclear Structure and Reactions.*

Nuclear structure and reaction research is progressing dramatically as a result of theoretical progress on the quantum many-body problem, computational advances, and simultaneous experimental developments. The challenges driving this field include determining the microphysics involved in creating all of the elements in the universe, examining general organizational principles and processes governing the structure and dynamics of nuclei, exploiting nuclei to reveal the fundamental symmetries of nature, and determining the scientific and practical uses of nuclei for society, including related fields of science. Nuclear many-body theory dramatically advances our ability to extract knowledge from experimental facilities, and then links the underlying theory of the strong interaction (QCD) to the study of fundamental symmetries, and to astrophysics and astrophysical environments.

■ *Cold Quantum Chromodynamics.*

The dynamics of quarks and gluons — the fundamental constituents of protons, neutrons, nuclei, dense matter, and exotic states of strongly interacting matter — are determined by the nonlinear and essentially quantum laws emerging from QCD. Precise knowledge and understanding of the quark-gluon structure of protons, neutrons, and nuclei; their excitations; the forces

between them; their response to external forces; and the nature of dense and exotic matter are critical to many aspects of subatomic science and provide essential inputs for additional NP Priority Research Directions. The prime areas of research in cold QCD are hadron structure and spectroscopy including the search for exotic states, nuclear interactions and structure arising from QCD, and applications of QCD to fundamental symmetries in nucleons and nuclei. Large-scale lattice QCD calculations are critical to addressing the scientific objectives in each of these areas.

■ ***Hot Quantum Chromodynamics.***

Moments after the big bang, the matter in the early universe was a very hot and dense plasma of quarks and gluons — the Quark Gluon Plasma (QGP). As the universe cooled, this plasma transitioned through a recently discovered liquid phase to a gas of protons, neutrons, and other particles, eventually forming the light elements in the Periodic Table. Research efforts in hot QCD are focused on developing a precise understanding of the structure and behavior of matter under extreme conditions of temperature and density, both in and out of equilibrium, and thereby complementing and supporting the extensive international experimental effort in heavy-ion collisions. Lattice QCD calculations and real-time Hamiltonian simulations play an essential role in this area.

2.6 Priority Research Directions for Advanced Scientific Computing Research (ASCR)

For ASCR researchers, computing facilities are not just a means to an end — the facilities themselves are the subject of the research: for instance, ASCR’s applied mathematicians are at the cutting edge of defining new algorithms to run on new machines.

For these reasons, the ASCR Exascale Requirements Review was scheduled to follow reviews by the other SC program offices so that ASCR researchers could consider the needs expressed in the previous reviews with a view toward (1) assessing the ability of the ASCR HPC facilities to support the needs of SC researchers in an exascale computing environment; and (2) identifying areas where that support could be enhanced, potentially opening new opportunity areas for computational research.

ASCR’s top research areas are as follows:

■ ***Computer Science.***

This program pursues innovative advances in programming languages and environments; system software; computer architectures; performance and productivity tools; and data management, analysis, and visualization, among many others areas. The primary focus is on effective use of very large-scale computers, networks, and data.

■ ***Applied Mathematics.***

This program supports basic research for delivering greater scientific computing and predictive capabilities through next-generation algorithm, modeling, data analysis, and software developments.

■ ***Computational Partnerships.***

This program supports research that will utilize or lead to partnerships with SC, the National Nuclear Security Administration (NNSA), or other DOE programs. This includes Scientific Discovery through Advanced Computing (SciDAC), which brings together experts from across application domains, including computer science and applied mathematics, and from universities and national laboratories to ensure that scientists are using state-of-the-art technologies to address their increasingly complex computational and data science challenges.

3 FINDINGS

The six Exascale Requirements Reviews presented big picture challenges and priority research directions. Findings across the reviews, both common and unique, were examined. During the Crosscut Meeting, participants identified common key challenges and needs, as well as any high-priority needs that may be unique to individual offices. The findings fell into four broad categories: computing, data, software and application development environment, and training and workforce development.

3.1 Computing

The requirements that are needed to support the science drivers from the six Exascale Requirements Reviews identify fast-growing computational needs and increasingly complex use cases. Compute requirements are driven by a broad range of scales and multiple capabilities. Findings across all reviews identify increasing needs for a rich computational ecosystem with improved access across all user facilities.

3.1.1 Hardware

A number of overarching observations and findings related to hardware emerged from the reviews:

- The Office of Science has substantial needs for both extreme-scale computing and computing at a somewhat smaller scale that still exceeds the capability of local resources outside the ASCR facilities.
 - See the complete reports for BER, BES, FES, HEP, and NP at exascaleage.org.
- The computational expertise possessed by a scientist does not always match the complexity of the hardware needed for his or her research [NP 3.4.4]³ [BES 3.5.1, 3.8.2].
- Scientists' computational tasks vary widely in computational complexity and have a growing range of workflow needs, both of which must be supported by exascale hardware.
- Usable and appropriate systems greatly enhance scientists' productivity; in contrast, hard-to-use systems and those poorly matched to computational needs can be a hindrance to scientific progress.
- Researchers need wide and early access to testbed and prototype systems to support development of codes and workflows prior to deployment on leading-edge production systems at scale.

The Office of Science research community articulated requirements for substantial computational resources at both extreme and moderate scales. These requirements are not always separate within a science campaign, but can be tightly coupled for a science campaign. A robust ecosystem capable of delivering computational and data capabilities will be essential to delivering on the science mission. While hardware solutions are potentially premature, the deployed systems need to maintain usability and appropriateness for research goals. Continued focus on balancing different compute capabilities and data capabilities is crucial for scientific success [HEP], [NP], [FES], [BES]. Here, coupling could include co-scheduling, easier access across compute resources, and workflow management, etc.

³ Material in brackets indicates specific sections in the six Exascale Requirements Review reports from which findings have been drawn for this Crosscut Report. For ease of comparison, Appendices G through L contain (or reprint) the specific sections cited, as follows: Appendix G contains sections cited from BES; Appendix H contains sections cited from BER; Appendix I contains sections cited from FES; Appendix J contains sections cited from HEP; Appendix K contains sections cited from NP; and Appendix L contains sections cited from ASCR. In addition, the complete reports can be accessed at exascaleage.org.

Complexity is increasing in most aspects of computing for the Office of Science. Applications are becoming more complex in scale, algorithms, and physics. Adequate computational resources are crucial to delivering the science vision described in the six reports. As an example, BES has simulations that currently require weeks to run on existing HPC systems, but will need to complete in minutes to meet decision-making requirements and provide useful results [BES 3.6.2.4.2]. BES also has vast ensembles of simulations of all sizes that are needed to perform validation, uncertainty quantification, and parameter sweeps that will enable scientists to draw robust, high-fidelity scientific conclusions [BES ES.1]. This need is common among the program offices.

Scientific campaigns are increasing in complexity, coupling multiple resources at multiple scales at local sites, on community resources, at experimental facilities, and at ASCR HPC centers. Given the variety of computational environments and the generational hardware changes that are coming with increasing frequency, scientists need a consistent, reliable way to support effective development of scientific applications and workflows.

3.1.2 Allocations, Access, and Policies

Scientists need policies at HPC centers that accommodate their requirements. The program offices had substantial needs in this category, including the following:

- A substantial increase in hours allocated to each project over what is available today. For example, computational needs in terms of today's core-hours are expected to increase by a factor of 1,000 by 2025 in some of the High Energy Physics frontiers [HEP 4.3].
- Diverse scheduling. Real-time, pseudo-real time, co-scheduling, variable job requirements, allocations based on other resources like disk and memory, etc., are needed [HEP 4.1.2] [FES ES.4, 4.4] [NP 3.1.4] [BES 3.6.2.1.3, 3.6.2.2.3] [BER 3.1].
- Improved engagement with scientists about allocation programs and policies, including the possibility of long-term commitments to support extended science campaigns and missions [HEP 1.2] [FES 3.1.2.4.3; page C-9, Ernst] [BES 4.2].
- Access to HPC systems' operational data for research purposes [ASCR 4.1.2].
- Seamless access across facilities, especially to support data-intensive workflows [BER 3.2.3.4].
- HPC facility policies that support easier testing and deployment of new software technologies developed by ASCR researchers.

Researchers at all discipline reviews noted that their use of computing resources is growing into nontraditional domains, in addition to increasing demands from traditional domains. With the breadth of scientific challenges, scheduling needs are not all the same. Some projects have long-running, large-scale jobs, whereas others have smaller, very-long-running ensemble runs or process large amounts of experimental data. All six offices described a need for faster access to compute resources, whether for development, debugging, analysis, or real-time computing [BER 3.2.1.1.4] [BES ES.1, 3.1.4] [ASCR 3.2.1.1, 3.4.2.1] [FES ES.4] [NP ES.1] [HEP 4.1.2]. Whether a need is truly real-time or not, the need exists to couple multiple resources — compute, data, and experiments — to accomplish science, with significant implications for DOE ASCR's networking facility, ESnet. Even in more traditional use models, researchers expressed concern that substantial queue wait times can significantly impede scientific progress.

As the use of computing resources is changing, participants at several reviews identified a need to evolve allocation programs and usage policies. Allocation programs like the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) at the Leadership Computing Facilities (LCFs) and the Energy Research Computing Allocations Process (ERCAP) at the National Energy Research Scientific Computing Center (NERSC) are the mechanisms to gain access to the

ASCR compute facilities. Policies for using the computing facilities are driven by each facility's program mission, allocation programs, and metrics through which facility performance is evaluated. For example, users at the HEP review expressed concern about their ability to conduct long-range planning when they must re-compete for compute resources on a regular basis. In addition, researchers expressed concern that performance criteria and programmatic value are not playing an appropriate role in the review process.

Scheduling and the difficulty of access across ASCR HPC facilities was a common concern across all six reviews. As workflows and use-models of facility resources become more complex, it is important that policies be in place to support them. Security policies at facilities need to be sufficiently open to accommodate automated workflows operating between multiple sites, as well as sufficiently secure to accomplish DOE's translational research goals. Different facilities have different access mechanisms and scheduling software, adding to developer cost and requiring additional implementation in workflow automation tools [BES 3.5.4]. Generally, the review participants agreed that finding a mechanism to improve access to the facilities will be important to their scientific progress.

3.2 Data

Opportunities for scientific discovery through the analysis of extremely large datasets are growing quickly in fields of research supported by the Office of Science. To realize these opportunities, science teams are looking to facilities and expertise within ASCR to address challenges of scale and complexity that cannot be addressed by individual science teams or individual DOE program offices. For example, the executive summary of the BES report highlights the fact that improvements and new capabilities at BES facilities are creating "offshoot" challenges — specifically, unprecedented growth in data volumes, complexity, and access; the need for curation; and integration of diverse datasets — that the community must address to realize science objectives made possible by enhanced facilities. Drivers for BES come from the Advanced Photon Source Upgrade (APS-U) at Argonne and the Linac Coherent Light Source (LCLS-II) at the SLAC National Accelerator Laboratory, among others. In HEP, the CMS and ATLAS experiments at the LHC are turning to HPC to help solve data analysis challenges associated with LHC upgrades. BER has extreme data and analysis needs in genomics and Earth system research. These are just a few of many examples where HPC-class resources are required to realize the scientific benefits of DOE investments in next-generation facilities and experiments.

New data challenges are driven not only by experimental facilities; the requirements reviews also emphasized that growth in data created on HPC systems by simulations is becoming increasingly difficult for individual principle investigators to store, analyze, and manage [HEP ES] [FES 2.2] [BER 3.1.4]. This finding highlights the similar challenges shared by both the experimental science and modeling/simulation communities.

Across all of the exascale requirements reviews with the six Office of Science program offices, a number of common themes emerged as being particularly challenging areas:

- Large-scale, long-term data storage, analysis, and community reuse;
- Experimental and simulation workflows;
- Data management, archiving, and curation;
- Remote access, sharing, and data transfer; and
- Input/output (I/O) performance from simulation and analysis programs.

Addressing these challenges will be eased by having a close integration of HPC simulation and data analysis. Being able to perform simulations and analysis together in one locale will minimize data movement and facilitate interdependent workflows [HEP ES] [BER 2.2.1] [ASCR 3.2.1.1].

3.2.1 Large-Scale Data Storage and Advanced Analytics

Simulation, experiment, and observation are producing data at unprecedented rates and quantities, to the point that supercomputers are needed to analyze the data and draw meaningful scientific conclusions. Simulated datasets produced in cosmology, Earth system science, and fusion research are large and complex and require post-processing to find important effects like extreme weather events, the formation of structure in the early universe, and key plasma instabilities that affect fusion reactor performance [BER 3.2.1.2] [FES 3.1.4] [HEP 3.1]. On the experimental/observational side, the volume of data produced is rising rapidly [HEP 4.1.2] [BER 3.1.4.4.1] [BES 3.6.2.3.1]. For example, the current LHC output is on the order of 10 PB/year and is expected to rise to more than 150 PB/year in less than a decade, with a need to store exabytes' worth of data permanently [HEP 4.1.2]. Projects like CMB-S4 (cosmic microwave background) and the Large Synoptic Survey Telescope (LSST) (optical) will gather tens to hundreds of PBs' worth of data [HEP C.7, Borgland et al.].

The methods and software used to perform data analysis are significantly different than those used by “traditional” HPC simulation. The reviews revealed a need for new algorithms and approaches for analyzing data, including advanced statistics, machine learning, in-situ processing, image segmentation, anomaly detection, data filtering, and neural networks.

The activities of analyzing and drawing inferences from scientific data are being revolutionized by advances in machine learning and analysis frameworks that rely on software technologies like python, MATLAB, R, and Julia [FES pages C-24, 25; Jenkins et al.] [BER 3.2.7.2; page C-159, Urban; page D-21, Maranas]. At the same time, long-established analysis packages — particularly those used by LHC experiments — have strict dependencies on software and operating environments not typically found on HPC systems [HEP C.8.2, Viren and Schram]. Support for large-scale data analysis at HPC centers will require the centers to support nontraditional workflows, software, and access methods. It will require a close collaboration between science teams and facilities to enable the capabilities required for analysis of this “big data.” For example, the following were noted in the requirements reviews:

- Improved tools are needed for data analysis and visualization including machine learning and deep learning techniques for pattern recognition and data quality assessment [BES ES.4, 4.3].
- The overhead associated with tuning a complex particle accelerator can be significantly reduced, and the performance of the accelerator can be significantly improved, by employing machine learning. The community would benefit greatly from the development of better algorithms (such as machine learning methods) and data-processing tools for lossless real-time data reduction near the beamline [NP 3.2.1, C.2.4].
- The growth rate of data produced by simulations is overwhelming the current ability of both facilities and researchers to store and analyze it. Additional resources and new techniques for data analysis are urgently needed [HEP ES].
- New approaches are needed to make more effective scientific use of the enormous volumes of data generated within BER's biological science programs [BER 2.1].

3.2.2 Experimental and Simulation Workflows

Workflows that manage everything from data acquisition to transfer, analysis, archiving, and sharing are a key need in both experimental analysis and simulation. Several workflow use cases, in addition to those associated with analysis of large data mentioned in the previous section, were identified in the reviews.

3.2.2.1 Steering of Experiments and Co-scheduling

As data volumes and acquisition rates grow at DOE experimental facilities like light sources and particle accelerators, analysis and storage needs will outstrip — or have already — the capabilities of local resources [BES ES.2]. Efficient and effective use of these facilities often requires that scientists receive immediate feedback to steer subsequent measurements to be taken during the same session [FES 3.4.1.1] [BES ES 3.6]. There is a widespread need for connecting HPC facilities to experimental facilities for data analysis during experiments [BES ES.1, ES.3.6] [BER 3.1.4.2.3] [FES 3.4.1.1] [NP ES.3.2]. Data streamed from the experiment to an HPC facility, analyzed there, and returned in near real time will provide invaluable information for guiding the next measurement. For this workflow to be successful requires co-scheduling of experimental runs and HPC resources and a cross-site workflow that seamlessly transfers the data, performs an analysis, and returns the results during the allotted time at the experiment [BES ES.3.6].

3.2.2.2 Direct Streaming of Large Datasets from Experiments

Because time at light sources and accelerators is tightly scheduled — with scientists often having only a short time at a facility — scientists are looking to save as much data as possible from each run, and reduce or eliminate many of the “triggers” that are currently used to filter data near the detector [NP 3.2.1]. Combined with the explosive growth in the amount of data produced with each run, this set of pressures will in the future completely overwhelm the ability of local resources to collect, save, and analyze these data [BES 3.6.1]. HPC facilities offer the potential ability to store large repositories of experimental data that can be streamed and analyzed immediately, stored for later offline analysis, leveraged for planning future experiments, or served to the domain community to amplify the impact of the work.

3.2.2.3 Integration of Simulation and Experimental Data for Analysis

A close integration of HPC simulation and data analysis is often needed to interpret the results of experimental and observational data. Software interfaces to access and compute on both types of data are needed [BER 3.2.3.3]. The sheer amount of data poses severe challenges for both data storage and network bandwidth. Analysis needs to operate on data produced by both simulation and experiment, and having them co-located greatly facilitates performance and efficiency. Such an integration minimizes data movement and facilitates interdependent workflows [HEP ES] [FES 3.1.4]. In cases where this co-location is not possible, scientists need efficient and easy-to-use procedures to transfer the data [HEP C.2, Almgren et al.].

3.2.2.4 Facilitating Analysis Workflows across Multiple, Distributed Facilities

Researchers often work at multiple HPC facilities and wish to analyze data from multiple facilities to accomplish science goals [BES 3.6; 3.7.2.2]. Common tools, practices, interfaces, and environments across the facilities would greatly enhance productivity [NP 3.1.4; 3.2.3.2, item 3] [BES 3.7.2.2].

Some data analysis workflows require concurrent access to data from multiple sources that may be geographically distributed. For example, HEP Energy Frontier applications require real-time remote access to resources such as remote databases while jobs are running [HEP 4.1.2].

3.2.3 Data Management, Archiving, and Curation

Long-term archiving is needed for big datasets that cannot be easily recreated or reacquired, for example, LHC data that are on the order of hundreds of petabytes [HEP 4.1.2]. Nuclear astrophysics researchers will need persistent storage of 100-PB datasets per simulation by 2025 [NP 3.1.4]. In many cases, these data have value for decades.

Because of the complexity of data, new methods for managing it and its associated metadata need to be developed [NP 3.2.3.2, item 3] [BES 3.6.2.3] [FES 3.4.2.2]. Scientists must be able to find what is of high importance to address immediate scientific tasks and inquire about its provenance.

To enable these features, the development of community standards and federated databases is needed [FES 4.3] [BES 3.3.4]. Automated provenance capture and improved dataset browsing capabilities are a priority research area for verification and validation and uncertainty quantification [FES 3.4.2.4].

Further collaborations will be needed to develop new methods for managing, storing, and analyzing data, particularly owing to its heterogeneous and large-scale nature [BER 3.2.5.4.5].

3.2.4 I/O Performance

For many data analysis applications, I/O performance is a limiting factor [BER 3.1.3.4]. For simulation, I/O performance controls the amount of data that can be saved for future analysis and limits the number of checkpoint files that can be written to guard against system component failure, which will become more crucial as HPC systems grow in size and complexity [BER 3.2.4.4] [HEP D.3.2, Finkel et al.]. Simulation scientists use a “rule of thumb” that I/O should not consume more than 20% of the compute time [NP 3.4.4] [HEP D.1.3, Ge et al.], which implies system I/O rates of terabytes (TBs) per second [FES 3.1.1.4]. I/O challenges on current platforms already limit what can be saved for post-run analysis, leaving valuable physics on the table [FES 2.2]. In addition, I/O performance has a significant impact on data-intensive analysis applications, and improvements in I/O performance would be of significant benefit to data-intensive applications [HEP 4.1.2; C.6, Wurthwein et al.] [BER 2.3.2; 3.1.1.4.2] [BES page C-36, Dixon; page C-94, Hexemer and Parkinson].

The ability of storage systems to keep up with the performance of exascale computing systems was identified as a concern in multiple reports. The growing divergence between computing capability and I/O bandwidth in large-scale systems could have significant impact on overall application performance [BER 3.2.1.1.4] [FES 3.1.2.4.1, 3.4.2.2] [HEP D.4.3, Gnedin].

3.2.5 Remote Access, Sharing, and Data Transfer

Increases in demand for computing by data-intensive science programs will increase the demand for networking and data transfer in and out of HPC facilities [HEP 1.2]. Data transfer and storage will become significant issues for some programs in the future, and indeed are already pain points today [HEP 4.1.1]. High-performance, reliable, and predictable data transfer across multiple storage endpoints is a requirement for the future [HEP 4.1.2] [BER 3.2.6.1.1].

Data sharing is important for multiple science communities, with challenges to be overcome as datasets scale up [BER 2.1.1, 3.2.1.1.4, 4.3] [BES 3.6.2.3] [FES 4]. Data access needs to be controlled, such that data can be private to a project, shared with the world, and/or following policies in between (a la Google Docs sharing).

Data transfer between experimental facilities and HPC facilities is increasingly important. Access to HPC facilities by scientists at experimental facilities requires federated authentication, as well as high-performance data transfer and management capabilities [BES 4.3].

In addition to performing analysis on data, scientists need to archive these datasets and make them available to the community for additional analysis. Like data from experiment and observation, simulated data are also becoming a valuable resource to the entire community [HEP 4.1.1] that, as such, need to be archived, analyzed, and made accessible to broader communities.

3.3 Software and Application Development Environment

In their entirety, the six exascale reports are a call for DOE to advance a common, sustainable, and productive computational science and engineering (CSE) software ecosystem. Moreover, the requirement for a cohesive, cross-ASCR-facility software ecosystem is identified in all previous reports [HEP 4.1, 4.2] [BES 4.2] [NP 4.1] [FES 4.2] [BER 4.2] [ASCR 4.1.2]. DOE’s strategy of advancing diverse computing architectures to more fully meet the needs of a broad and diverse mission workload and better manage programmatic risk — the risk associated with deploying unique, “serial-number one” class systems — is only fully realized in the context of performance-portable application codes; system software; and a programming environment based on open, broadly supported, and sensible standards.

The challenge of developing scientific applications and translating these into portable, sustainable software for high-performance computers is not a new one. However, it is a challenge that is growing more difficult as the complexity and diversity of HPC grows. In addition, there are trends articulated in the six requirements reports that add to the urgency of realizing sustainable software development, such as these:

- An accelerating growth of nontraditional HPC users and use cases from experimental, observational, and otherwise data-intensive applications [HEP 3.2, 4.1.2] [ASCR 3.2.1.1] [NP 3.2] [BES 3.6] [BER 3.1.4, 3.2.6] [FES 3.4].
- Diverse responses from various academic departments to the challenge of defining training in computing and computational methods either within or outside core disciplinary education (Section 3.4). This diversity results in a wide range of capabilities for researchers executing the Office of Science mission.
- Various levels of support for the development of scientific software (a.k.a., “intangible scientific instrumentation”) among relevant research sponsors.

Widely used applications need organized support for portability and sustainability. Investments become more pressing because of the increasing complexity of computer architecture. Cutting across the six reports, a grand vision for scientific software development in the future emerges of an ecosystem that (1) allows efficient programming of tomorrow’s machines that is as straightforward as programming the personal computers and laptops of today, and (2) enables the substantial investment in legacy code to be transformed, expressing all available parallelism, and thereby sustained into future missions.

3.3.1 Models, Methods, and Algorithms

Mathematical, software, and algorithm development is required in all Office of Science programs to take advantage of exascale computing architectures and to meet data analysis, management, and workflow needs [BES ES Abstract, bullet 2]. In the requirements reviews, researchers identified a need for sustainable and performant software tools and applications to enable effective use of exascale systems. New algorithms are needed to enable codes to run efficiently on upcoming HPC architectures so scientists will be able to model larger systems with greater fidelity and predictive power (see, e.g., [BES 3.5] [FES 3.1.1.4] [NP 3.3, 3.4] [BER 3.2]). In specific disciplines, such as in Earth system, biology, and material science, some applications are used by many scientists who may or may not be developers of the applications. In addition, in the biology community, new models are continuously being developed in response to the rapid evolution of scientific drivers.

New mathematics are needed to enable order-of-magnitude improvements in speed and accuracy in predictions [BES 3.7] [NP 3.1] [BER 3.1]. Revolutionary software and algorithms are needed to take advantage of exascale capabilities and meet data analysis and management needs. Scalable data processing, data analysis, machine learning, discrete algorithms, and multiscale/multiphysics

simulations are crucial for reducing and understanding the large-scale data that will be produced by exascale systems.

Improved mathematical, algorithmic, and workflow methods for verification, validation, and uncertainty quantification (UQ) are needed for a wide array of Office of Science missions. A systematic validation of models and approximations is a fundamental aspect for BES in the development of reliable computational discovery tools [BES 3.3.3.1]. New methods are needed to greatly accelerate the verification, validation, and UQ process. UQ resulting from model parameters will require large numbers of simulations to cover high-dimensional parameter spaces. Participants in the FES workshop named application codes and verification and validation techniques, as well as models and algorithms, as key factors requiring significant methods development investments [FES 3.4]. The key findings from the BER workshop included that the fusion of model simulations and observational data must take place for better model initialization, uncertainty analysis, validation, and tuning [BER ES.1, Key Findings; 3.2.5; 3.2.6]. The need for UQ and rigorous model-observation comparison manifests in many aspects of Earth system modeling. On a high level, this need for tight coupling of theory, observation, and simulation is driven by the need to quantitatively assess the predictive fidelity of ESMs, as well as the use of scenario analysis to inform decision makers in integrated assessment modeling. In this context, the term “UQ” is very broadly defined and incorporates activities ranging from observational data characterization and calibration to sensitivity analysis, surrogate construction, forward propagation, and attribution. The NP workshop participants indicate that collaborations with statisticians are forming, and these collaborations are expected to become essential considering the need for uncertainty quantification in theoretical computations and the dramatic increases in data volume and complexity that are expected in the exascale era [NP 4.1]. Participants in the ASCR workshop called out the need for support for automated software testing on all systems at all computing facilities, including unit, regression, verification, integration, and performance testing [ASCR 3.4.2.1].

3.3.2 Software Portability, Performance, and Sustainability

All SC program communities urgently require sufficiently performance-portable and expressive programming models [HEP 3.3] [FES 3.1.3] [ASCR 2.1] [BES ES.3.7] [NP ES.1] [BER ES.4]. Investment in software for new architectures is becoming even more pressing because of the complexity of the new architectures. Scientists identified a need for a broader environment to develop applications in collaborative efforts for new architectures. Some communities (e.g., traditional HEP users) recommend broad support for a limited set of basic primitive operations [HEP 4.1.1]. Others, including nontraditional, data-intensive user communities, will require greater integration of HPC with grid-based or cloud-based technology [HEP 4.1.2]. Many reports strongly articulated requirements focusing on improving the ease-of-use of future system architectures [BER 3.2.4]. These researchers prefer to avoid spending years’ worth of time rewriting their software to utilize tomorrow’s computing platforms, thereby limiting time devoted to advancing the progress of their scientific endeavors [BES 3.7.2.3]. However, they recognize that some code rewriting is inevitable because different algorithms are optimal on different hardware platforms [BER 3.1.1.4].

Even with an assumption that a single new language or programming model (PGAS, C++ meta programming, etc.) provides the best path to portability, scientists identified a need for methods to confidently rewrite large legacy code models in that language or model. A thorough suite of unit tests is a practical necessity for systematically rewriting a large legacy code (e.g., one unit at a time) and provides a way to verify the rewritten portion. Finding efficient, semiautomatic ways to retrofit unit tests to a legacy code is a needed area of research [BER 3.2.6].

The requirements reviews highlighted the need for improved developer productivity and the ability to package the artifacts of DOE research for broad reuse [ASCR 3.4.2]. The challenges to be addressed to meet this goal include intertwined technical, policy, and social issues. For example, there is a recognized need for community planning, policies, and processes to improve software quality, interoperability, testing, and deployment within the DOE exascale ecosystem. The strategy must be flexible to embrace diverse business models for software sustainability, as, for example, some mission-critical research activities are enabled by proprietary and/or commercial vendor software. Other software is not widely distributed but remains the property of individual research groups or collaborations. In each case, there will be the need for funding and investment to sustain and grow key DOE research software products. In every case, there is the need for increased collaboration among DOE science communities, ASCR Research, ASCR Facilities, and the Exascale Computing Project (ECP).

Requirements reviews attendees highlighted the need for the following:

- Programming models to reduce the complexity of scientific programming [BES 3.5.4] [BER 3.2.7.2].
- Investments maintained in legacy applications [HEP 3.3] [FES 3.1.3].
- The ability to more easily port applications to new architectures without requiring heroic efforts [BES 3.7.2.3] [ASCR 3.6.1] [BER 3.2.4.4].

The idea of “sustainability tension” captures the following:

- The software engineering effort should be focused on extending capabilities of science codes rather than on the interface between science codes and HPC machines.
- Millions of lines of existing code were written for previous architectures and represent a valuable scientific capability that will be needed for the exascale ecosystem.
- There is a need to maintain and advance this software instrumentation and these codes through transformation to express all available parallelism on exascale architectures in a manner that can still be understood by domain scientists [HEP 2.2] [BER ES.3.1.1].

3.3.3 Common Environment/Sensible Standards

Broad consensus exists among the reports that the sustainability of DOE software infrastructure and the productivity of DOE researchers can be advanced by establishing a common computing and data analysis environment across platforms and facilities and through the promotion of sensible, community-based standards.

A common environment across diverse compute, data, experimental, and observational facilities will enable the advancement of mission goals. Many of these requirements were repeated throughout the requirement reviews, such as the need to:

- Address both large-scale simulations and data analysis on flagship and leadership facilities, as well as development activities on testbeds and other platforms.
- Coordinate persistent support for policies across facilities: for example, sustained access, authentication, queue policies, data federation, and transport of large data.
- Advance goals of portable performance, testing, continuous integration, and software quality through common tools, environments, and modules.
- Standardize management of HPC codes on DOE compute and data systems.
- Implement common documentation and system defaults for libraries, job launchers, job schedulers, etc.

Standard mathematical libraries and frameworks must transition to new architectures in a predictable and timely fashion, including via:

- Deeper involvement and sharing in open-source activities: compilers, libraries, committees, etc.
- Improved standards (languages, tools, and programming support) to facilitate implementation in exascale ecosystems.
- Improved programming models to reduce complexity in scientific programming.

Urgency is being expressed through these reports to define which common tools and libraries will be supported and sustained, especially tools and libraries for which DOE has been the sponsor of the software development. The concern that today's important libraries could disappear in the absence of sustained support is inhibiting their broader adoption among DOE mission scientists. There needs to be a model in which sustained maintenance of tools and libraries is assured.

3.4 Training and Workforce Development

The need for significant investments in workforce development and training was a common theme in all six requirements workshops and generated significant discussion in the Crosscut workshop as well. In the executive summaries of four of the six reports (BER, BES, HEP, and NP), workforce development or training were key findings. For FES, there was significant discussion in several sections, and the issue was mentioned many times in the white papers and case studies. For ASCR, training was a major area of interest in the survey that led up to the workshop and was referenced in several sections of the report.

The following examples from the FES, BER, and HEP reports shed light on the critical issues: an “insufficient application workforce would impede progress” [FES page C-21, Umansky], “few [institutions] train interdisciplinary computationally sophisticated scientists” [BER page C-63, Brown], “our workforce as it stands does not possess the skill set to tackle [refactorization of codes]” [HEP 4.1.2].

3.4.1 Workforce Development

The BES review found that across BES, workforce development is a significant bottleneck:

“Providing the kinds of education and training deliverables we need to prepare these future computational and domain scientists and computational software developers for exascale computing presents a major challenge” [BES 3.8.1].

This section goes on to describe the challenges with next-generation architectures and the need for cross-disciplinary education.

Directly from the first page of the NP report in the key findings, the community states that workforce development should include “enhanced collaboration between NP, ASCR, and the National Science Foundation (NSF)” and “new positions at the ASCR/NP interface” [NP ES]. NP mentions that “particularly critical for the long-term success of this program (NP) is the growth in the number of permanent positions for scientists working at the interface of applied math and computer science — with NP, both at national laboratories and universities” [NP ES.1]. The report goes on to say that long-term career paths are needed and that curating these career paths is a place of collaboration between ASCR and NP. NP specifically cited SciDAC and the ASCR facilities as key elements of workforce development [NP 4.1].

The Biological and Environmental Research report states in its key findings that “creation of the necessary system components requires a workforce trained deeply not only in the core computational, data scientific, mathematical, and natural scientific disciplines that underlie the above technologies but in how to co-design and develop tools that support open-community development and research” [BER ES.1].

3.4.2 Training

For HEP, training would be critical for developing better workflows and transitioning codes to next-generation platforms. In tandem, a workforce would be grown that is capable of developing simulations and providing new analysis, thereby leveraging the facility platforms. This development, along with the other findings, would be critical in moving the community closer to realizing its scientific goals [HEP ES, 4.1.2].

FES was concerned about the ability to train the next generation of students coming into their workforce pipelines. FES also explored the problem of connecting these students with standardized programming models and environments across the facilities [FES 3.1.3.4.5]. Facilities would need not only to support the software but also be able to effect the use of best practices and act as facilitators for new approaches for the community [FES 3.3.2.4].

The ASCR report cited a need for consolidated information and greater depth of training offered at facilities [ASCR ES.3]. There is specific interest in how the facilities can better train the researchers and developers in ASCR Research to use emerging systems. Finally, ASCR Research saw the facilities as a platform for software training and coordinated training management. For instance, there is a need for training programs on continuous integration and library development [ASCR 3.6.1].

A need for training and development in data analysis and machine learning was noted in the BER report, where it was mentioned that student internships and summer workshops are good mechanisms to promote workforce development [BER 3.2.6.2.3].

3.4.3 Existing Programs

BES named existing programs such as Computational Science Graduate Fellowship (CSGF), Argonne Training Program for Extreme-Scale Computing (ATPESC), and SciDAC as current exemplars on how to provide this development and noted expansion of these programs as a priority. BES also notes the CECAM (Centre Européen de Calcul Atomique et Moléculaire) workshop model as another way to approach training and education [BES 3.8.3.1]. FES, like BES, mentions that successful programs such as SciDAC are critical for developing these teams [FES 4.1]. NP specifically cited SciDAC and the ASCR facility training programs as key parts of workforce development [NP 4.1].

This page is intentionally left blank.

4 PATH FORWARD

In addition to requirements and findings presented in the previous section of this report and in the six previous reports, attendees identified a number of areas in which deeper collaborations and relationships among DOE SC offices will be needed to advance their research and to effectively compute in an exascale environment. In each of the reports, attendees asked for deeper relationships with ASCR Facilities and ASCR Research to address their needs and challenges [HEP 4.2] [BER 4] [ASCR 4, 4.1.2] [NP 4.1] [BES 4] [BER 4.2].

A clear theme was a call for increased communication among all stakeholders and the dissolution of any organizational barriers or compartmentalization that might hinder building effective collaborations. Another ubiquitous desire was to enhance efforts to educate and train a sustainable workforce stationed at the interface of traditional application scientists and HPC facility experts and ASCR researchers.

4.1 Areas for Collaboration

The following themes emerged as areas in which strong collaborations are needed and were believed to be potentially the most productive in establishing an exascale ecosystem.

4.1.1 Co-evolving ASCR Facilities and DOE Experimental Science

As the data-related needs of experimental science at DOE user facilities and other DOE-supported missions/projects continue to grow quickly, a co-evolution of capabilities, services, and practices between experimental scientists and ASCR Facilities will be required. Collaborations are needed to:

- Ensure that appropriate ASCR facility resources are available for the lifetime of experiments [HEP 4.2].
- Ensure that local computing at user facilities is appropriate and can integrate with ASCR facilities as needed [NP 4.1].
- Enable, where required, real-time computing, flexible queuing structures, support for complex workflows, and the development and implementation of other policies to support data science and policies that may be different from those in place at HPC facilities [BES 4] [FES 4.2], while maintaining appropriate support for traditional HPC use cases.
- Develop the tools and operational models to support the coupling together of experimental and computational facilities [BES ES.1, ES.3.6] [BER 3.1.4.2.3] [FES 3.4.1.1] [NP ES.3.2].
- Identify, configure, and support testbed systems that might be needed to assure the success of the ASCR Facilities' mission.

Scientists working on a number of experiments stated that planning to ensure resource availability is essential because experimental projects and missions have extended lifetimes of many years [HEP 4.2].

4.1.2 Addressing Data-related Needs and Developing New Methods of Analysis and Data Management

Analysis and management of data both from simulation and experiment are needed [HEP 4.3] [NP 4, 4.1] [BES 4.3], and there is an opportunity for application teams and ASCR to collaborate on creating new, scalable algorithms; build better memory management techniques; explore data-related methods; and develop end-to-end analysis pipelines to solve workflow, visualization, and optimization concerns.

4.1.3 Pursuing Formal, Close Collaborations between Application Scientists and ASCR Experts

The advancing complexity of computer hardware presents performance and portability challenges to application programmers, but also opens new opportunities for addressing scientific challenges that have heretofore been impossible to solve. Success in using cutting-edge ASCR facilities requires close collaborations among application scientists, applied mathematicians, computer scientists, and experts in computer architectures to sustain and build new scalable, performant algorithms and applications [BES 4.1]. Many existing codes are not ready to run efficiently on next-generation machines, and future systems may be more challenging. Old algorithms and methods, developed for traditional architectures, may not be those that are most appropriate for modern energy-efficient processors that are extremely parallel. Participants named application codes as well as models and algorithms as key factors requiring significant methods development activity [NP 4.1] [BES 4.1] [BER 4.1] [FES 4.1]. The SciDAC program can serve as a model for such collaborations that will focus specifically on creating the exascale ecosystem [BER 2.3.2].

The reports also expressed urgency about the need to define which common tools and libraries will be supported and sustained so as to transition to new architectures in a predictable and timely fashion.

4.1.4 Creating an Easy-to-use HPC Environment That Is Integrated across ASCR Centers and Other Work Environments

Science teams use a diverse set of systems and facilities, and they would like their computing environments to be as uniform as possible to ease development efforts and promote portability of applications and workflows [FES 4.2]. Researchers would like to work with ASCR to develop an exascale HPC environment that optimizes their productivity, including practices that enable ease of use, application portability, and workflow integration across environments [BES 3.7.2.3]. In addition to working at ASCR HPC facilities, researchers perform a variety of computations and tasks at other user facilities and on local resources. They would like to move among these environments easily without undue effort.

Development of common identification and authorization mechanisms would greatly help integrate computing environments [BES 4.3] across scientists' computing landscapes.

4.1.5 Sharing the Scale-of-Data “Burden” among Facilities

The scale of data generated in current scientific research creates challenges in performance, locality, analytics, access, and curation.

4.1.6 Developing Closer Collaborations among ASCR Research and Facilities to Support DOE Science

ASCR can play a major role in preparing application teams and their codes for exascale architectures. In the ASCR report, meeting participants identified a number of potential areas of collaboration between ASCR Facilities and ASCR Research that would assist in this effort. Areas in which ASCR Research and Facilities can work together to develop additional capabilities include the following:

- Instrumentation and monitoring,
- Flexible software deployment processes,
- System reconfiguration agility,
- Facilitation of hardware and software configuration changes,

- Deployment of next-generation network technologies,
- Development of data movement tools,
- Enablement of on-ramp capabilities for new system software,
- Strategies for transitioning research project successes into deployed software, and
- Data repository and archiving.

4.1.7 Recruiting and Providing Training to DOE Researchers Who Are Late Adopters

During the Crosscut breakout session, attendees recognized that there were both early adopters of HPC and others who adopted technologies much later. This situation creates opportunity gaps among current practitioners and reduces the possibility of successfully realizing cross collaboration between offices. There was general agreement that DOE should take advantage of the lessons learned and software developed by the early adopters and use that to accelerate the adoption of these HPC facilities and tools by DOE offices that are later adopters. In addition, it was found that connecting program managers in those offices to ASCR program managers would help streamline these initiatives.

4.1.8 Supporting Training and Workforce Development

The reports identified a current obstacle to collaboration and integration of application science teams, facilities experts, and computer scientists and mathematicians: a lack of organizational structure to support these collaborations [HEP 4.2] [FES 4.4]. A possibility expressed at the reviews is to enhance or redirect partnerships where some entity — like joints institutes with the program offices and ASCR — would provide multifaceted partnerships with ASCR Research and Facilities divisions [HEP 4.2] [NP 4.1].

Collaborations such as these are not entirely new, and review participants identified SciDAC and the application readiness programs at the facilities — the NERSC Exascale Science Applications Program (NESAP), the Center for Accelerated Application Readiness (CAAR) at the Oak Ridge Leadership Computing Facility (OLCF), and the Early Science Program (ESP) at the Argonne Leadership Computing Facility (ALCF) — as successful programs to copy or extend [NP 4.1]. Facility training programs, which train new project members and disseminate lessons learned, were also listed as valuable, and expanding them was believed to be valuable [NP 4.1] [BES 3.8.2]. The Argonne Training School on Extreme-Scale Computing was cited as an excellent program. SciDAC Application Partnerships or a similar program that brings together application scientists and experts in applied mathematics and computer science were listed as crucial [NP 4.1] [BES 3.8.2].

Across DOE offices, participants stated that a closer integration and collaboration among DOE experimental facilities and ASCR HPC centers was desired, as in this example: “ASCR and BES facilities must create more data-centric environments with highly effective data analytics tools for their users” [BES 4.3]. However, few specifics on how this integration would be implemented were proposed; it was not clear whether a path forward would involve joint facility-to-facility efforts or efforts from ASCR Facilities to individual science teams or from ASCR Facilities to individual experiments (e.g., beamlines) — or whether all of the above would be most productive.

The ASCR Computational Science Graduate Fellowship (CSGF) program is an example of a highly successful program that integrates domain scientists with experts in computation. The postdoctoral programs associated with the NESAP, CAAR, and ESP programs at NERSC, OLCF, and ALCF, respectively, are further examples of productive workforce development programs. Participants at the Crosscut Meeting called for an expansion of these programs.

The requirements review attendees noted that a vastly increased effort in broadening awareness of these programs through directed communications at all levels of DOE is needed to recruit the next-generation workforce — especially articulating, with examples, how these programs provide not only a rich scientific working environment but also marketable skills to applicants (and indicating where those jobs are in the marketplace).

4.1.9 Continuing to Collect Requirements, Identify Needs, and Deliver HPC Resources in Support of Science

The ability of research teams to effectively and efficiently use computational and data resources at scale is enhanced by continuing collaborations among domain experts, DOE managers, and staff from the national laboratories to help identify science-based requirements — as in these requirements reviews — and then deliver the resources. Close collaborations will identify needs for resources at different scales and levels of complexity, for example, to explore parameter space and quantify uncertainties or integrate with emerging data-centric workflows [NP 4.1].

4.2 Ongoing Facility Efforts and Programs Already Under Way to Reach an Exascale Ecosystem

The ASCR facilities form a solid foundation upon which to build toward the exascale era, and the computing and networking capabilities provided by the ASCR facilities will remain at the forefront as the exascale transition unfolds. The Exascale Computing Project will enable the procurement of exascale systems and is building on the capabilities and collaborations (both research and production) at the facilities. The facilities and ECP are briefly described below.

4.2.1 ESnet

The Energy Sciences Network (ESnet) is the U.S. Department of Energy’s high-performance networking facility, engineered and optimized for large-scale science. ESnet interconnects the entire national laboratory system, including its supercomputer centers and user facilities — enabling tens of thousands of scientists to transfer data, access remote resources, and collaborate productively. ESnet serves more than 50 DOE sites, connecting them to more than 150 research and commercial networks worldwide. As part of its mission, ESnet provides a state-of-the-art, high-speed network to the DOE laboratories and facilities that will deploy and use the exascale systems enabled by the Exascale Computing Project.

In the context of the deployment of pre-exascale systems and the transition to the exascale era, ESnet will be a key contributor to the beneficial impact of these next-generation HPC systems on the DOE science mission. This effort will be especially true in the data space, as evidenced by the large number of cross-cutting data requirements, which explicitly call out multisite workflows (Section 3.2). Data transfers to and from exascale systems, analysis of experiment and observational data using exascale applications, and access to exascale systems by expert science teams addressing the most important mission challenges will all be enabled by ESnet. ESnet staff work closely with scientists and science collaborations to optimize multisite workflows, and ESnet’s research and development efforts in advanced networking, network automation, and next-generation network services help ensure that the network will continue to contribute to the success of the DOE science mission in the exascale era.

4.2.2 LCF Programs

In 2013, the LCFs' 10-year mission was stated as

...continuing to accelerate scientific discovery and engineering progress by providing multi-petaflop, approaching exaflops, of sustained computational performance coupled with a dedicated team of computational scientists that enable broad use and applicability of HPC technologies coupled with advanced data science infrastructure at the highest level of capability and with exceptional price/performance (DOE-ASCR 2013).

The vision was that in 2023, the two LCF centers — the ALCF at Argonne and the OLCF at Oak Ridge — will be delivering world-changing science on sustained exascale performance systems. Following are descriptions of exascale-oriented programs hosted at each LCF at present:

- **The Early Science Program (ESP) at ALCF.** As part of the process of bringing a new supercomputer into production, the ALCF hosts ESP to help ensure that its next-generation systems are ready to hit the ground running. The intent of the ESP is to use the critical pre-production time period to prepare key applications for the architecture and scale of a new supercomputer and to solidify libraries and infrastructure to pave the way for other production applications to run on the system. In addition to fostering application readiness, the ESP allows researchers to pursue innovative computational science projects not possible on today's leadership-class supercomputers.
- **Center for Accelerated Application Readiness (CAAR) at OLCF.** In preparation for the next-generation supercomputer, Summit, OLCF's CAAR program enables application development teams and staff from the OLCF Scientific Computing group to collaborate. CAAR is focused on redesigning, porting, and optimizing application codes for Summit's hybrid CPU-GPU architecture. Through CAAR, code teams gain access to early software development systems, leadership computing resources, and technical support from the IBM/NVIDIA Center of Excellence at Oak Ridge National Laboratory. The program culminates with each team's scientific grand-challenge demonstration on Summit.

4.2.3 NERSC

In the fall of 2014, NERSC launched the NERSC Exascale Science Applications Program (NESAP), a collaborative effort in which NERSC has partnered with code teams and library and tools developers to prepare for the NERSC-8 Cori many-core architecture.

During this period, 20 project teams, selected by DOE reviewers and guided by NERSC, Cray, and Intel, have undertaken intensive efforts to adapt software that could take advantage of the new Cori system's Knights Landing many-core architecture. They also receive access to early hardware, special training, and preparation sessions. The teams are using the resultant codes to produce path-breaking science on an architecture that may represent an approach to exascale systems, helping to advance DOE SC missions.

4.2.4 Exascale Computing Project (ECP)

The ECP is a collaborative effort of two DOE organizations — the Office of Science and the National Nuclear Security Administration. As part of the National Strategic Computing Initiative (NSCI), ECP's multi-year mission is to maximize the benefits of high-performance computing for U.S. economic competitiveness, national security, and scientific discovery. By pursuing development of a capable exascale ecosystem — one that encompasses applications, system software, hardware technologies and architectures, and workforce development — ECP seeks to meet the scientific and national security mission needs of DOE in the mid-2020s timeframe.

In September of 2016, ECP announced its first funding awards totaling nearly \$40 million for 22 projects. This first round includes funding for a broad set of modeling and simulation applications with a focus on portability, usability, and scalability. A key consideration in the selection process was each team's emphasis on co-design of the applications with the ECP's ongoing development of hardware, software, and computational capabilities, including physical models, algorithms, scalability, and overall performance. Given this emphasis on co-design, teams from 45 research and academic organizations received awards.

5 REFERENCES

Bonoli, P.T., R. Parker, and S.J. Wukitch, et al., 2007, “Wave-particle studies in the ion cyclotron and lower hybrid ranges of frequencies in Alcator C-Mod,” *Fusion Science and Technology*, Volume 51 (3), p. 401.

Chowdhury, J., W. Wan, Y. Chen, S.E. Parker, R.J. Groebner, C. Holland, and N.T. Howard, 2014, “Study of the L-mode Tokamak Plasma ‘Shortfall’ with Local and Global Nonlinear Gyrokinetic δf Particle-in-Cell Simulation,” *Physics of Plasmas* 21, 112503; <http://dx.doi.org/10.1063/1.4901031>, accessed May 13, 2016.

DOE-ASCR (U.S. Department of Energy-Advanced Scientific Computing Research), 2013, *DOE Leadership Computing Facility’s Strategic Plan for Advancing Fundamental Discovery and Understanding in Science and Engineering over the Next Decade 2014–2023*, Office of Science, March 1.

DOE-BESAC (Basic Energy Sciences Advisory Committee), 2015, *Challenges at the Frontiers of Matter and Energy: Transformative Opportunities for Discovery Science*, November, http://science.energy.gov/~media/bes/besac/pdf/Reports/CFME_rpt_print.pdf, accessed April 7, 2016.

DOE-SC (Office of Science), 2015, *The Office of Science Fusion Energy Science Program: A Ten-Year Perspective (2015–2025)*, Office of Fusion Energy Sciences, December, http://science.energy.gov/~media/fes/pdf/program-documents/FES_A_Ten-Year_Perspective_2015-2025.pdf, accessed October 14, 2016.

DOE and NSF (National Science Foundation), 2015, *Reaching for the Horizon: The 2015 Long Range Plan for Nuclear Science*, October, DOE Office of Science, Office of Nuclear Physics and NSF Division of Physics, Nuclear Physics Program, http://science.energy.gov/~media/np/nsac/pdf/2015LRP/2015_LRPNS_091815.pdf, accessed September 2016.

FES (Office of Fusion Energy Sciences), 2016, *Fusion Energy Sciences (FES)*, U.S. Department of Energy, Office of Science, <http://science.energy.gov/fes/>, accessed June 23, 2016.

Minamoto, Y., and J. Chen, 2016, “DNS of a Turbulent Lifted DME Jet Flame,” *Combustion and Flame*, Volume 169, pp. 38–50, July, <http://www.sciencedirect.com/science/article/pii/S0010218016300542>, accessed May 11, 2016.

National Academies Press, 2013, *Nuclear Physics: Exploring the Heart of Matter*, National Research Council Report, Paperback, ISBN: 978-0-309-26040-4, DOI: 10.17226/13438, <https://www.nap.edu/catalog/13438/nuclear-physics-exploring-the-heart-of-matter>, accessed October 2016.

Shimamura, K., F. Shimojo, R.K. Kalia, A. Nakano, K. Nomura, and P. Vashishta, 2014, “Hydrogen-on-Demand Using Metallic Alloy Nanoparticles in Water,” *Nano Letters*, Volume 14, Issue 7, pp. 4090–4096.









Splettstoesser, T., 2016, SCISTyle, Berlin, Germany, <http://scistyle.com>, accessed March 6, 2017.

This page is intentionally left blank.

6 ACRONYMS AND ABBREVIATIONS

ALCF	Argonne Leadership Computing Facility
Argonne	Argonne National Laboratory
ASCR	Advanced Scientific Computing Research
ATLAS	Argonne Tandem Linac Accelerator System
ATPESC	Argonne Training Program for Extreme-Scale Computing
BER	Biological and Environmental Research
BES	Basic Energy Sciences
BESAC	BES Advisory Committee
BNL	Brookhaven National Laboratory
BSSD	Biological Systems Science Division
CAAR	Center for Accelerated Application Readiness (OLCF)
CESD	Climate and Environmental Sciences Division
CMB-S4	Cosmic Microwave Background-Stage 4
CMS	Compact Muon Solenoid (CERN detector)
CSGF	Computational Science Graduate Fellowship
DNA	deoxyribonucleic acid
DOE	U.S. Department of Energy
DUNE	Deep Underground Neutrino Experiment
E3SM	Energy Exascale Earth System Model
ECP	Exascale Computing Project (DOE)
EIC	Electron Ion Collider
EM	electron microscopy
EMSL	Environmental Molecular Sciences Laboratory
EP	energetic particle
ESM	Earth System Model
ESP	Early Science Program (ALCF)
FES	Fusion Energy Sciences
FRIB	Facility for Rare Isotope Beams
GPU	graphics processing unit










HEP	High Energy Physics
HPC	high-performance computing
HTC	high-throughput computing
I/O	input/output
IAM	integrated assessment model
JGI	Joint Genome Institute
JLab	Thomas Jefferson National Accelerator Facility
LHC	Large Hadron Collider (CERN)
MHD	magnetohydrodynamic
MPI	message passing interface
MSU	Michigan State University
NERSC	National Energy Research Scientific Computing Center
NESAP	NERSC Exascale Science Applications Program
NP	Nuclear Physics
NSF	National Science Foundation
OLCF	Oak Ridge Leadership Computing Facility
PB	petabyte(s)
PRD	priority research direction
QCD	quantum chromodynamics
QGP	quark-gluon plasma
RF	radio frequency
RHIC	Relativistic Heavy Ion Collider
SC	Office of Science (DOE)
SciDAC	Scientific Discovery through Advanced Computing
UQ	uncertainty quantification



APPENDICES

Executive Summaries

This page is intentionally left blank.



APPENDIX A

Executive Summary from the Basic Energy Sciences Exascale Requirements Review Report

This page is intentionally left blank.

EXECUTIVE SUMMARY

Abstract

Computers have revolutionized every aspect of our lives. Yet in science, the most tantalizing applications of computing lie just beyond our reach. The current quest to build an exascale computer with one thousand times the capability of today's fastest machines (and more than a million times that of a laptop) will take researchers over the next horizon. The field of materials, chemical reactions, and compounds is inherently complex. Imagine millions of new materials with new functionalities waiting to be discovered — while researchers also seek to extend those materials that are known to a dizzying number of new forms. We could translate massive amounts of data from high precision experiments into new understanding through data mining and analysis. We could have at our disposal the ability to predict the properties of these materials, to follow their transformations during reactions on an atom-by-atom basis, and to discover completely new chemical pathways or physical states of matter. Extending these predictions from the nanoscale to the mesoscale, from the ultrafast world of reactions to long-time simulations to predict the lifetime performance of materials, and to the discovery of new materials and processes will have a profound impact on energy technology. In addition, discovery of new materials is vital to move computing beyond Moore's law. To realize this vision, more than hardware is needed. New algorithms to take advantage of the increase in computing power, new programming paradigms, and new ways of mining massive data sets are needed as well. This report summarizes the opportunities and the requisite computing ecosystem needed to realize the potential before us.

In addition to pursuing new and more complete physical models and theoretical frameworks, this review found that the following broadly grouped areas relevant to the U.S. Department of Energy (DOE) Office of Advanced Scientific Computing Research (ASCR) would directly affect the Basic Energy Sciences (BES) mission need.

- **Simulation, visualization, and data analysis are crucial for advances in energy science and technology.**
- **Revolutionary mathematical, software, and algorithm developments are required in all areas of BES science to take advantage of exascale computing architectures and to meet data analysis, management, and workflow needs.**
- **In partnership with ASCR, BES has an emerging and pressing need to develop new and disruptive capabilities in data science.**
- **More capable and larger high-performance computing (HPC) and data ecosystems are required to support priority research in BES.**
- **Continued success in BES research requires developing the next-generation workforce through education and training and by providing sustained career opportunities.**

ES.1 Summary and Key Findings

The findings presented throughout this report are the result of a joint requirements review by BES scientists and DOE ASCR facilities teams. The mission of BES is to support fundamental research to understand, predict, and ultimately control matter and energy at the electronic, atomic, and molecular levels in order to provide the foundations for new energy technologies. The BES scientists focused on current scientific challenges in the areas of novel quantum materials and chemicals; catalysis, photosynthesis, and light harvesting; combustion; materials and chemical

discovery; soft matter; advances in algorithms for quantum systems; and computing and data challenges at BES user facilities (each described in greater detail in the following sections). These scientific challenges were then evaluated in the context of the computational science and computing ecosystems required to achieve the BES mission goals.

In order for BES science to take advantage of exascale computing architectures, software and algorithm development is required in all areas. Researchers need a new suite of sustainable and performant software tools, programming models, and applications to enable effective use of exascale systems. New algorithms are needed to enable codes to run efficiently on upcoming HPC architectures, allowing scientists to model larger materials and chemical systems with greater fidelity and predictive power. New mathematics are required to reduce the computational scaling of current algorithms and to develop new physical models. New software tools are needed to analyze and manage data produced by experiments, including tools that can drive complex workflows. New programming models are required to enable facile development of software that performs across multiple, distinct computing platforms. Attendees agreed that there is an opportunity for mathematics and computer science to have a transformational impact on BES science.

BES and the ASCR facilities are experiencing a pressing need to mature their capabilities in data science. Improvements and new capabilities at BES facilities are creating challenges that the community is not prepared to address. These include unprecedented growth in data volume, complexity, and access requirements; the need for curation of the massive amounts of data that are retained; and integration of diverse datasets from different experiments to enable new scientific conclusions. Efficient and effective use of BES facilities requires real-time access to ASCR HPC facility-class resources to support streaming analysis and visualization to guide experimental decisions.

Last, and of equal importance, a strong vision for workforce development is crucial for scientific success within future computing ecosystems. BES science will require cross-cutting, multidisciplinary teams of domain scientists, applied mathematicians, and computer scientists who can work together to address the broad range of challenges. Apparent to both BES and ASCR participants is a need to vastly improve the current training strategies and opportunities to develop the next generation of scientists and engineers.

Participants in the BES Exascale Requirements Review were asked to articulate the BES vision and grand challenges, identify priority research topics and computing needs, and outline a path forward for BES and ASCR. The following subsections summarize each of these topics in detail.

ES.2 Basic Energy Sciences Vision and Grand Challenges

BES supports fundamental research to understand, predict, and ultimately control matter and energy across scales from atoms to molecules to materials. This research probes the electronic and magnetic properties of matter and energy — an effort that is foundational to discovery of future energy solutions for the nation. Theory, computation, and the user facilities have played key roles in the BES portfolio to provide these solutions and have been recognized as essential for the future of the program. In 2015, the DOE's Office of Science charged the BES Advisory Committee (BESAC) to develop a series of “grand challenges” that would inspire and guide BES research. That tasking resulted in a report entitled *Challenges at the Frontiers of Matter and Energy: Transformative Opportunities for Discovery Science* (DOE-BESAC 2015), which identified five new transformative opportunities. If realized, these would lead to the control of matter and energy at the molecular, atomic, and quantum levels and could spark revolutionary changes in technologies to help us meet some of humanity's most pressing needs, including the need for renewable, clean, and affordable energy. The transformative opportunities are expressed in the BESAC report as three themes:

- Mastering Hierarchical Architectures and Beyond-Equilibrium Matter
- Beyond Ideal Materials and Systems: Understanding the Critical Roles of Heterogeneity, Interfaces, and Disorder
- Harnessing Coherence in Light and Matter

and two cross-cutting opportunities:

- Exploiting Transformative Advances in Imaging Capabilities across Multiple Scales
- Achieving Revolutionary Advances in Models, Mathematics, Algorithms, Data, and Computing

The energy systems of the future — whether they tap sunlight, store electricity, or make fuel from splitting water or reducing carbon dioxide — will revolve around materials and chemical changes that convert energy from one form to another. These advanced materials and chemical processes are not found in nature; they must be designed and fabricated to exacting standards using principles revealed by basic science. Acquiring the ability to reveal this fundamental knowledge has largely come about as a result of making extremely rapid progress in our ability to analyze and simulate experimental data and systems.

Not surprisingly, one of the transformative opportunities is to develop revolutionary advances in models, mathematics, algorithms, data, and computing. Today we are on the verge of standing up 100-petaflop machines. Furthermore, the Office of Science envisions reaching the exascale level of computing within the next decade (with machines capable of performing a million trillion floating-point calculations per second). This enormous growth in computational power, coupled with major advances in other fields, such as the development of coherent light sources and increased imaging resolution, is accelerating the pace of materials and chemical sciences research that will enable basic understanding and control at the quantum level.

The BES program also operates major scientific user facilities to serve researchers from universities, national laboratories, and private institutions. The availability of state-of-the-art BES user facilities has opened new avenues of research that were not available just a few years ago. The planned and potential upgrades of several BES facilities in the next five to ten years will support more detailed and expanded experiments, producing an exponential growth of data. To manage and extract useful scientific information from those experimental data will require computational resources beyond the capabilities of current resources.

In addition, a strong connection is required between new computing landscapes and experimental and theoretical toolsets to realize transformational opportunities. This review articulated the requirements of BES research and user facilities for models, mathematics, algorithms, data, computing, and workforce development. Implicit in these requirements is the need to increase engagement with ASCR facilities and resources with the goal of developing sustained partnerships to make full use of the ASCR resources to advance the BES science mission.

ES.3 Priority Research Topics and Computing Needs

Review participants focused on eight areas in which advancement in the transformative opportunities can be achieved through key and sustained efforts in computation, simulation, and advanced tool design.

ES.3.1 Novel Quantum Materials and Chemicals

Technologies of the future will rely upon the specific combinations of elements, compounds, materials, and phases whose behaviors are “emergent”: synergistically enhanced and not predictable from studying the components in isolation. The range of technological needs is vast, from

alternatives to silicon electronics, including spin- and Mott-tronics, optical and neuromorphic device components that take computing past the end of the “Moore’s Law” barrier; to high-transition-temperature superconductors that move electricity with no loss of energy; to better thermoelectrics that capture and utilize waste heat; to new, strong, earth-abundant magnets for turbine engines; to extractants for the separation of heavy elements in waste mixes; and to better photovoltaics. The need for new and improved functionalities can be met by new generations of “quantum materials and chemicals,” with greatly enhanced responses emerging from the “quantum chemistry” of strongly interacting electrons. Critical to the design of chemicals and materials with the properties we need is the development of new predictive theories, efficient and adaptive software, and exploitation of the full capabilities of new and future computing architectures at ASCR facilities, coupled synergistically with advanced experimental facilities at BES nanocenters and photon and neutron sources.

ES.3.2 Catalysis, Photosynthesis and Light Harvesting, and Combustion

Catalysis is the essential technology for accelerating and directing chemical transformation and is a key to efficient fuel production and industrial processes. To realize the full potential of catalysis for energy applications, scientists must develop a profound understanding of catalytic transformations so that they can design and build effective catalysts with atom-by-atom precision and can convert reactants to products with molecular precision, requiring a fundamental understanding of catalytic processes occurring in multiscale, multiphase environments. Likewise light harvesting and photosynthetic pathways hold the promise of efficient, inexpensive power sources if the processes can be understood and manipulated. Until these alternative energy sources are economically available, combustion will continue to be a dominant mode of energy conversion for transportation, power generation, and industrial thermal processes. High-performance computing, in particular exascale computing, is playing and will play a central role in providing the insight needed to design these energy-efficient transformations that involve processes at multiple length and time scales under real-world — rather than the idealized — conditions that constitute current simulation and modeling scenarios. However, methodologies that deterministically describe the coupling between different scales must be solved to make full use of future HPC platforms.

ES.3.3 Materials and Chemical Discovery

In an ideal world, researchers could create a new material or chemical with exacting properties to meet the need at hand, thus saving extensive costs associated with experimental trial and error. However, predictive modeling of properties, tailored synthesis of chemicals, and control of materials require advances in modeling capabilities and hardware resources and computational interpretive software for experimental techniques — to enable the characterization of spatial and temporal fluctuations and of short-lived intermediates along synthesis and degradation pathways. New computational tools will enable (1) computational discovery of novel materials and chemicals with target properties (including hierarchical structures with multiple functionalities); (2) prediction of pathways to the synthesis of these materials and chemicals (with consideration of sustainability and green chemistry principles); and (3) prediction of their kinetic or thermodynamic stability and degradation pathways.

ES.3.4 Soft Matter

Soft matter provides unique and critical materials behavior in a wide range of industrial products. Polymers, surfactants, electrolytes, and microheterogeneous fluids have long been key components in a multitude of applications, including energy storage (e.g., batteries and capacitors) and energy production (e.g., photosystems), chemical separations, enhanced oil recovery, food packaging, chip manufacturing, and health care products. Soft materials composed of molecular and/or modular building blocks can provide the hierarchical complexity and tunability for making paradigm-shifting materials that can accomplish multiple tasks. The complexity of soft materials presents

scientific as well as computational challenges that make exascale computing a pivotal resource in achieving the goal of designing functional matter, which requires not only orders-of-magnitude greater scalability in both dimensional and time scales, but also seamless integration with exabyte big data analytics and mining so as to extract maximal scientific knowledge.

ES.3.5 Advances in Algorithms for Quantum Systems

Implicit in many of the BES mission phenomena is the need to develop truly multiscale methods that can span multiple time and length scales in a seamless and self-consistent manner. However, realistic simulations of complex materials and chemical problems remain out of reach due to their high cost. Truly predictive simulations will require development of robust hierarchical theories and algorithms to treat electron correlations across all relevant length scales. To take advantage of exascale systems, it is necessary to develop highly parallel, low-scaling algorithms for each scale of the system (quantum and classical) and for multiscale methods. In particular, because computer configurations have ever-increasing numbers of computational elements on each node, the newly designed algorithms must be able to take advantage of multilevel parallelism with multiple layers of memory and communication hierarchies.

ES.3.6 Computing and Data Challenges at BES Facilities

Computing and data challenges may be characterized as streaming analysis and steering of experiments, multimodal analysis of results from different instruments, and long-term data curation. The growing complexity of the analysis process (mixing fast data analysis and numerical modeling) will require capabilities beyond the petascale-level capabilities that the ASCR facilities currently offer. Addressing these challenges by providing for the analysis, management, and storage of user data signals a fundamental change in the operation and responsibility of BES user facilities. In addition, the computational capabilities of the future will provide a platform for real-time modeling and simulation so that experiments can be augmented and understood as they are in progress. This coupling, along with the need to manage data and the ability to steer and make decisions during the experiment to optimize the scientific outcomes, will require significant “on-demand” exascale types of computational resources to deliver the necessary feedback and insights in real time as the experimental process unfolds.

ES.3.7 Mathematics and Computer Science Transforming BES Science

Bridging the gap between BES scientific goals and ASCR computing capabilities will fundamentally depend on the collective abilities of our science domains and facilities to deliver transformative breakthroughs in mathematics and computer science — an objective that will require investing in and capitalizing on evolving, state-of-the-art mathematics and computer science. Mathematics will need to be developed that enables order-of-magnitude improvements in speed and accuracy in predictive materials and chemistry modeling. Mathematical algorithms and unified software environments must be delivered to allow fast, multimodal analysis of experimental data across different imaging modalities and DOE facilities. Software tools must be built that will make efficient programming of tomorrow’s machines as straightforward as programming today’s laptops. Finally, these three advances must be tied together to significantly advance our understanding in scientific domains.

ES.3.8 Next-Generation Workforce

The complexities and multiple layers of hierarchy of next-generation programming environments will include developing the physical and mathematical models, expressing the scientific workflow, developing the numerical algorithms, decomposing the algorithm to the optimal level of task granularity, expressing fine-grained parallelism in domain-specific languages (DSLs), and ensuring

that all of the layers of the programming model and runtime have the right abstractions to enable flexibility and performance. Providing the kinds of education and training deliverables to prepare future computational and domain scientists and computational software developers for exascale computing presents a major challenge.

ES.4 Path Forward









The support and development of our evolving computing ecosystem relies on continued collaboration between BES and ASCR. Rooted in the discussions about the BES vision, research directions, and computing needs, four categories grew out of the review: methods development, computational environment, data, and cross-community engagement.

Methods development includes researching new, scalable algorithms; building better memory management techniques; exploring data-related methods; incorporating the full physical and chemical environment in a theoretically sound and yet affordable manner; providing parallel-in-time and new sampling approaches; and developing end-to-end analysis pipelines to solve workflow, visualization, and optimization concerns.

The **computational environment** addresses requirements for access, scheduling, and software environments. A greater demand for near real-time (or on-demand) computing and analysis is driven by HPC needs at experimental facilities during or shortly after experiments. At the same time, the HPC market appears to be delivering exascale and affordable petascale within the same time frame. This creates opportunities for improvement of tools not only for exascale systems but also for the petascale mid-range of tomorrow.

The **scale of data** generated these days creates challenges in performance, locality, analytics, access, and curation. Today, we find that with the increase of data outputs, analysis has moved from the workstation back to the supercomputing center. Analysis now includes machine learning and deep learning techniques. Curation comes into play because communities want the capability to provide access and knowledge of the location of the data in a push to the community, as well as the ability to combine data and analysis to discover new insights into experiment and simulations.

Finally, **cross-community engagement** brings together three fronts: workforce development, collection and feedback with standards, and the development of training materials and best practices. While there are existing efforts, there is an opportunity to create cohesive approaches on one or more of these fronts. The formation of new institutes to connect mathematicians and computer scientists and engineers with BES researchers and facilities is essential toward progress in addressing the challenges of exascale computing.



APPENDIX B

Executive Summary from the Biological and Environmental Research Exascale Requirements Review Report

This page is intentionally left blank.

EXECUTIVE SUMMARY

Abstract

Understanding the fundamentals of genomic systems or the processes governing impactful weather patterns are examples of the types of simulation and modeling performed on the most advanced computing resources in America. High-performance computing and computational science together provide a necessary platform for the mission science conducted by the Biological and Environmental Research (BER) office at the U.S. Department of Energy (DOE). This report reviews BER's computing needs and their importance for solving some of the toughest problems in BER's portfolio.

BER's impact on science has been transformative. Mapping the human genome, including the U.S.-supported international Human Genome Project that DOE began in 1987, initiated the era of modern biotechnology and genomics-based systems biology. And since the 1950s, BER has been a core contributor to atmospheric, environmental, and climate science research, beginning with atmospheric circulation studies that were the forerunners of modern Earth system models (ESMs) and by pioneering the implementation of climate codes onto high-performance computers.

ES.1 Summary and Key Findings

This review found the following broadly grouped areas relevant to the BER mission.

- Scalable data processing, data analysis, machine learning, discrete algorithms, and multiscale multiphysical simulation are crucial for advancement of biological and environmental systems science.
- Innovations in representation, search, and visualization of large-scale, heterogeneous, ontologically rich primary and derived biological and contextual data (e.g., abiotic environmental information) are crucial for input to and validation of these methods.
- New architectures, data transport protocols, software libraries, and languages are necessary to create a platform for community tool development and use supporting interactive and seamless interoperation of both mid- and large-scale cluster resources and enterprise-class computing environments.
- Algorithms are needed for Earth system processes such as atmospheric dynamics and clouds, oceans, tracer transport, coastal processes, and land that scale effectively on advanced computer architectures.
- Capability is needed for large ensembles, together with methods to effectively capture statistical information on Earth systems and climate variability beyond brute-force ensembles.
- Fusion of model simulations and observational data must take place for better model initialization, uncertainty analysis, validation, and tuning.
- Earth system model complexity requires exascale systems built with powerful general purpose nodes with large amounts of high-bandwidth memory.
- Creation of the necessary system components requires a workforce trained deeply not only in the core computational, data scientific, mathematical, and natural scientific disciplines that underlie the above technologies but in how to co-design and develop tools that support open-community development and research.

ES.2 BER Vision and Grand Challenges

The scope of BER programs is vast, and the challenges to integrate, analyze, test, and simulate processes, data, and information obtained from widely diverse disciplines are daunting. Computational capabilities at the exascale offer an opportunity to study and couple the most critical processes through model simulation and to combine and analyze diverse data sets collected across multiple disciplines and a range of spatial and temporal scales, as well as to run numerical models and simulate interactions among biological, biogeochemical, and physical processes from molecular to Earth system scales and from current to future states to solve critical scientific research problems in support of DOE missions.

BER's programs are divided between two divisions, the Biological Systems Science Division (BSSD) and the Climate and Environmental Sciences Division (CESD). Combined, the research programs and the user facilities within each division encompass laboratory- to field-based research and observation, as well as numerical modeling across the full range of spatial and temporal scales of interest to BER, DOE, and DOE's Office of Advanced Scientific Computing Research (ASCR). While the focus of each division's efforts differs, many of the computational, experimental, data management, analysis, and simulation challenges are similar. BER seeks to develop advanced experimental and observational methods, together with advanced computational approaches, to connect large and diverse data sets within multiscale modeling frameworks and thus enable understanding, prediction, and testing of complex systems.

ES.2.1 BSSD Vision and Grand Challenges

Biological systems science integrates multidisciplinary discovery- and hypothesis-driven science with technology development on plant and microbial systems relevant to national priorities in sustainable energy and innovation in life sciences. Biology as a research endeavor is changing rapidly from a traditionally qualitative science to a much more quantitative science. This change has been driven by revolutionary advances in molecular measurement technologies (including genome sequencing) and imaging and new tools and methods for biotechnology.

Major goals for BSSD from the 2015 strategic plan include efforts to:

- Provide a basic understanding of plant and microbial biology to lay the foundation for the production of biofuels and bioproducts from sustainable plant biomass resources.
- Develop the fundamental understanding of genome biology needed to design; modify; and optimize plants, microbes, and biomes for beneficial purposes.
- Gain a predictive understanding of biological processes controlling the flux of materials (e.g., carbon, nutrients, and contaminants) in the environment and how these processes affect ecosystem function.
- Develop the enabling computational, visualization, and characterization capabilities to integrate genomic data with functional information on biological processes.
- Exploit new technologies and correlative approaches to image, track, and measure key processes occurring at the molecular and cellular level within plant and microbial cells.
- Broaden the integrative capabilities within and among DOE user facilities to foster interdisciplinary approaches to BER-relevant science and aid interpretation of plant, microbe, and microbial community biology (DOE-BER 2015).

ES.2.2 CESD Vision and Grand Challenges

The Climate and Environmental Sciences Division supports fundamental science and research capabilities that enable major scientific developments in the coupled Earth system. These capabilities focus on atmosphere, ocean, cryosphere, and terrestrial processes understanding and

modeling in support of DOE's mission goals for basic science, energy, and national security. CESD leads important international process and modeling research involving clouds, aerosol-cloud interactions, subsurface biogeochemistry and hydrology, terrestrial systems, and integrated human-Earth system modeling. Unique capabilities include research designed to integrate models and measurements, a focus on scale-dependent process representations in models, investigations into the dominant uncertainties in Earth system behavior, and development of Earth system codes to run efficiently on advanced computer architectures.

The Climate and Environmental Sciences goals articulated in the 2012 CESD strategic plan are to:

- Synthesize new process knowledge and innovative computational methods advancing next-generation, integrated models of the human-Earth system.
- Develop, test, and simulate process-level understanding of atmospheric systems and terrestrial ecosystems.
- Advance fundamental understanding of coupled biogeochemical processes in complex subsurface environments to enable systems-level environmental prediction and decision support.
- Enhance the unique capabilities and impacts of the Atmospheric Radiation Measurement (ARM) and Environmental Molecular Sciences Laboratory (EMSL) scientific user facilities and other BER community resources to advance the frontiers of Earth system and environmental science.
- Identify and address science gaps that limit translation of CESD fundamental science into solutions for DOE's most pressing energy and environmental challenges.

ES.3 Priority Research Directions (Topics) and Computing Needs

ES.3.1 Biological Systems Sciences

The core ambitions of Biological Systems Science to predict, control, and design the function, environmental health, and productivity of plants, microbes, and the biomes they support require discovery and characterization of the causal linkages in biomolecular networks within and across organisms and abiotic interfaces. These, in turn, control the dynamics of cellular and organismal populations that respond and adapt to changing environmental conditions and reciprocally affect small and large changes in the environment and Earth systems. Their activities can be harnessed for the production of energy and other renewable resources or the mitigation of energy production processes. There are nearly 400,000 species of plants known on Earth with genomes far more complex than those of humans and with metabolic capabilities we have not nearly explored. There are more than 10^{30} microbes on Earth — smallish genomes that have solved the problem of living anywhere life can survive — and they support a dizzying array of other organisms and environmental processes. We have only scraped the surface of the study of microbes.

In this review, we focused on how scaling computational resources could accelerate the discovery and application of biological knowledge effectively. What emerged is the need for highly data-aware systems that serve scaling and integrating mid-scale data analytical efforts into large multiscale modeling codes that could exploit the next-generation architectures emerging at the exascale. Driven by the pressing increase in the size, heterogeneity, and structural complexity of biological data sets, participants set specific goals for high-risk, high network capacity machines and high memory for many of the key types of algorithms involved.

ES.3.1.1 Multiscale Biophysical Simulation from Molecules to Cells

A truly causal understanding of biological function requires an understanding of the often-complex heterogeneous physics that drives living processes. If we are to effectively harness the genetic potential of the Earth for novel catalysts for energy and renewable chemical production, or be able to interpret the genomes of organisms glimpsed only through the power of deoxyribonucleic acid

(DNA) sequencing, we need to obtain a deep mechanistic description of how proteins, ribonucleic acids (RNAs), and other biomolecules perform their functions. We are becoming more adept at using a combination of direct comparisons of new molecular structures to old ones, and applying modeling approaches ranging from quantum mechanics through molecular dynamics to coarser-grained simulations all the way up to the cellular level. However, many of the core individual codes need to be upgraded to exploit the new proposed architectures. New algorithms for incorporating the vast amounts of data derived from innovations in genomics, molecular imaging, structural biology, and spectroscopy to accelerate and expand the predictive capacity of these codes need to be developed. More rigorous tools are necessary for scaling modeling beyond single macromolecules or even their complexes to whole cells, which will allow whole communities of researchers to work together. Finally, algorithms that allow multiresolution simulations, which may be mechanistically detailed in some parts while more phenomenological in others while tracking and quantifying uncertainty, are considered key elements of success. We are looking for new approaches that enable orders-of-magnitude jumps in the time and space scales that can be simulated by such models, along with the architectures that also support ensemble methods so that training and uncertainty/sensitivity calculations can be performed more effectively.

ES.3.1.2 Mapping Sequence to Models

At least three disruptive experimental technologies are driving the need for extraordinary computational innovation and scaling in understanding the function encoding in the genomes of organisms: sequencing, molecular imaging methods such as cryo-electron microscopy (EM), and sophisticated molecular functional assays such as those made available by innovations in mass spectroscopy. Here, the data rates and data sizes are scaling exponentially; and algorithms that simply “process” the raw data into genes and genomes, protein structures, and chemical activities, respectively, are computationally intense, require large memory, and require a large amount of disk space. When enriching the derived sequences and structures with functional data through sophisticated phylogenomic analyses, large-scale molecular docking, and machine learning on large functional measurement data sets and natural language-processed literature, for example, the need for new algorithms that can move onto exascale machines becomes more critical. One of the key capabilities is the need to rerun prediction algorithms and phylogenetic estimations constantly as new data become available. Thus, it will be critical to develop and maintain key resources such as constantly updated taxonomic and gene family phylogenetic trees; open-access publications; and data in large-scale functional genomic resources, such as the genome portals at the Joint Genome Institute (JGI), the DOE Systems Biology Knowledgebase (KBase), and their collaboration with international repositories such as the Protein Data Bank (PDB), the National Center for Biotechnology Information (NCBI), and UniProt (Universal Protein Resource). These tools will provide the critical linkages, along with the new algorithms in cellular modeling identified in Section ES.3.1.1 to map genotype to phenotype in organisms with biotechnological and environmental significance and to understand the ecology and evolution of biological populations significant for Earth processes and biomass production and conversion for energy and renewables.

ES.3.1.3 Microbes to the Environment

It is here that these two major divisions within DOE’s BER office (CESD and BSSD) come most strongly together with the ambition to incorporate the predictions and data described above into integrated models of biomes and their environmental functions and to propagate their effects up to the Earth system scale. Plants and microbes are significant players in the global carbon, nitrogen, sulfur, and phosphorous cycles; and understanding how they respond to and affect environmental change at all scales — and ultimately the health and resilience of our own species — is critical given the large footprint that energy-derived processes have in these phenomena. Predictively modeling the effects of plant/microbe activities to Earth processes requires representing the

processes by which plants uptake carbon and sequester it in the soil wherein it is converted and processed back into carbon dioxide and other gases by microbes. These processes depend on mechanisms occurring at the scale of pores in soil particles and through reactive transport in flowing watersheds. The output from these models, which are developed based on process research and modeling across BER, can then feed the large Earth system models from CESD. Here the detailed modeling community activity regarding microbe-microbe, microbe-plant, and plant-plant interactions becomes critical, along with interaction with other environmental creatures (nematodes, insects, etc.). The computational requirements include new algorithms for (1) large-scale network inference, model selection, and training on micro- and geospatially resolved genomic and biological functional data; and (2) multiscale, multiphysical simulation of genome-informed reactive transport models. These algorithms must be able to run the new exascale architectures with many of the same requirements as the molecular simulations in Section ES. 3.1.1. New community codes for integrating biological data and simulation models — having a power similar to what the Earth scientists have done with Earth system models — will be crucial to innovation in this area. Success means deriving defensible predictions of environmental change and more efficient explorations of routes to achieving beneficial outcomes.

ES.3.1.4 Biological Big Data Challenges

One of the special challenges in biological systems science is in the size, quality, and structure of the data that need to be analyzed and understood to make effective predictions about biological identity, function, and behavior. Even a single bacterial cell is a spatially structured, mechanically and electrically coupled system composed of approximately ten billion molecules drawn from around 10,000 different chemical species. The interactions among all of these species define dynamic chemical networks, mechanical engines such as motility apparatuses, active structure platforms such as membranes, and motor-driven construction systems such as DNA replication systems. The complexity grows when talking about groups of cells in tissues or groups or organisms in communities. Sequence data, molecular structure, molecular abundance and activities, spatial imaging, and population numbers are all being measured — among other things — with more or less precision at rapidly increasing rates and scales. There is immense pressure to improve the efficacy and efficiency of algorithms that are used to cluster, reduce dimensionality, compute graphs of probabilistic dependencies, and generally find models of the data. This effort needs to be complemented by knowledge systems that also incorporate ontological systems for classifying data and relating different types together. These all must be easily accessible and transportable across computational infrastructures both for these primary data analytical algorithms and to service the modeling tools above. Innovations in data transport, data management, data processing, knowledge representation, machine learning, and mixed mechanistic and statistical simulation will be necessary to realize the full value of the biological data derived from DOE's flagship facilities like the JGI, Advanced Light Source (ALS), EMSL, and others. Data science libraries suitable for working both in cluster and enterprise environments, database systems accessible on both, and sophisticated interactive data visualization all require significant innovation in this area.

ES.3.2 Climate and Environmental Sciences

Increasingly, environmental systems simulation represents physical, chemical, and biological processes spanning ever-broader ranges of spatial and temporal scales. The interactions among scales produce complex feedbacks that determine the system behaviors. The types of modeling and simulation tools range from molecular and process-scale models to fully coupled global Earth system simulation codes. New measurement and observational capabilities from laboratory and field studies inform the development of these simulation capabilities and provide critical data to test their fidelity. Although reductionism requires complex modeling to be broken into tractable-sized efforts, the interaction and feedbacks among the parts in integrated modeling systems have emerged

as a central challenge in simulation, data collection, data management, and scientific analysis. Exascale-class simulation, data management, and network capabilities are essential to the fusion of simulation and data analysis to advance scientific discovery and the rapid use of new understanding to solve real-world problems.

ES.3.2.1 Atmospheric Simulation and Data Assimilation within the Earth System

In the next 5–10 years, advanced computing resources may enable researchers to substantially improve the representation of clouds and their impacts on the energy and water cycles of the Earth system. Improving atmospheric simulation within Earth system models, such as the Energy Exascale Earth System Model (E3SM), will require integrating observations with a hierarchy of models ranging from direct numerical simulation through explicit turbulence fluid dynamics models, large-eddy simulation (LES) models, and regional to global atmospheric models to improve parameterizations for key processes (e.g., cumulus cloud convection, aerosol microphysics, and cloud physics) to properly represent subgrid-scale processes across models of differing resolutions. With increasing resolution and scale-aware physics parameterizations, the atmospheric and coupled Earth system models are better able to realistically model heavy precipitation events, droughts, floods, and other low-frequency, high-impact events with important consequences. Improved representation of clouds, aerosol-cloud interactions, and land-atmosphere interactions will enable more accurate simulation of the formation, maintenance, and dissipation of clouds that play crucial roles in determining cloud feedback and climate sensitivity. Data-model fusion, best exemplified by data assimilation, serves two critical purposes for atmospheric models. First, it confronts models with observational data so that the fidelity of model performance can be continuously evaluated, calibrated, and validated. Second, it provides the capability for the models to be initialized so that realistic simulations can be conducted. Advances in atmospheric data assimilation in the last several years have shown promising approaches for building data assimilation systems for Earth system models with minimal new algorithm and software engineering investments. Future research needs to develop data assimilation techniques at the appropriate spatial scales and with the appropriate targeted observations that will address specific science needs.

ES.3.2.2 Terrestrial and Subsurface Research

Mechanistic understanding of terrestrial and subsurface processes continues to improve, driven by hypothesis testing in a coupled framework of experimentation, observation, and modeling. Many land processes of importance to the integrated functioning of the Earth system operate on spatial scales much finer than those represented in the current generation of ESMs. Looking ahead 10 years, we expect the horizontal resolution of land processes in ESMs to increase from current high-resolution grids at 10–20 km toward resolutions of 1 km or finer with surface meshes structured around watersheds and related landforms. However, even at that future target resolution with the global land surface resolved as hundreds of millions of grid elements, many land processes still reside at subgrid scales. For example, hillslope hydrology representing lateral surface and subsurface flows occurs at scales of meters to tens of meters; surface inundation and associated biogeochemical dynamics connected to microtopographic variation in flat and gently sloping landscapes occur at scales of centimeters to meters; interactions among plants and microbial communities occur in the rhizosphere at scales of millimeters to tens of centimeters; interactions among microbial communities and the soil physical and chemical environment are localized on mineral surfaces at scales of microns to millimeters; and a host of biological processes operate at the cellular and subcellular scales in plants and microbes, and they interact directly with the physical and chemical environment with significant impacts at all larger scales up the entire globe. A major scientific challenge is to develop comprehensive and robust theories allowing process knowledge to migrate effectively up in scale from process-resolving to process-parameterized to improve the modeling of terrestrial and subsurface processes. Developing and exercising process-

resolving models in many geographic and functional spaces and developing rigorous approaches for up-scale knowledge migration founded on fine-scale models represent important challenges and opportunities relevant to exascale systems.

ES.3.2.3 Oceans and Cryospheric Research

The ocean-cryosphere system comprises the global ocean, including the main deep basins, marginal seas, coastal ocean, and estuaries along with all of the sea-ice and land-ice systems. It is estimated that over the twentieth century, the ocean system has absorbed approximately 90% of the heat trapped by greenhouse gases and that oceans have absorbed more than one-third of all anthropogenic carbon emissions. The cryosphere is undergoing the most rapid recent changes within the entire Earth system. This transition is particularly evident in the Arctic, where a transition toward a summertime sea-ice-free condition is under way. Meanwhile, abrupt sea-level rise could emanate from ocean/land-ice interaction around Antarctica. Gaining understanding and the ability to project ice-free conditions in the Arctic, potential changes in the rate of ocean uptake of heat and carbon, and projecting sea-level rise in the twenty-first century remain grand challenges. Both process-based studies and the geometry of ice cavities suggest that subkilometer resolution is needed in both the ocean and ice models in order to accurately represent the melting process at the ocean-ice interface that contributes to abrupt sea-level rise. The impacts of sea-level rise occur primarily during storm surges when an additional volume of ocean water finds its way into human and ecological systems at elevations above the high-tide elevation. Accurate simulation of inundation extent and depth during extreme weather events, such as hurricanes, also requires representation of processes occurring at the terrestrial-aquatic interface at subkilometer scale. Overall, access to exascale computing resources has the potential to dramatically improve the fidelity of the simulation of the ocean-cryosphere system and its coupling with other Earth system components. This improvement, in turn, will allow us to better understand the role of these systems in a variable climate, as well as to quantify the impact of a changing ocean and cryosphere on human systems.

ES.3.2.4 Earth System Models

Fully coupled Earth system models or ESMs integrate the physical and biogeochemical components of the Earth's climate to capture the many feedbacks in this complex system. Resolving processes at relevant space and time scales and providing decision-relevant information are driving requirements for very high spatial resolution and an increased use of integrated ensembles of simulations that can be enabled only by exascale computing systems. Some examples of high-priority, coupled-system research include projection of sea-level change, impacts of weather extremes, and improvements in estimation of climate sensitivity. Estimating the rate of sea-level rise and understanding the coastal impacts require integration across Earth system components, as well as high spatial resolution. Most of the economic impacts of climate variability result from extreme weather events. Predicting changes in the frequency of extreme events requires large ensembles of ESM integrations at high spatial resolution (~ 1 km) to resolve cloud and convective processes and generate probability distributions of weather events. Better simulations of climate sensitivity require improved representation of cloud changes, and an understanding of carbon uptake by land and ocean ecosystems and how vegetation changes in response to the physical climate. The computing requirements for generating large ensembles using atmospheric, oceanic, and Earth system models on the order of 100 members cut across many of these research areas. Exascale computing allows climate models to increase resolution, and high-resolution models help to remove the need for approximate parameterizations in favor of directly resolving processes. However, the added complexity of process-resolving modeling can paradoxically introduce more uncertainties and does not remove the requirement for large numbers of ensembles. Even with exascale computing, it will not be possible to run large simulation ensembles at the highest possible resolution, and therefore,

new computational and theoretical methods will be required to evaluate and reduce the uncertainties in the system. New uncertainty quantification (UQ) techniques will need to be developed to combine the results of different classes of simulation, which could also leverage the testing and tuning of simulations routinely performed during model development. Therefore, progress in Earth system modeling will require coordinated progress in better statistical and ensemble methods, increased model and process resolution, and improved UQ methods.

ES.3.2.5 Integrated Assessment and Impacts-Adaptation-Vulnerability Modeling

Integrated assessment models (IAMs) have historically focused on understanding the implications of human activity on the Earth system at the global scale. However, IAMs are increasingly coupled to other models, including both Earth system models and impacts, adaptation, and vulnerability (IAV) models. These model couplings present a variety of theoretical, operational, and computational challenges as they include multiple scales, processes, sectors, disciplines, institutions, and sets of heterogeneous data. As different questions may require different suites of models, the coupling infrastructure needs to be flexible, modular, and extensible. Three use cases are illustrative of the interactions among human and Earth system processes that warrant new coupled modeling applications and/or the integration of heterogeneous data across a range of spatial and temporal scales. First, DOE has recognized the need for an integrated science approach to informing the resilience of managed water and energy in the face of climate variability and other global and regional change drivers such as population growth and technological change. Climate effects may alter water availability for hydropower, nuclear power, and fossil power production and irrigation, while deployment of renewable energy technologies and biofuels have systemwide implications for water demands and uses. Second, most of the world's population lives in cities, so understanding how climate effects will affect the urban environment and its infrastructure systems is critical to developing effective strategies for water and energy resilience, as well as predicting climatic conditions at decision-relevant scales. Third, many cities and important infrastructure assets are located near the coast where near-coastal ocean dynamics drive mesoscale climate phenomena, and the dynamics of wave propagation determines the impacts of storm surge, tsunamis, and sea-level rise for human infrastructure. Understanding human-natural process interactions in coastal regions is essential for evaluating coastal vulnerability and assessing adaptive measures to mitigate flood risk. Addressing problems such as these will require careful selection of appropriate models and coupling strategies.

ES.3.2.6 Transforming Science through Exascale Capabilities: Model-Data Fusion and Testbeds

A model development testbed is a systematic, automated framework involving a combination of model and observations used to understand physical processes and to evaluate and identify sources of error in a model during its development. During the workflow of model testbeds, model simulation output can be compared to observations of the Earth system in order to identify errors in the model simulations and determine the specific model processes that need improvements. Testbeds can also be used to provide scientific insights into dominant processes and process interactions, as well as to increase our understanding of the role of various physical processes involved in a particular case study or meteorological event. The common challenge of testbeds in terms of computational needs is the large amount of simulation output with very high temporal frequencies required in order to study detailed processes. Efficient post-processing for model output, as well as application of proper metrics and diagnostics packages, is another challenge. Unique opportunities that would come about with increased computational capability for current DOE model testbeds include the possibility to study processes using high-resolution simulations and multiyear hindcasts, test computationally expensive parameterizations in high-resolution simulations, and incorporate more frequent use of instrument simulators. Three areas of priority

computer-intensive research directions that are most crucial for improving our atmospheric model development capability for the next generation of DOE's ACME model include developing the capability to generate initialized and coupled hindcasts for cloud parameterization testing and addressing coupled model bias; applying a nonhydrostatic regional refinement modeling framework for very-high-resolution cloud processes, cloud-aerosol interactions, and extreme events studies; and applying UQ techniques for routine and systematic model physics parameters estimates or tuning in these model testbeds.

ES.3.2.7 Transforming Science through Exascale Capabilities: Algorithms and Computational Science

Advances to exascale must wrestle with the “granular” building blocks of algorithms, programming languages and programming models, and software engineering to achieve portability. As Earth system models approach deployment on exascale architectures, the algorithms that will enable their use will necessarily be as broad and diverse as the problems these models seek to address. To this end, scientists have outlined a suite of requirements for Earth system models that should be addressed with new and expanded algorithms that must span model testing, integration with multiple scales, coupling, analysis, and understanding.

Simulation in the climate community is dominated by the nexus of the Fortran/C/C++ languages and the OpenMP/OpenACC programming models. While both Fortran and C/C++ are interoperable on almost all computing platforms available today, there is a strong desire for other productivity-oriented scripting languages in a distributed environment (e.g., Python/pyMPI), as well as the desire to move away from the flat MPI model, which may be functional on the next generation of systems but will lose a factor of 10 to 50 times on GPU-accelerated systems. To that end, pragma-based, hybrid programming models are currently being explored.

In the simplest characterization, the complexity of Earth simulation modeling would benefit most from exascale systems with powerful general purpose nodes and large amounts of high-bandwidth memory. However, there are equally complex challenges facing the continued evolution of today's petascale machines, most notably more complex system and central processing unit (CPU) architectures. In order to mitigate the risk associated with uncertainty in the future directions of machine architecture, the codes will need to be portable across two or more different exascale architectures. Here, portability has many aspects: compiler portability, performance portability, and scientific portability, among others. DOE is currently investigating two different pre-exascale architecture “swim lanes”: (1) a modest number of compute nodes with multiple multicore CPUs and multiple accelerators; and (2) a large number of compute nodes, each with many-core CPUs. Performance at the node level will come from efficient exploitation of these different architectural approaches, which do not conform with the idea of general purpose nodes. However, although the detailed approaches to achieving high performance are quite different, at some level these two architectural paths forward share many “general purpose” characteristics related to exposing the parallelism available in the modeling and simulation methodologies. Each will have different implementation requirements for achieving good computational performance and will require a robust programming and tools environment for facilitating portability across architectures and the ability to exploit future architectures beyond what is currently envisioned for exascale capability.

ES.4 Path Forward









The collaboration between BER and ASCR scientists and facilities will be crucial for the deployment of effective computational infrastructures that enable new scientific advances in the DOE biological systems science and climate and environmental science missions. The requirements have been categorized in the broad areas of methods development, computational environment, data, and communication and community involvement.

Methods development includes development of highly scalable and portable algorithms with accurate physical models; integration of models with physical measurements, including the associated uncertainties in both; and the integration of large-scale heterogeneous data sources with theoretical and mathematical approaches capable of representing multiscale, multiphysics system descriptions. The tight coupling of well-supported, hardened, distributable, reusable, and modular components into end-to-end modeling, analysis, and visualization workflows will enable deployment to a larger community.

The *computational environment* will need to address requirements for both the large-scale, production-class simulations and data analysis efforts on leadership facilities, as well as algorithm and application development on smaller platforms and persistent support for access and transport of large data sets. Common needs include approaches for integrating observational and experimental data in the development of descriptive and predictive modeling capabilities, and these capabilities have components in the software stack focused on data integration, weakly coupled systems for statistical methods, interactive testbeds, and flexible scheduling tools and policies.

Data from observational and experimental sources are proliferating, and the integration of these data into the development of physical models is a major and common concern for the entire BER community that depends increasingly on programs and facilities that generate large amounts of data in a wide variety of forms. Tools and workflows are required for capturing, representing, curating, and providing provenance for data, as well as the long-term, large-scale, and distributed storage and delivery capabilities needed to serve a large and diverse scientific community.

Communication and community involvement are crucial to dealing with the level of expertise required for the development of methods and algorithms for future architectures, which, in turn, underscore the need for dedicated developers as part of a computationally focused workforce within the user community with adequate experience and a reward structure that supports strong development efforts. Connected to this need is the realization that proposal mechanisms should focus not only on research but also on software development, testing, and performance optimization.



APPENDIX C

Executive Summary from the Fusion Energy Sciences Exascale Requirements Review Report

This page is intentionally left blank.

EXECUTIVE SUMMARY

Abstract

The additional computing power offered by the planned exascale facilities could be transformational across the spectrum of plasma and fusion research — provided that the new architectures can be efficiently applied to our problem space. The collaboration that will be required to succeed should be viewed as an opportunity to identify and exploit cross-disciplinary synergies.

ES.1 Summary and Key Findings

To assess the opportunities and requirements as part of the development of an overall strategy for computing in the exascale era, the Exascale Requirements Review meeting of the Fusion Energy Sciences (FES) community was convened January 27–29, 2016, with participation from a broad range of fusion and plasma scientists, specialists in applied mathematics and computer science, and representatives from the U.S. Department of Energy (DOE) and its major computing facilities. This report is a summary of that meeting and the preparatory activities for it and includes a wealth of detail to support the findings. Technical opportunities, requirements, and challenges are detailed in this report (and in the recent report on the Workshop on Integrated Simulation). Science applications are described, along with mathematical and computational enabling technologies. This review generated the following key findings:

- Progress in computation across the range of fusion and plasma topics in recent years has been dramatic. Advances can be attributed to coordinated improvements in theory, computational and mathematical algorithms, performance engineering, computing hardware and software platforms, and uncertainty quantification (UQ).
- Broader and deeper integration into multiphysics and multiscale domains is a critical next step and will be necessary to address many important questions. These are exascale-level problems. Dramatically increased computing needs are also driven by ensemble runs in support of uncertainty quantification.
- The technical implementations for practical and affordable exascale platforms will present a number of significant challenges to approaches and algorithms used in today's codes.
- Additional challenges are presented in the areas of fault tolerance, software engineering, workflows, data management, *in-situ* analytics, and visualization.
- Close collaboration among stakeholders in various communities will be crucial to overcoming these challenges and realizing the advantages afforded by the new platforms. To that end, a large and specific set of needs for improved computational techniques, programming models, tools, software libraries, and algorithms have been identified.
- Predictable and stable access to high-performance computing resources is essential if the returns from major programmatic investments in code development are to be realized. In general, closer integration of processes for funding people and projects on the one hand and provisioning computer time on the other could lead to more efficient and optimal outcomes.

ES.2 Fusion Energy Sciences Vision and Grand Challenges

ES.2.1 Background and Context

Plasmas, the fourth state of matter, are ubiquitous in nature, making up 90% of the visible universe. The study of plasmas is key to our basic understanding of the cosmos, including the sun and other stars, the interaction of the solar winds and the earth's magnetic field, and the cores of large planets. Plasmas are important for practical applications in modern civilization, notably in the achievement of plasma processing for semiconductors and new materials, lighting, biomedical applications, and — as a scientific grand challenge for the benefit of future civilizations — in the realization of fusion energy. A closely related topic is the interaction of burning plasmas with ordinary matter, including the effects of fusion products on the first wall and structural materials in a fusion reactor. The physics of plasmas is an application of nonequilibrium statistical mechanics governed by the Boltzmann equation (in three spatial dimensions, three velocity space dimensions, and time) coupled to Maxwell's equations for the evolution of magnetic and electric fields. In many practical problems, these also couple to equations for atomic, molecular, and nuclear processes. The theoretical challenges arise from the intrinsic nonlinearity, high dimensionality, and extreme range of mutually interacting temporal and spatial scales that interact with each other in a typical problem. In a strong magnetic field, as in the case with fusion energy applications, extreme anisotropy and sensitivity to geometric details are also present.

ES.2.2 The Role of Advanced Computing

Advanced computing has been an integral part of plasma science and fusion research from its beginnings and is motivated by the recognition that numerical methods would be a necessary complement to analytic theories. As a result, plasma researchers have been in the forefront of scientific computing for the last four decades. This expertise has been a particular strength of the U.S. research program, and advanced computing is a key element in the DOE-FES strategic plan (DOE-SC 2015). In developing the strategy, the program has enumerated a well-defined set of computational challenges that will require exascale computing. It is important to note that historically, improvements in algorithms and software have been as important as speedup in hardware in setting the pace of progress. At the same time, we note the requirement for expanded capabilities on computational platforms at all scales in support of program objectives. Speedy advanced computing for experimental operation planning and next-shot improvement has also been affecting fusion experiments and code validation. The ever-increasing speed and size of experimental data generation from a variety of physics diagnostic measurements have also benefitted from advanced computing. The fusion community has been developing synthetic diagnostics tools for high-fidelity comparison between experimental data and advanced computing data. As ITER will produce larger-than-ever amounts of data in the future exascale computing era, with research and operation participants collaborating all over the world, an advanced remote data management workflow for computational pre-operation planning, run-time experimental steering, and real-time scientific discovery and code validation will be an essential feature.

ES.3 Priority Research Topics and Computing Needs

Review participants focused on the following areas in which advancement in the transformative opportunities can be achieved through key and sustained efforts in computation, simulation, and advanced tool design.

ES.3.1 Fusion Energy Science

ES.3.1.1 Turbulence and Transport in a Fusion Reactor

Plasma turbulence and transport determine the viability of a fusion reactor. To obtain a high enough temperature and particle density for efficient fusion reactions, the plasma is self-heated

by the fusion-born energetic alpha particles with assistance from external heating, such as radio frequency (RF) or neutral beam heating. If the plasma energy and particles are lost too quickly, which is mostly determined by plasma turbulence, the fusion burn cannot occur or be sustained. If the plasma is not confined at the edge, then core confinement is not achieved. Advances in high performance computing (HPC) have enabled enormous progress in direct numerical simulation of fusion plasma turbulence. Today, simulations of turbulence and transport routinely compare well with experimental measurement — enabling predictive capability to move forward. These simulations involve complex magnetic geometries, multiple temporal and spatial scales, and multiple physical effects. Comprehensive numerical simulations solve the nonlinear gyrokinetic equations that are rigorously derived from first principles using well-agreed-upon ordering parameters. Multiphysics in the complicated nonthermal edge plasma are not scale-separable and require extreme-scale computing. These simulations are already fully utilizing the capabilities of leadership-class supercomputers (up to the maximal Titan and Mira cores). Considering that the world's first burning plasma, ITER, currently under construction in southern France, will have a plasma volume of about an order of magnitude larger than today's largest existing tokamaks, well-resolved full-torus gyrokinetic simulations for ITER will definitely require exascale-class supercomputers. The ultimate goal is the reliable prediction of the radial temperature and density profiles in ITER and future fusion reactors in order to determine and optimize their fusion performance.

ES.3.1.2 Energetic Particles and Magnetohydrodynamic Instabilities in a Fusion Reactor

Confinement of energetic particles (EPs) is a critical issue for burning plasma experiments because the ignition in ITER relies on the self-heating by energetic fusion products (α -particles). Energetic particles can readily excite mesoscale instabilities that drive large EP transport, which, in turn, can degrade overall plasma confinement and threaten the machine's integrity. Because EPs constitute a significant fraction of the plasma energy density in ITER, energetic particles will also strongly influence the microturbulence responsible for turbulent transport and macroscopic magnetohydrodynamic (MHD) instabilities potentially leading to disruptions. In fact, plasma confinement properties in the ignition regime of self-heating by α -particles is one of the most uncertain issues when extrapolating from existing fusion devices to ITER. Predictive capability requires exascale-level, integrated first-principles simulation of nonlinear interactions of multiple kinetic-MHD processes. For example, the excitation, dynamics, and control of the neoclassical tearing mode (NTM), the most likely instability leading to disruption in a tokamak, depend on nonlinear interaction of MHD instability, microturbulence, collisional (neoclassical) transport, energetic particle effects, and RF waves.

The MHD studies may have different computational requirements than exascale can provide and can be accommodated by mid-range computing: Fusion energy science will continue to require many mid-scale computations in the exascale era. While exascale computing will enable a small number of “heroic” runs of the models in our community that are closest to first-principles, there are other classes of computations (e.g., those that take fluid moments of the kinetic distribution function) that play an essential role in interpreting and guiding experiments. However, the large timescale separation assumption between stability and transport phenomena necessitates very long time calculations to study the onset and eventual saturation or other termination of a global event, such as a disruption. Even though the codes use advanced and fully implicit time-stepping, a single initial value simulation can require hundreds of wall-clock hours and millions of CPU hours to perform a realistic (experimentally relevant) simulation. These jobs normally need to be carried out as a series of restarts, each taking 10–20 wall-clock hours. Because the National Energy Research Scientific Computing Center (NERSC) typically has hundreds of jobs waiting in the queue at any given time, each new restart has to get in line and wait again for its time to run, which

alone can take typically several days or longer. This situation leads to periods of months to run a single job to completion. In addition, addressing a scientific or engineering objective often requires scanning parameters and hence many mid-scale calculations. Applying modern UQ techniques also requires performing many simulations of the same event where parameters are systematically varied. It would greatly improve productivity if more hardware were available for running many computations at the 10,000- to 50,000-processor level — an approach that is often called capacity computing. An additional consideration is that many codes (such as the implicit MHD codes) run more efficiently with larger amounts of memory per node than are available on the leadership systems. Dedicated capacity systems with larger amounts of memory may be very cost effective for these codes (which will also help free up the leadership machines for problems that really need their capability).

ES.3.1.3 RF Heating in a Fusion Reactor

The success of next-generation fusion devices and subsequent commercial power plants will rely critically on the robust and efficient application of high-power RF systems in the ion cyclotron, electron cyclotron, and lower hybrid ranges of frequencies. Achieving these goals will depend on the development of a predictive simulation capability of sufficient fidelity for how RF waves interact with the tenuous edge plasma of a fusion device where they are subject to a variety of deleterious interactions, including the formation of RF sheaths at plasma-material surfaces, parametric decay instability, and scattering from turbulence. Once RF power has been successfully coupled to the plasma core, the RF heated plasma species in the core can form nonthermal distributions of ions or electrons, thereby heating specific species or enabling control of the current and pressure profiles, and RF waves can also interact with energetic populations of fusion alpha-particles. A predictive simulation capability is therefore needed to study the stability of these wave-particle interactions and to understand how these energetic populations can be used most effectively to heat and control a burning plasma. A successful collaboration with the Advanced Scientific Computing Research (ASCR) community would make it possible to implement simulation models for coupled antenna-to-core wave-particle interactions on emerging exascale architectures that fully account for the multiscale and multiphysics nature of RF heating and current drive.

ES.3.1.4 Whole-Device Fusion Modeling

A special property of the hot fusion plasma in a toroidal geometry is that there are several multiphysics processes working together, and most of them are scale inseparable and interacting nonlinearly with each other in a self-organized manner at a fundamental physics level. Thus, the fusion reactor plasma must be understood and predicted in a whole-device modeling approach. The best way to simulate the whole-device plasma is to use the 6-dimensional (6-D) Boltzmann or 5-dimensional (5-D) gyrokinetic Boltzmann equation coupled to Maxwell equations for the electromagnetic fields. All of the multiphysics phenomena are included in the 6-D Maxwell-Boltzmann equation system. However, whole-device modeling with a 5-D or 6-D Boltzmann equation cannot be realized until an exascale (or beyond) computational capability is available. For this reason, only kinetic physics addressing individual phenomena — or a combination of only a few multiphysics phenomena — has been studied in today's leadership-class computers, with corresponding scale separation assumptions. An exascale computer can truly enhance the whole-device modeling capability at high fidelity. An alternative method for whole-device modeling is to componentize the multiphysics using reduced transport equations and fluid models with scale separation assumptions. The fidelity of the reduced components can be improved from the knowledge obtained from kinetic simulations. The latter method has different computational requirements from the former in that it requires only a small to mid-size computer with a quick turnaround time, and thus is preferred by experimental modelers for quick experimental data analysis. A larger-scale computer is needed to run many jobs simultaneously.

ES.3.2 Plasma Surface Interactions and Structural Materials

The realization of fusion as a practical, twenty-first-century energy source requires improved knowledge of plasma-material interactions and the materials engineering design of component systems that can survive the incredibly extreme heat and particle flux exposure conditions of a fusion power plant. The traditional trial-and-error approach to developing first-wall material and component solutions for future fusion devices (ITER, DEMO) is becoming prohibitively costly because of the increasing device size, curved toroidal geometry, access restrictions, and complex programmatic priorities. This set of conditions requires changing from an engineering emphasis toward a more fundamental approach, grounded in a multiscale modeling methodology capable of simultaneously attacking the plasma-material interface problems from both a bottom-up and a top-down approach. The dynamic modeling of the kinetic processes occurring at the near-wall layer requires the coupling together of different physical models and codes, namely:

1. A multi-species kinetic model of the plasma sheath/presheath region, handling the evolution of the distribution function of electrons, ions, neutrals, and material impurities from the quasi-neutral region to the first surface layer; the target equations are the Boltzmann-Poisson and the Boltzmann-Maxwell.
2. A kinetic model of the material wall, handling ion-matter interaction and including relevant phenomena such as sputtering, backscattering, and implantation, on a material surface having dynamic composition and evolving morphology; the target equation is the classical multibody problem for a given (known) interaction potential.
3. A proper collision operator accounting for the interaction among species, handling the relevant atomic physics such as ionization, charge exchange, ion and impurity recycling, and more. The target equations are the Fokker-Planck and nonlinear collision operator.

We anticipate that exascale computing will enable the fusion and ASCR community to achieve an integrated and first-principles-based suite of advanced codes to predictively model the boundary plasma and material surface. Such codes will incorporate rigorous treatment of the turbulent transport, along with kinetic and sheath effects in the plasma, and will be efficiently coupled to a multiscale materials modeling framework. The codes will also predict the evolving plasma-facing components' (PFCs') performance, in terms of erosion, PFC lifetime, and tritium inventory, such that the plasma boundary models can provide feedback to the codes modeling the plasma pedestal and the burning plasma core performance.

ES.3.3 Discovery Plasma Science

ES.3.3.1 General Plasma Science

In this review, we focused on magnetic reconnection and turbulence and their role in particle acceleration and heating in space and astrophysical plasmas. Our understanding of these plasma processes is mature in two dimensions (2D), thanks to advances in theory and the availability of petascale computers; however, much remains to be understood in three dimensions (3D). In high-Lundquist-number (S) plasmas, the computational cost to resolve the reconnection layers and follow the macroscopic evolution on the global Alfvén time increases as $S^{5/2}$ for 3D explicit simulations. For $S \sim 10^6$, these requirements can quickly surpass the capabilities of petascale computers, thereby requiring exascale-level resources. *The priority research directions in this area are: (1) the influence of the electron and ion kinetic scales on the large-scale evolution, (2) reconnection and magnetic island dynamics in 3D geometries, (3) energetic partition and particle acceleration, and (4) relativistic reconnection.* Because turbulence mediates the transport of energy, momentum, and particles through motions spanning many orders of magnitude in scale, the modeling of plasma turbulence is an inherently multiscale problem, formally beyond the reach of even today's most advanced computers and sophisticated algorithms, so exascale computing shows the path forward

to making transformative progress in the field. In addition, the problem of space and astrophysical plasma turbulence is made yet more complex by the fact that, at the typically low densities and high temperatures of these plasmas, the turbulence dynamics are often weakly collisional, requiring the application of kinetic plasma theory to follow the evolution and dissipation of the turbulence. For turbulence research, the key question to answer is: *How does turbulence in a kinetic plasma mediate the conversion of the energy of plasma flows and magnetic fields at large scales to plasma heat, or some other form of particle energization?* Over the next decade, through a coordinated program of spacecraft measurements, theoretical calculations, and nonlinear kinetic numerical simulations, the scientific community is poised to make transformative progress on this problem. Exascale computing will play an essential role in this research effort, enabling direct numerical simulations of the high-dimensional, nonlinear turbulent dynamics.

ES.3.3.2 High-Energy-Density Laboratory Plasmas

High-energy-density laboratory plasmas (HEDLPs) are extreme states of matter characterized by pressures in excess of 1 Megabar. Such systems are routinely created by powerful lasers at many university-scale facilities around the world; they span a wide range of physical phenomena, from microscopic instabilities and laser-driven particle accelerators, to millimeter-scale inertial confinement fusion (ICF) capsule implosions at the National Ignition Facility, to cosmic ray acceleration. The physics of laser-plasma interactions and HEDLP is multiscale, highly nonlinear, and often needs to be described by a kinetic modeling approach. For these reasons, computer modeling of HEDLP experiments requires extreme HPC resources.

Opportunities for HEDLP physics modeling on exascale computer systems arise from several factors, including the ability to (1) increase the problem size to more closely approximate realistic systems than currently possible; (2) increase the grid resolution to improve credibility; (3) run ensembles of runs for error sensitivity/UQ or for providing trends; and (4) reduce turnover time for interactivity with experimental campaigns. To take full advantage of extreme-scale HPC systems, it is essential to have robust I/O tools and *in-situ* analysis/visualization capabilities, as well as the support of high-level languages (e.g., Python) as a front end to number-crunching modules.

HEDLPs hold the promise of leading to breakthrough discoveries in fundamental science, such as discoveries concerning the origin of cosmic rays to applications like inertial confinement fusion and compact X-ray sources for homeland security, provided that exascale computing resources can be leveraged to better understand the multiscale aspects involved in the underlying systems.

ES.3.3.3 Low-Temperature Plasmas

Low-temperature plasmas (LTPs) are partially ionized gases with electron temperatures in the range of 1–10 eV. Atomic and molecular processes play a key role in LTPs, as does interaction with solid state or liquid surfaces, where surface current and charge, particle fluxes, and many geometric effects on fields can play a dominant role. LTPs are involved in about 70% of the steps in the manufacture of the ubiquitous electronics components that drive modern civilization. One of the fastest-growing areas in LTPs is that of biomedical plasmas, which have current applications in surgery, wound healing, and sterilization, with the promise of many future applications yet to be discovered. LTPs are used to modify thin film material for packaging and solar panels, and the ozonizing processes used for water treatment and plasma-based physical vapor deposition coatings, with these markets together amounting to tens of billions of dollars annually. Driven low-temperature plasmas are typically in a strongly nonequilibrium state. To obtain a high-fidelity understanding of these nonequilibrium plasmas, large-scale computational research is key, even though few researchers are utilizing these resources for LTP study at the present time. Exascale resources will provide a capacity improvement of many orders of magnitude. The capability to provide high confidence models will transform the applied use of LTPs in industry, where

capital investments and subsequent business success depend upon finding the correct answer and increasingly on knowing the error bars. Exascale-enabled verification, validation, and UQ techniques are one such game changer. Elimination or decreased reliance on ad hoc and simplified physics models is another such disruptor.

ES.3.4 Verification, Validation, and Uncertainty Quantification

A major program element in the FES strategy is to develop “validated predictive models,” so it is important to note that confidence in our models can only be earned through systematic confrontation with experimental data and a sharp focus on careful and quantitative estimates of errors and uncertainties. Fortunately, the disciplines of verification, validation, and uncertainty quantification (VVUQ) are rich areas of research in many technical fields. Our challenge is to find the methodologies and algorithms best suited to our problems; to identify gaps where additional research in applied mathematics and computer science is needed and to apply those techniques to specific codes and simulations; and to secure large-scale computational resources for UQ. Overall, these efforts will have a significant impact on future research directions and computational requirements. A key question, and perhaps the most important source of uncertainty in our domain, is the sensitivity of the models to assumptions and inputs used for any particular problem. Challenges are particularly acute for multiphysics integration, which presents mathematical obstacles, and for multiscale integration, which drives a need for large numbers of production runs of computationally expensive codes. The priority research directions compiled in the body of this report summarize the challenges, enabling us to recommend potential approaches to address those challenges. Notable among these are the need for improved methodologies for code verification, especially for coupled/integrated physics models and the extension of existing intrusive and nonintrusive methods for uncertainty quantification and sensitivity analysis to our particular codes.

ES.4 Path Forward










The support and development of our evolving computing ecosystem relies on continued collaboration between the FES and ASCR communities. Rooted in the discussions about the FES vision, research directions, and computing needs, four categories grew out of the review: methods development, computational environment, data, and communication and community involvement.

Regarding **methods development**, the advancing complexity of computer hardware requires FES researchers to have more scalable, performant algorithms and applications that are capable of efficient execution on future computing architectures fielded by ASCR facilities. Meeting participants discussed those computing ecosystem aspects that will accelerate or impede their progress in the next 5–10 years. Participants named application codes and verification and validation techniques, as well as models and algorithms, as key factors requiring significant methods development activity, as well as additional representative methods identified in Section 3 (and listed in Section 4).

Regarding the **computational environment**, requirements for the access, scheduling, and software ecosystem identify an evolving use-model. The “traditional” HPC model, defined as running a large simulation that generates data that are then post processed, is no longer the only primary use-model for many FES projects. Emerging demands, such as for complex workflows and near-real-time computing, are changing the landscape.

The scale of **data** generated from FES simulations and the requirements needed for verification and validation have created an opportunity and a challenge. ASCR and FES facilities must create more data-centric environments with highly effective data analytics tools for their users. Development of such environments and tools will require expertise from domain scientists, data scientists, and applied mathematicians. Continued collaboration will be required to assess proper deployment of the environments as computing resources evolve.

Activities related to **communication and community involvement** are ongoing today in multiple institutions; however, efforts to connect them to the larger science community have been attempted on an “ad hoc” basis to date. ASCR facilities can explore new or improved communication channels and activities. In addition, experience has shown some of the best impact from strong collaborations.



APPENDIX D

Executive Summary from the High Energy Physics Exascale Requirements Review Report

This page is intentionally left blank.

EXECUTIVE SUMMARY

Abstract

The U.S. Department of Energy (DOE) Office of Science (SC) Offices of High Energy Physics (HEP) and Advanced Scientific Computing Research (ASCR) convened a programmatic Exascale Requirements Review on June 10–12, 2015, in Bethesda, Maryland. This report summarizes the findings, results, and recommendations derived from that meeting. The high-level findings and observations are as follows.

- Larger, more capable computing and data facilities are needed to support HEP science goals in all three frontiers: Energy, Intensity, and Cosmic. The expected scale of the demand at the 2025 timescale is at least two orders of magnitude — and in some cases greater — than that available currently.
- The growth rate of data produced by simulations is overwhelming the current ability of both facilities and researchers to store and analyze it. Additional resources and new techniques for data analysis are urgently needed.
- Data rates and volumes from experimental facilities are also straining the current HEP infrastructure in its ability to store and analyze large and complex data volumes. Appropriately configured leadership-class facilities can play a transformational role in enabling scientific discovery from these datasets.
- A close integration of high-performance computing (HPC) simulation and data analysis will greatly aid in interpreting the results of HEP experiments. Such an integration will minimize data movement and facilitate interdependent workflows.
- Long-range planning between HEP and ASCR will be required to meet HEP's research needs. To best use ASCR HPC resources, the experimental HEP program needs (1) an established, long-term plan for access to ASCR computational and data resources, (2) the ability to map workflows to HPC resources, (3) the ability for ASCR facilities to accommodate workflows run by collaborations potentially comprising thousands of individual members, (4) to transition codes to the next-generation HPC platforms that will be available at ASCR facilities, (5) to build up and train a workforce capable of developing and using simulations and analysis to support HEP scientific research on next-generation systems.

ES.1 High Energy Physics: Vision and Grand Challenges

High energy physics is entering a challenging and exciting period over the next decade with new insights into many of the fundamental mysteries of the Universe moving tantalizingly within reach. The discovery of a Higgs boson in 2012 was just the first of many anticipated discoveries as new Large Hadron Collider (LHC) runs are performed at greater energy and luminosity, giving rise to myriad questions: Are there many Higgs? Is supersymmetry correct? Are there new fundamental forces of nature yet to be discovered? Is there physics beyond the Standard Model, and if so, what form does it take? Next-generation experiments aim to solve the current mysteries of neutrino physics: What is the origin of neutrino masses and how are the masses ordered? Do neutrinos and antineutrinos oscillate differently? Are there additional neutrino types or interactions? Are neutrinos their own antiparticles? As cosmological surveys and dark matter experiments come online, a flood of new data will become available to help answer fundamental questions: What is dark matter? What is dark energy? What is the origin of primordial fluctuations?

The high energy physics community — both theoretical and experimental researchers — is actively seeking answers to these questions, and there is a strong anticipation and sense of excitement that we are on the threshold of paradigm-shifting discoveries. Success in these endeavors, however,

requires many new tools of discovery. Forefront among them, as in so many scientific fields today, is a need to perform massive computations on an unprecedented scale and to collect, store, and analyze complex datasets at never-before-seen rates and volumes.










ES.2 Computing Needs/Requirements

The HEP community has historically leveraged HPC capabilities primarily to address theory and modeling tasks. Energy Frontier-related experimental efforts are now also beginning to employ HPC systems. The estimated computational needs for these experiments far exceed the expected HEP investment in computational hardware that will be required to execute the envisioned HEP science program by 2025. Thus, partnering with ASCR is essential to the scientific success of the HEP community.

The 2025 timeline will define a phase transition in terms of HEP data flows. Three new facilities will come online and the landscape will change dramatically. The high-luminosity LHC will be operating, producing data samples at unprecedented scales. Fermi National Accelerator Laboratory's (Fermilab's) flagship Deep Underground Neutrino Experiment (DUNE) will be operating, as will the next-generation dark energy survey experiment, the Large Synoptic Survey Telescope (LSST). Data sizes produced each year could be 200 times greater than what is being produced by today's operating experiments. In addition, these new experiments will each require a simulation program that dwarfs what we are doing today in order to take advantage of the expected improvement in statistical precision.

ES.3 Path Forward

One of the primary goals of computing within HEP will be to transition to HPC capability when appropriate. This will require success at a number of levels. Critical components of the HEP code base will have to be refactored to take advantage of HPC architectures. Equally importantly, the HPC environment will need to adapt to HEP use cases, including the adoption of "edge" services, networking modifications and optimizations, available storage cache at the facilities to optimize data transfer, and associated workflow management and diagnostic tools. Taking these steps will be essential to addressing escalating computing and data demands in the HEP community and, more precisely, to informing HEP scientific programs about computing opportunities and challenges in the future exascale environment.



APPENDIX E

Executive Summary from the Nuclear Physics Exascale Requirements Review Report

This page is intentionally left blank.

EXECUTIVE SUMMARY

Abstract

Imagine being able to predict — with unprecedented accuracy and precision — the structure of the proton and neutron, and the forces between them, directly from the dynamics of quarks and gluons, and then using this information in calculations of the structure and reactions of atomic nuclei and of the properties of dense neutron stars (NSs). Also imagine discovering new and exotic states of matter, and new laws of nature, by being able to collect more experimental data than we dream possible today, analyzing it in real time to feed back into an experiment, and curating the data with full tracking capabilities and with fully distributed data mining capabilities. Making this vision a reality would improve basic scientific understanding, enabling us to precisely calculate, for example, the spectrum of gravity waves emitted during NS coalescence, and would have important societal applications in nuclear energy research, stockpile stewardship, and other areas. This review presents the components and characteristics of the exascale computing ecosystems necessary to realize this vision.

Nuclear physics research and its applications, more than many other areas of science, rely heavily on large-scale, high-performance computing (HPC). HPC is integral to (1) the design and optimization of an extensive and vibrant experimental program, (2) acquisition and handling of large volumes of experimental data, and (3) large-scale simulations of emergent complex systems — from the subatomic to the cosmological. Dramatic progress has already been made in enhancing our understanding of strongly-interacting matter across many areas of physics: from the quark and gluon structure of the proton and neutron, the building blocks of atomic nuclei; to relating the hot dense plasma in the early universe to the near-perfect fluid produced at the Relativistic Heavy Ion Collider (RHIC); to the structure of rare isotopes and the reactions required to form all the elements of the universe; to the creation and properties of the dense cold matter formed in supernovae and NSs.

In addition to pursuing new analytical and experimental techniques and algorithms and more complete physical models at higher resolution, the following broadly grouped areas relevant to the U.S. Department of Energy (DOE) Office of Advanced Scientific Computing Research (ASCR) and the National Science Foundation (NSF) would directly affect the mission need of the DOE Office of Science (SC) Nuclear Physics (NP) program.

- **Exascale capability computing and associated capacity computing will revolutionize our understanding of nuclear physics and nuclear applications.**
- **Closely tied to the needs for exascale hardware are the requirements for new software (codes, algorithms, and workflows) appropriate for the exascale ecosystem, which can be developed through collaborations among ASCR and NSF mathematicians and computer scientists and NP scientists.**
- **There is a need to read, write, manage, analyze, curate, and track data of a complexity and scale never before encountered. This need is important in all areas of nuclear physics.**
- **Growing and sustaining a workforce to carry nuclear physics through the exascale era is vital to meeting the NP mission; workforce development should include enhanced collaboration between NP, ASCR, and NSF and, in particular, new positions at the ASCR/NP interface.**

ES.1 Summary and Key Findings

The mission of the NP program is to discover, explore, and understand all forms of nuclear strongly interacting matter. The fundamental degrees of freedom in nuclear physics are quarks and gluons, whose dynamics create diverse emergent phenomena and nuclear matter over an enormous range of length scales and energies. HPC plays an essential role in understanding these emergent systems and in establishing predictive capabilities in the associated energy regimes. In this review, NP scientists focused on the scientific challenges in the areas of Nuclear Astrophysics, Experiment and Data, Nuclear Structure and Reactions, Cold Quantum Chromodynamics (QCD), and Hot QCD. These areas pertain to understanding the high-temperature quark-gluon plasma (QGP) in the early universe; the collective QCD phenomena that form nucleons, describe their interactions, and provide laboratories to probe the fundamental laws of nature; the structure and reactions of a diverse range of atomic nuclei; the evolution of the dense matter formed in supernovae and NSs; and environments that create the elements in the universe. Scientists in each NP area have worked extensively to maximally exploit the leading computational and software capabilities to advance nuclear science. The computational, data, software, and workforce needs were evaluated in terms of their ability to meet the mission of the NP program. Many of the requirements identified in this report have their precursors in previous reports, including the original *Forefront Questions in Nuclear Science and the Role of Computing at the Extreme Scale* meeting (DOE 2009), the National Academies decadal study report (National Academies Press 2013), and the 2015 *Long Range Plan for Nuclear Science* (Long-Range Plan) (DOE and NSF 2015).

Wide availability of exascale computing is key to critical progress in nuclear physics across all areas. Software, algorithm, and workflow development are required to allow NP to take advantage of exascale computing architectures.

- Nuclear scientists need sustainable and high-performance software tools, programming models, and applications to enable effective use of exascale systems.
- New algorithms are needed to enable codes to run efficiently on upcoming HPC architectures, allowing scientists to model larger systems at higher resolution with greater accuracy.
- New applications are required to perform the suites of simulations necessary to explore new parameter regimes and to quantify uncertainties.
- New hardware and software tools are needed to analyze, track, and manage data produced by experiments and in simulations, including developments in databases, and to move data efficiently between sites for appropriate analysis.


The NP community is convinced that advanced computing, mathematics, and computer science can have a transformational impact on nuclear physics.

The growing needs of NP experiments, and the data-related needs of the community in general, have been identified and characterized. The data analysis, storage, archiving, curation, movement, and tracking requirements in both the experimental and theoretical components of the NP program are increasing very rapidly. Further, the experimental program would benefit from real-time access to the advanced computing capabilities of ASCR and NSF.

An exascale ecosystem for nuclear physics requires a well-trained, diverse, and capable workforce to achieve the vision outlined in this review. Such a workforce is crucial to meeting the scientific challenges that face nuclear physics, and it should include applied mathematicians, computer scientists, statisticians, and computational nuclear physicists. Particularly critical for the long-term success of this program is the growth in the number of permanent positions for scientists working at the interface of applied math and computer science with NP, both at national laboratories

and universities. Successful and long-term career paths for such scientists are only now being developed, and further advances require a strong coordination and collaboration between ASCR, NSF, and NP.

Participants in the NP Exascale Requirements Review were asked to identify the grand challenges facing nuclear physics, determine the priority research topics and computing needs required to successfully meet these grand challenges, and outline the most important steps in the path forward to reach the goals.



Results of simulating a Type-1a supernova from merging white dwarfs, each with 1.06 solar mass. (Moll et al. 2014)

ES.2 Nuclear Physics Vision and Grand Challenges

Nuclear physics is a diverse field with strong internal connections and with external relations to many areas of science — from the hot dense matter of the early universe and the properties of protons at the smallest length scales, to the properties of nuclei and the reactions that form all the elements, to the dense matter and violent explosions of NS mergers and supernovae. It reaches from the most fundamental physics, including searches for physics beyond the Standard Model and the properties of neutrinos, to critical societal problems including the production of energy through nuclear fission and fusion and applications in nuclear medicine. The overarching questions addressed by the nuclear science community have been succinctly formulated in the National Academies nuclear physics decadal study report (National Academies Press 2013):

1. How did visible matter come into being and how does it evolve?
2. How does subatomic matter organize itself and what phenomena emerge?
3. Are the fundamental interactions that are basic to the structure of matter fully understood?
4. How can the knowledge and technical progress provided by nuclear physics best be used to benefit society?

Computational nuclear physics plays a key role in answering each of these questions. It addresses, both through large-scale simulations of astrophysical objects and through studies of nuclei and their reactions, the questions of how all the elements came into being. Computational nuclear physics is required to interpret experimental data and thus return maximum results from DOE investments in experimental facilities, and to translate these results into new knowledge about subjects as diverse as the quark-gluon liquid created at the RHIC at Brookhaven National Laboratory (BNL), the quark and gluon structure of protons, the interaction between nucleons arising from these structures, and the properties of nuclei and nucleonic matter arising from the interaction between nucleons.

For all these reasons, computational nuclear physics was highlighted in all areas in the community process leading up to the 2015 Long-Range Plan, which states: “We recommend new investments in computational nuclear theory that exploit the U.S. leadership in high-performance computing” (DOE and NSF 2015).

Programs such as Scientific Discovery through Advanced Computing (SciDAC) and the Exascale Computing Project (ECP) are essential to these new transformational opportunities in nuclear science. Revolutionary advances in computing, mathematics, algorithms, and data will dramatically alter the field of nuclear science. DOE SC envisions reaching the exascale level of computing within the next decade (by 2026). With simultaneous advances in applied mathematics, computer science, software and data, and NP itself, we can expect a transformation of nuclear science with dramatic impact in all areas, including experiment and applications.

ES.3 Priority Research Directions

The NP Exascale Requirements Review participants focused on five main areas for which exascale resources are required to achieve the goals of the nuclear physics community. There are common elements in the needs in these five areas, but there are also requirements that are unique to each.

ES.3.1 Nuclear Astrophysics

Nuclear astrophysics research offers exciting new opportunities, tying together new laboratory experiments and new astrophysical observations with dramatically increased capabilities in HPC. Cutting-edge efforts in computational nuclear astrophysics are fundamentally multi-physics in nature, and large-scale numerical simulations are at the heart of this science.

Key objectives in the Long-Range Plan (DOE and NSF 2015) include exploring the origin of the elements (e.g., ascertaining the site[s] where the heavy r-process elements, such as uranium and gold, are synthesized), the physics of ultra-dense neutron-rich nuclear matter, and the nature of neutrinos and their interactions with nuclei and in dense nuclear matter.

Selected Highlights of Exascale-enabled Nuclear Astrophysics

- Complete three-dimensional calculations of NS mergers, including realistic treatments of strong interaction and neutrino microphysics and nucleosynthesis including uncertainties.
- Complete three-dimensional plus temporal evolution of core collapse supernovae (CCSNe) with realistic microphysics to determine explosion mechanism, neutrino physics, and nucleosynthesis.
- Calculate flame spreading across the NS surface in x-ray bursts (XRBs), including a robust subgrid flame model and coupling to photon transport.
- Complete more realistic studies of neutrino quantum kinetic equations in both CCSNe and NS mergers to address potential quantum impacts in astrophysical evolution and neutrino physics.

Exascale computational resources are required to meet these objectives through large-scale simulations, demanding expanded access to new architectures. Advanced software and analysis techniques are also required, because the simulations need extensive parameter studies. Unique and challenging new kinds of physics will also be modeled — for example, neutrino flavor evolution with neutrino scattering and the gravitational signals from binary NS mergers. Another example is modeling the electromagnetic signals accompanying binary NS mergers, or the r-process nucleosynthesis of CCSNe — a key part of “multi-messenger” probes of these events. We can expect exciting new developments along many of these fronts in the exascale era.

ES.3.2 Experiment and Data

Experimental nuclear physics computing and data needs were explicitly considered in this review, a new arena of intersection between ASCR and NP. Nuclear physics experiments are driven by precision to access the multi-dimensional and multi-channel problem space, requiring beam intensity, polarization, and careful treatment of backgrounds and systematics. DOE NP experimental efforts at major facilities include nuclear studies at the Argonne Tandem Linac

Accelerator System (ATLAS) at Argonne National Laboratory, electron scattering at Thomas Jefferson National Accelerator Facility (JLab), and proton and heavy-ion collisions at the RHIC at BNL. Needs for increased computational resources are being driven by the Facility for Rare Isotope Beams (FRIB), currently under construction at Michigan State University (MSU); by an anticipated Electron Ion Collider (EIC); and by fundamental physics, including the Fundamental Neutron Physics Beam Line at the Spallation Neutron Source at Oak Ridge National Laboratory and neutrinoless double-beta ($0\nu\beta\beta$) decay experiments, and research at smaller university and laboratory facilities. The ever-increasing data and analysis needs of NP experiments can be addressed, in part, through collaboration with ASCR.

Selected Highlights of Exascale-enabled Experiment and Data


- Implementation of end-to-end software frameworks that employ best practices for integrating data transport, messaging, compute services, and visualization components.
- Development of new software frameworks for fine-grained parallelism for detector simulations and data analysis.
- Creation of end-to-end software frameworks that operate at different data-processing stages to efficiently use available computing resources and adapt to real-time environments.

The major experimental needs outlined in this report are data streaming and near-line validation,⁴ detector simulation and data analysis on HPC resources, the combination of work flows, data management and network infrastructure, and accelerator simulation. The capability and capacity computing needs in this area are modest compared with those in other areas, but the requirements for both immediate (hot) and long-term (cold) data are significant. Also there is a significant need to promote a close association between the experimental nuclear physics facilities and HPC resources, driving both hardware requirements in areas like networking and data storage and software efforts in data validation, analysis, and visualization. The major new nuclear physics experiments could greatly benefit from an increased connection with the forthcoming exascale ecosystem.

ES.3.3 Nuclear Structure and Reactions

Nuclear structure and reaction research is progressing dramatically as a result of exciting theoretical progress in the nuclear many-body problem and simultaneous experimental developments. The challenges driving this field include the following: determining the sites and processes creating all the elements in the universe, examining general organizational principles and processes governing the physics of nuclei, exploiting nuclei to reveal the fundamental symmetries of nature, and determining the scientific and practical uses of nuclei for society in general and for related fields of science.

⁴ Near-line validation is the extraction of physics-based observables from stored data on a timescale typically within a few minutes of collection to validate correct experimental operation, enable optimization, and provide input for online data analysis.



In partnership with applied mathematics and computer science, modern nuclear structure theory strives to estimate uncertainties in predictions and assess extrapolations. Shown is the landscape of bound nuclei as a function of proton and neutron number predicted in nuclear Density Functional Theory (Erler et al. 2012). Mean drip lines, where the nuclear binding ends, and their systematic uncertainties (red) were obtained by averaging the results of different theoretical models. The inset shows the posterior distributions for the parameter set of the energy density functional UNEDF1 (McDonnell et al. 2015); this is an essential element of the statistical uncertainty estimate.

The quantum many-body methods used in nuclear physics all have analogues in other fields dealing with complex systems (e.g., condensed matter and atomic and molecular physics), but the unique features of nuclear interactions have particular requirements in the present and future eras of computational science. The requirements are, in most cases, similar to those in other areas of nuclear physics: access to significant resources on the emerging architectures and data storage and analysis needs. Some algorithms — for example, the large-scale linear algebra used in configuration interaction methods — have particular needs for large memory per node as well. Also, here and in other areas, there are significant workforce needs, particularly for permanent positions at the interface of nuclear science, computer science, and mathematics.

Selected Highlights of Exascale-enabled Nuclear Structure and Reactions

- Complete ab initio calculations of neutron-rich nuclei and electromagnetic transitions, including those to be studied at FRIB and beyond, with quantified uncertainties.
- Identify global nuclear physics inputs, with theoretical uncertainties, for r-process nuclei and develop the universal energy density function for all nuclei.
- Make quantified predictions of the alpha-capture reactions that synthesize carbon, oxygen, and heavier elements in stars.
- Calculate $\beta\beta$ decay matrix elements of ^{76}Ge , ^{136}Xe , and ^{130}Te with uncertainties of less than 25% and determine ab initio electroweak responses of ^{40}Ar and nucleonic matter in astrophysical environments.

The questions being addressed in this field are very closely tied to large-scale experiments and facilities in nuclear science and beyond. In some cases, these connections are obvious, like those in the lower-energy nuclear facilities such as ATLAS at Argonne, the National Superconducting Cyclotron Laboratory (NSCL) and FRIB at MSU, and other university laboratories. The science of nuclei also benefits greatly from connections to JLab and important accelerator neutrino experiments like the Deep Underground Neutrino Experiment (DUNE). Another major new

initiative of nuclear science is the search for $0\nu\beta\beta$ decay, the observation of which would provide conclusive proof of the violation of lepton number and the Majorana nature of the neutrino.

An improved theoretical understanding is essential in optimizing the physics resulting from the greatly enhanced experimental program. Nuclear many-body theory holds great promise for dramatically advancing our understanding and exploitation of nuclei at laboratories for fundamental physics research through exascale computing.

ES.3.4 Cold Quantum Chromodynamics

Protons, neutrons, and their interactions arise from quarks and gluons and their dynamics, as dictated by QCD. A precise understanding of the proton, the neutron, and the inter-nucleon forces is critical to many aspects of subatomic science. The three-dimensional structure of the nucleon and its excitations is being explored with precision by the experimental programs at JLab and RHIC. These programs will be extended to studies of the gluonic structure of nucleons and nuclei with construction of the planned EIC. Understanding/refining the forces between the neutrons and protons is important to building a more comprehensive picture of nuclei and their reactions; such research is a focus of the FRIB facility.

Selected Highlights of Exascale-enabled Cold QCD

- Complete precision calculations of nucleon couplings and associated form factors to all quark and gluon interactions. Obtain a detailed understanding of the gluonic contributions to the mass and spin of the proton, nuclei, and exotic states.
- Complete a quantitative exploration of the strong-interaction glue, and excitations, that binds quarks into hadrons, to support and complement an EIC experimental program.
- Perform precision calculations of light nuclei and chiral nuclear forces at the physical quark masses, including electromagnetism, to support and complement a FRIB experimental program.
- Complete precision calculations of the two- and three-nucleon short-distance interactions contributing to $0\nu\beta\beta$ -decay of nuclei and of interactions contributing to a neutron electric dipole moment (nEDM) and nuclear EDMs in support of an enhanced program in Fundamental Symmetries.

The prime areas of research in cold QCD are hadron structure and spectroscopy, nuclear interactions and structure arising from QCD, and applications of QCD to fundamental symmetries in nucleons and nuclei.

Each of these areas has close connections to the NP experimental program. HPC is essential in this effort. Indeed, the drive toward precision theoretical understanding of cold QCD systems has traditionally spurred the development of HPC. The computational requirements are large at both the capability and capacity level. Capability computing is required to generate the gauge configurations in large spacetime volumes that can later be analyzed for specific observables using capacity computing. Hot and cold data requirements are important, and workforce needs are significant for simultaneously advancing on these multiple fronts.

The future of lattice QCD at the exascale is very exciting. Lattice QCD will allow precise calculations of the quark and gluon structure of the proton and neutron and their excitations, and enable nucleons to be used as precise laboratories for studies of fundamental symmetries through measurements such as the nEDM. Advances into multi-nucleon systems at physical parameters, from which the nuclear forces can be refined through the use of effective field theories and modern phenomenological interactions, are important. They offer the promise of directly connecting QCD to nuclei, thereby providing a rigorous underpinning for calculations throughout the periodic table of elements and of dense hadronic matter.

ES.3.5 Hot Quantum Chromodynamics

The matter in the early universe was initially very hot and dense. As it cooled to temperatures in the vicinity of 1.5×10^{12} K, it underwent a transition from a state composed of unbound quarks and gluons — the QGP — to a gas made of hadrons, including protons, neutrons, and pions. The RHIC facility at BNL and the Large Hadron Collider (LHC) at CERN (European Organization for Nuclear Research) reverse this transition, colliding nuclei at very high energies to probe the properties of the QGP. Lattice QCD calculations use first-principles methods to study the properties of the QGP at these high temperatures. The hydrodynamics of the collisions probe the transport properties of the plasma, and they have revealed that the QGP is, in fact, nearly a perfect fluid, with an extremely low shear viscosity.

Studies of QCD at high temperatures address questions related to the nature of the equation of state (EoS) and transport properties of QCD at finite temperature and density, how the early-time dynamics of the nuclear collisions approach equilibrium conditions, and the nature of the elementary excitations in the QGP. They probe the essence of matter at the extremes of temperature and density.

Selected Highlights of Exascale-enabled Hot QCD










- Complete lattice QCD calculations of thermodynamic quantities related to RHIC and LHC heavy-ion experimental programs.
- Complete fully realistic calculations of meson spectral functions and transport coefficients.
- Achieve precise extraction of QGP properties at low and high baryon density, from first principles and from the large body of available high-precision experimental data.
- Complete fully realistic classical statistical simulations of early-time dynamics and equilibration with dynamical chiral fermions.

Hot QCD efforts span an important range of equilibrium (lattice QCD) and non-equilibrium (early-time dynamics and thermalization) efforts, leading the way to a much more complete understanding of all the properties of QCD in the hot and dense regime. The computational requirements of hot QCD emphasize the need for capability computing resources from present day through the exascale era. The data requirements are also significant, particularly for lattice QCD studies aimed at locating the QCD critical point at high temperature and relatively modest baryon density.

ES.4 Path Forward

The scientific challenges facing the U.S. nuclear physics research program are unprecedented in scope and impact. These challenges require an extensive and diverse exascale high-performance computing and data ecosystem (HPCDE). The diversity of the computing requirements emerges from the need to be able to solve QCD in both the hadronic and low-energy regimes and at extremes of temperature and density, calculate the structure and reactions of atomic nuclei, perform large-scale nuclear astrophysics simulations, and acquire and analyze large data sets generated in the nation's laboratories with real-time data-processing capabilities.

The required exascale environment includes capability computing resources to be delivered with heterogeneous architectures and commensurate hot- and cold-data storage capabilities; enhanced infrastructure for data-handling capabilities; growth of a highly skilled and sustainable workforce (developed in close collaboration with ASCR and NSF); and increased capacity computing resources at universities and laboratories. It would be advantageous for DOE NP and NSF to develop coherent plans for nuclear physics computing in the exascale era that complement and enhance those of ASCR. The integration of this exascale ecosystem into the nuclear physics research program will provide unprecedented predictive capabilities, sparking scientific discoveries across the wide spectrum of nuclear science, including many that we cannot imagine today.



APPENDIX F

Executive Summary
from the
Advanced Scientific Computing
Research Executive Summary
Exascale Requirements
Review Report

This page is intentionally left blank.

EXECUTIVE SUMMARY

Abstract

The widespread use of computing in the American economy would not be possible without a thoughtful, exploratory research and development (R&D) community pushing the performance edge of operating systems, computer languages, and software libraries. These are the tools and building blocks — the hammers, chisels, bricks, and mortar — of the smartphone, the cloud, and the computing services on which we rely. Engineers and scientists need ever-more specialized computing tools to discover new material properties for manufacturing, make energy generation safer and more efficient, and provide insight into the fundamentals of the universe, for example. The research division of the U.S. Department of Energy's (DOE's) Office of Advanced Scientific Computing and Research (ASCR Research) ensures that these tools and building blocks are being developed and honed to meet the extreme needs of modern science.

ES.1 Summary and Key Findings

The ASCR Research Requirements Review brought together scientific and computational experts with interests in ASCR-supported computer science, applied mathematics, and next-generation network research programs. The objectives of the review were to (1) assess the ability of the ASCR high-performance computing (HPC) facilities to support the needs of these researchers in an exascale computing environment; and (2) identify areas where that support could be enhanced, potentially opening new opportunity areas for computational research. Given the broad spectrum of ASCR Research activities, the organizing committee surveyed potential attendees to identify discussion topics. On the basis of the survey results and organizing committee discussions, the breakout sessions were structured to cover a diverse set of HPC technology areas: HPC architectures, data management/visualization/analysis/storage, high-performance distributed computing, software development, systems software, deployment, and operations.

The following broadly grouped areas directly affect the mission need of the DOE ASCR Research program to pursue exploratory R&D in programming systems, operating systems, architectures, data management systems, and performance-edge hardware and software. Section ES.3 (ASCR Computing Needs) describes the identified areas of need in more detail.

- In contrast to the domain science areas, the ASCR Research focus on computer science, applied mathematics, and next-generation networking means that current and future HPC computing systems themselves are frequently the subject of the research, and thus, researchers need *frequent, comprehensive access* to these systems.
- Close collaboration and effective communication between HPC facilities and ASCR researchers are critical to facilitating researcher access to HPC resources, as well as enabling researchers to *study and test HPC systems in detail*.
- Because HPC facilities are often the venue for ASCR researchers to develop, test, and deploy their tools, these facilities need to provide a computing ecosystem that supports a *development environment for ASCR research efforts* and timely deployment of ASCR research products, as well as an environment for HPC users.

In aggregate, participants in the breakout sessions identified 77 needs, each framed by specific science driver(s). Of these, 52 were categorized as “high-priority.” Because many of the needs identified were similar, the organizing committee combined logically related needs into higher-level areas of need. The intention was to provide a holistic framework to address areas of need, rather than seeking point-by-point solutions for the large number of individual needs.

The HPC environment is a computational ecosystem with diverse, complex, and tightly inter-related components — both for production use of the facilities and for the ASCR-funded computational R&D activities that rely on those facilities. Enhancing that ecosystem benefits all aspects of the ASCR Research program: computer science, applied mathematics, and next-generation networking programs. Thus, our findings are best viewed in the context of enhancing the system as a whole, rather than just its individual components.

ES.2 ASCR Mission and Community

Mission

The mission of the ASCR program is to advance applied mathematics and computer science; deliver the most advanced computational scientific applications in partnership with disciplinary science; advance computing and networking capabilities; and develop future generations of computing hardware and software tools for science, in partnership with the research community, including U.S. industry. The strategy to accomplish this mission has two thrusts: (1) developing and maintaining world-class computing and network facilities for science; and (2) advancing research in applied mathematics, computer science, and advanced networking.

The ASCR Research Division underpins DOE’s world leadership in scientific computation by supporting research in applied mathematics, computer science, high-performance networks, and computational partnerships (SciDAC).

The *Computer Science* program pursues innovative advances in a broad range of topics, including programming systems, system software, architectures, performance and productivity tools, and many others. In particular, the program focuses on effective use of very large-scale computers and networks, many of which contain thousands of multi-core processors with complex interconnections and data movement.

The *Applied Mathematics* program supports mathematical and computational research that facilitates the use of the latest HPC systems to advance our understanding of science and technology. More specifically, this program develops mathematical descriptions, models, methods, and algorithms to describe and understand complex systems, often involving processes that span a wide range of time and/or length scales.

The *Next-Generation Network for Science* program in ASCR conducts R&D activities to support distributed high-end science. It focuses on end-to-end operation of high-performance, high-capacity, and middleware network technologies needed to provide secure access to distributed science facilities, HPC resources, and large-scale scientific collaborations.

Scientific Discovery through Advanced Computing (SciDAC) brings together computational scientists, applied mathematicians, and computer scientists from across application domains and from universities and national laboratories to ensure that scientists are using state-of-the-art technologies to solve their increasingly complex computational and data science challenges.

Community

ASCR researchers are often software developers, as well as software users. They require low-level access to the hardware to study the platform itself, and they may need access to computing systems in a variety of modes that domain scientists — who employ computing resources for large-scale

production runs and require access to many nodes for many hours — may not need. ASCR research explicitly supports the development of software that introduces cutting-edge advances in applied mathematics and computer science, which in turn supports domain science applications. Software tools produced by ASCR computer science and applied math researchers enable physicists, biologists, and other domain scientists to exploit the full power of machines at the ASCR facilities and thereby increase their scientific output.

ASCR researchers' output ranges from a fundamental understanding of hardware systems, software and programming systems, and numerical algorithms to the design and development of fully functioning production software that enables cutting-edge domain science. In addition, because ASCR researchers work so “close to the metal,” there are significant opportunities to realize mutual benefits for the ASCR research and computing facility missions.

For the ASCR community, facilities need to provide support for a true developer ecosystem in addition to supporting HPC *users*. Although facilities have application readiness programs to ensure that applications can make effective use of new hardware when it goes into production, there is currently no similar program that targets programming systems, libraries, or tools.

ES.3 ASCR Computing Needs

Because the *mission and goals* of the ASCR Research community differ significantly from those of other DOE Office of Science researchers, the computing *needs* of ASCR researchers in computer science, networking, and applied mathematics also differ substantially from those in other DOE domain science areas. Computational biologists, chemists, and other domain scientists rely on HPC to perform high-fidelity simulations, analyze observational and experimental data, and connect simulation and data. The computing facilities provide the means by which they perform their research.

For computer scientists, however, the computing facilities are not just a means to an end — they are, in fact, the subject of the research. Similarly, while other Office of Science researchers rely on high-performing networks to carry their data, networking researchers study the detailed qualities and performance of the network itself. Finally, while scientists in many areas of study rely on efficient, scalable algorithms embedded in their simulation codes or available at the facilities, applied mathematicians are at the cutting edge of defining new algorithms to run on new machines. Thus, the computing ecosystem is not just a tool to be used in their research, but often the subject of their research, which presents an entirely different set of computing needs.

For the ASCR Research Exascale Requirements Review, the available computing environment was decomposed into three tiers to clarify the differences in platform readiness and provide a common language to address the different types of systems and their associated needs. We employed the Technology Readiness Level (TRL) scale developed and used by several U.S. Government agencies: the U.S. Department of Defense (DOD), DOE, and National Aeronautics and Space Administration (DOE 2009). These agencies have applied TRLs to a broad range of technologies, including aircraft, electronics, and computing. Section 2 provides further detail about TRLs. For our review, the TRLs were calibrated to several important computing metrics and elements, including components (e.g., central processing unit [CPU], graphics processing unit [GPU], field-programmable gate array [FPGA]), level of integration, scale, software readiness, and others. On the basis of the attendee survey and committee discussions, ASCR Research Exascale Review attendees were asked to categorize their findings and recommendations using the following three categories:

- **Emerging Systems** (TRLs 1–6) — innovative computing technologies that are currently not deployed to the HPC community at scale (e.g., FPGAs, quantum, neuromorphic, and experimental devices like carbon nanotube [CNT]-based integrated circuits).

- **Early Delivery Systems** (TRLs 7–8) — computing technologies scheduled to be deployed in a production system in the near future (e.g., Cori Phase 2, including early science period).
- **Production Systems** (TRL 9) — computing technologies deployed and operated at scale in the HPC community as a production resource (e.g., Titan, Edison, and Mira).

To organize the breakout sessions at the ASCR Research review, we combined (1) the three tiers of the computing environment with (2) the topics identified from the survey responses (HPC architectures, data management/visualization/analysis/storage, high-performance distributed computing, software development, systems software, deployment, and operations). Thus, each session focused on a particular topic or set of topics in the context of one of the three computing environment tiers (emerging, early delivery, and production).

The review found that the following broadly grouped areas relevant to ASCR facilities would directly affect the ASCR Research mission.

- The model of facilities providing compute cycles and researchers using those cycles to perform their research breaks down for ASCR Research. Computer scientists need access to emerging architectures and early delivery systems not just to use the systems, but to study them.
- Researchers benefit from availability of and access to operational data about the systems and the computing infrastructure with which they work. To allow computer scientists and networking researchers to study the systems themselves, they need improved system instrumentation and monitoring, access to system operational and performance data, network infrastructure counter data, usage statistics, running jobs data, data movement logs, and tools that provide performance insights. To enable applied mathematicians to develop new algorithms and implementations, they need to understand in detail how the current algorithms perform on current machines. In some cases (particularly emerging architectures), they also need the ability to change the system in order to study it (e.g., system reconfiguration agility and facilitation of hardware and software configuration changes).
- ASCR researchers need access to systems to perform testing and development; they often need a true development environment in addition to a user facility. This need includes accessibility and availability of development systems and small test bed facilities, end-to-end test and development capabilities up through the application layer, system availability for at-scale testing, routine testing of system software, and vendor participation in testbed facilities.
- Computer scientists, mathematicians, and distributed computing researchers must contend with a wide spectrum of HPC facility and resource access issues to conduct aggressive, relevant research in their areas of expertise. Simplification and uniformity of access to HPC resources across the different DOE-funded computing centers (including standardized authentication technologies) would enhance the R&D environment for those researchers, enabling them to be more efficient and productive.
- Researchers across the ASCR spectrum spoke about the need for improved information sharing about available systems (particularly emerging and early delivery systems) and clear policies regarding who can get access to these systems and how. Such communication can take the form of online documentation, outreach, and staff support. Users expressed frustration about the difficulty in finding relevant information; in some cases, the information was there but hard to find, and in other cases (particularly for emerging and early systems), the information was not available. ASCR researchers would like consolidated information about access and available resources, staff training and training materials for developers, access to system hardware and software information, and improved communication and coordination among developers and system users.










- ASCR researchers, including applied mathematicians, computer scientists, and networking researchers, require a broad spectrum of deployment capabilities to conduct their research effectively. At HPC facilities, such deployment capabilities should include flexible software deployment processes, deployment of next-generation network technologies and data movement tools, software portability to run common application code easily across HPC facilities, easier on-ramp capabilities for new system software, and improved strategies for transitioning successful research artifacts into deployed software.
- The ability to modify a computing environment, either through software/hardware reconfiguration or upgrade, is essential to the research activities necessary to advance computing technology. ASCR researchers will be more successful — both in terms of their R&D efforts and in advancing their technologies to production — with greater software/hardware deployment and reconfiguration agility at HPC facilities.
- ASCR researchers need better data repository capabilities to support the wide spectrum of research activities they conduct at HPC facilities. Their storage needs range from readily accessible results from previous computations to short-term caching capacity for research involving data-intensive workflows.
- To address the increasing importance of data in all its forms, ASCR researchers need facilities for data collection and preservation, modern HPC-enabled data repositories, storage capacity at HPC facilities for data-intensive workflows, and long-term archiving capabilities.

The areas of need listed above are not presented in order of priority — every one of them includes needs identified as “high-priority” by at least four different breakout groups. For that reason, we have chosen not to attempt to prioritize them. We encourage readers of this report to view them as an interrelated set of needs that combine to form the key findings of the report.

ES.4 Path Forward

Because HPC systems are the object of research for ASCR researchers, the compute ecosystem needs identified in the ASCR Research Exascale Review are very different from those identified in other reviews. Primarily, ASCR researchers identified the following as crucial cross-cutting needs: increased access to all phases of hardware, increased agility in hardware and software configuration and deployment, and increased access to test and development hardware and system data. These needs require substantial new resources, new collaboration, and new workforce effort to explore how computer science and applied mathematics research can be better enabled in production and early delivery systems.

This page is intentionally left blank.



APPENDIX G

Sections Cited in This Report
from the Exascale Requirements
Review for Basic Energy Sciences
(BES)

This page is intentionally left blank.

SECTIONS CITED IN THIS REPORT FROM THE EXASCALE REQUIREMENTS REVIEW FOR BASIC ENERGY SCIENCES (BES)

The following sections from the Basic Energy Sciences (BES) Exascale Requirements Review report are cited at least once in this Crosscut report.

[BES ES.1], [BES ES Abstract, bullet 2], [BES ES.2], [BES ES.3.6], [BES ES.3.7], and [BES ES.4]

Appendix A contains the BES Executive Summary.

[BES 3.1.4]

3.1.4 Computing Needs and Requirements

Success will depend critically on the ability of BES scientists to fully leverage the more heterogeneous architectures — both in types of hardware and in resource size and distribution — to come. The scientific topics of concern to BES have a wide array of computational requirements. Computational demands may be “bursty,” requiring heavy access for durations of relatively short (few months’) time periods, interspersed with periods of low access while new approaches are developed. Nearly instantaneous access with very short turnaround times is required for software development. Proposal processes as well as queuing and access procedures at HPC facilities should incorporate mechanisms to meet these needs in addition to the more common large-scale and often batch-mode provision of resources. Different computational campaigns will require different combinations of the available hardware resources with varying demands on floating point operations, memory size and bandwidth, and interconnect bandwidth. Some problems may require very long runtimes on a relatively small number of processing elements and must be a key component in the long-term strategy for the computing ecosystem.

Software development practices will need to rely on general purpose and domain-specific, accelerator-enabled scalable libraries that hide as many of the hardware specifics as possible. This will also be the best strategy to achieve architectural portability, and perhaps performance portability of scientific applications. Development and use of portable domain-specific libraries should be encouraged, as well as efficient and scalable general purpose middle and lower-level software tools, including message passing interface (MPI) libraries; directive-based threading models such as OpenMP and OpenACC; multiple accelerator-enabled mathematical libraries; and efficient, high-level I/O libraries. Distributed equidistant and non-equidistant fast-Fourier transforms are used in multiple methods. A large variety of quantum many-body methods rely, at least partially, on linear algebra. Distributed diagonalization of large dense matrices can become a bottleneck. Better support for scalable block-sparse linear algebra is of demand, as well. Many post-SCF (self-consistent field) quantum many-body methods, including coupled-cluster theory, configuration interaction, DMRG, and tensor network state methods, rely heavily on distributed tensor algebra.

Addressing these issues is part of a larger objective of pursuing portable performance, both across different machines of the same era and as architectures evolve in time. The efforts required in software development are rapidly increasing, and it is essential to ensure that as much software as possible is still useful 5–20 years into the future. Portable performance will also help bring exascale computing to applications and research codes that have shorter lives and research communities, but which are nevertheless at the forefront in the exploration of new ideas.

Certain adaptive algorithms would benefit from the availability of scalable, asynchronous dynamic runtime systems that may be task based. The implementation of distributed linear algebra is sensitive to the interconnect bandwidth, which currently grows more slowly than the Flop/s. An increased interconnect bandwidth would be highly beneficial. The newer programming models are still in their infancy, and standard interfaces are lacking. Another issue is the performance of the MPI library, which is a common communication layer for many scientific codes. In particular, the efficiency of one-sided communications and nonblocking collectives, including neighborhood collectives, still needs improvement.

Current efforts include the development of reduced-scaling algorithms for many-body methods that have controlled accuracy, based on sparse methods, low-rank representations, fast algorithms, changing representations, nonlinear approximation schemes, and matrix reconstruction, etc. These will need to be cast in high-performance algorithms that must be specified in a manner that is abstracted from architecture-specific implementations.

With the ability to perform computational screening of large collections of chemical systems or representations of complex systems, uncertainty quantification is required to transition these tools from basic research to engineering applications. Currently, these are not systematically employed, except in a few fields such as combustion.

[BES 3.3.3.1]

3.3.3.1 Validation and Uncertainty Quantification in Predictive Simulations

A systematic validation of models and approximations is a fundamental aspect of the development of reliable computational discovery tools (Figure 3-6). An exascale computational environment will greatly accelerate the validation and UQ process and thus lead to computational methods of unprecedented fidelity and robustness.

Simulations involve a wide variety of approximations at all levels of description of matter. At the microscopic level, quantum mechanical approximations include various approximate models of electron-electron interactions, such as DFT, DMFT, and multiple forms of many-body perturbation theory (MBPT), as well as QMC methods. Atomistic simulation methods rely on accurate and transferable force fields and efficient sampling algorithms. The description of transport properties (charge, spin, or heat) requires elaborate models that involve both atomic trajectories and electronic properties. Furthermore, numerical approximations, such as finite basis set in electronic structure computations or potential truncation in force field-based simulations, also affect the reliability of predictions. In addition, length and time scales investigated by computations are often orders of magnitude smaller than those of the corresponding experimental systems.

A systematic validation of such a broad range of approximations requires considerable computational resources and involves a hierarchy of computations performed with models of increasing accuracy.

Current resources allow for the validation of relatively cost-efficient DFT approaches for the description of ground-state properties of simple periodic solids. MBPT approximations require much larger resources and have only been validated in a few systems of very limited size. Comparisons of DFT-based approaches with wave function methods used in quantum chemistry are only feasible for molecules and clusters. The description of liquids and solid/liquid interfaces — essential to understanding many energy conversion processes — requires MD or Monte-Carlo simulations, which further increase the computational cost of the validation process. The validation of hybrid-DFT methods for the description of liquids is currently limited to simple systems and requires the largest resources currently available, namely, via INCITE 1 awards on leadership-class supercomputers.

Validation of force fields also requires long simulations given that they involve the testing of transferability to various problems. UQ resulting from model parameters will require large numbers of simulations to cover high-dimensional parameter spaces.

Currently, a number of important validation procedures are simply inaccessible with existing computational resources. They include the validation of models of transport properties, as well as the validation of linear- or reduced-scaling electronic structure methods in condensed-phase problems. One limiting factor in the validation of linear-scaling algorithms is the need for reference results obtained with more costly algorithms.

An exascale computational environment will open the door to the validation of MBPT approaches in liquids and disordered solids, as well as nonadiabatic approximations for the description of chemical reaction pathways in condensed phases.

A successful implementation of the validation process will require the ability to run scalable AIMD or Monte-Carlo simulations on large computing platforms. Workflow tools and online comparison tools will be essential when making the resulting reference data available to the research community.

[BES 3.3.4]

3.3.4 Computing Needs and Requirements

Computing needs enabling advances in the area of predictive modeling and detailed characterization of synthesis/degradation pathways are inherently heterogeneous. In some cases, memory requirements are limiting multireference wave function calculations for highly correlated electronic systems. In other cases, it is complex workflows for self-consistent multiscale modeling with adaptive resolution where seamless and low-latency connections of rather different computer programs with different scalabilities are required. However, in many cases, ensemble computing already successfully exploits petascale capabilities and will be transferable to exascale resources. For such calculations, the speed of individual compute nodes is often most essential.

High-throughput computational databases are responsible for many successful examples of materials discovery today (e.g., OQMD, Materials Project, aflowlib). We have the capability to scan computationally through thousands of materials to find those with desired properties. However, the current status of these databases is limited in terms of the properties that are calculated (mainly $T = 0$ K energetic properties) and level of accuracy of the methods (mainly standard DFT).

Three distinct, significant directions in high-throughput materials databases are enabled by computational power that is 100–1,000 times greater than present capabilities:

1. More accurate methods (than $T = 0$ K DFT)
2. More complex properties
3. Combinatorial expansion of material space

These future computational resources will allow significant extensions of these databases, which will greatly enhance our ability to predict new materials with tailored properties:

1. The use of methods beyond DFT, such as hybrids, GW methods/many-body, or QMC typically require ~100–1,000 times more computational resources than standard DFT, and hence fit precisely into the scale of future computational architectures. These methods also often have more parallelism than standard DFT (e.g., parallel over bands vs. band pairs, walkers, etc.).

2. Inclusion of a much richer variety of properties in these databases (at either the DFT level or beyond), such as defects, alloys, phonons, disorder, free energies, interactions, and spectroscopy (e.g., photoemission, X-ray, neutrons, Raman, absorption, electron energy loss spectroscopy [EELS]). These properties are more expensive than a single DFT energetic calculation, often again in the 100–1,000 times range, and hence are well suited to future computational expansions.
3. Combinatorial expansion, as detailed above, encompasses the structure prediction problem, missing structures, predicting structures of interfaces, etc., and will each require approximate sampling of ensembles of structures, often in the range of hundreds or thousands of structures.

The issue of “execution management” of these ensemble jobs will also need to be addressed. The current high-throughput databases have developed their own management systems for standard $T = 0$ K DFT (e.g., qmpy for the OQMD, Fireworks for Materials Project, and aflow for aflowlib). However, the inclusion of new methods and properties will require development of new execution management strategies.

[BES 3.5] / [BES 3.5.1]

3.5 Advances in Algorithms for Quantum Systems

3.5.1 Challenges and Opportunities

Quantum mechanics (QM) touches virtually all of computational chemistry, biology, physics, and materials, either directly or indirectly and including science areas that are in other sections of this report. Popular methods such as Hartree-Fock (HF), DFT, second-order perturbation theory (MP2), coupled-cluster theory (CC), multireference methods and QMC are all examples of widely used quantum approaches. In addition, most modern classical force fields are derived entirely or partially from quantum mechanics. The bottleneck in applying the quantum methods noted above to important problems in areas such as catalysis, photochemistry/photobiology, and mesoscale simulations is the high computational cost of scaling the quantum methods with system size, especially the methods that account for electron correlation like MP2, CC, QMC, and multireference approaches. Consequently, in order to facilitate the use of these high-level QM methods for (for example) MD simulations, nonadiabatic dynamics to describe photophysical phenomena, the study of heterogeneous catalysis, the investigation of quantum materials, or the investigation of processes at the mesoscale, it is essential to reduce the scaling of these methods with system size and to develop multilevel parallel algorithms for them.

Implicit in many of the BES mission phenomena is the need to develop truly multiscale methods that can span multiple time and length scales in a seamless and self-consistent manner. For example, substantive improvements in the treatment of electron correlation in model systems has recently been achieved through the development of new algorithms (greatly improved solvers for dynamical mean-field theory [DMFT] methods and their extensions), aided by the adoption of agile, open-source software practices. However, due to their high cost, direct application to complex materials and chemical problems will remain out of reach in the foreseeable future. True predictive power for these complex problems will require development of robust hierarchical theories and algorithms to treat electron correlations across all relevant length scales.

Furthermore, developing new algorithms as discussed in the previous paragraphs will require the collaboration of application developers with applied mathematicians and computer scientists and engineers, as exemplified by the SciDAC (Scientific Discovery through Advanced Computing) program. As discussed in Section 3.8, Next-Generation Workforce, investments must be made in human resources for emerging scientists to integrate expertise in a chosen field such as chemistry or physics with applied mathematics and computer science to advance the field.

[BES 3.5.4]**3.5.4 Computing Needs and Requirements**

QM-based algorithms and applications have a long and successful history of effective utilization of new generations of supercomputers. The first teraflop and first petaflop-scale calculations were DOE-supported, QM-based applications. These won the Institute of Electrical and Electronics Engineers (IEEE) Gordon Bell prizes for fastest scientific application in 1998 (Locally-Selfconsistent Multiple-Scattering [LSMS]) and 2008 (DCA++), respectively. Both applications were not “one offs” crafted for peak performance but later were used productively for research. To extend this path to the exascale era and to address the discontinuous changes in the computational architectures will require significant developments in the computing and software ecosystems, as well as commitments to address the education, development, and availability of a skilled workforce.

Addressing the pace of hardware change is proving challenging to applications teams, even in the 10-petaflop era. To make the best use of their limited resources, many application teams are developing codes for specific architectures. In most cases, there is a substantial investment in effort required to develop or port a code to both CPU and GPU architectures, for example; and with limited resources, this investment must be traded off against the implementation of new theories and algorithms that would immediately lead to new science. This approach inevitably postpones the re-architecting of applications for all exascale architectures. These barriers must be reduced by taking action at all levels of the software stack: by development and education of practical performance portable languages; by porting of standard library functionality to new architectures; and by identifying and extracting common domainspecific functionality to new libraries or computational frameworks. Importantly, mechanisms to communicate these developments to the application teams should be developed. To ensure broad relevance, effectiveness, developer acceptance, and ready identification of deliverables, we recommend that these efforts should be supported by competitively awarded, combined theory development and software development centers that, already at the proposal stage, identify and work with multiple applications within one or more related domains. Utility for multiple applications are essential to avoid customized solutions that are not general solutions for the community. Assessment criteria should include the extent to which the aforementioned challenges will be addressed and improved software instantiated in production domain science codes.

Even in the exascale era, research teams will still have production computing workloads that will require running jobs at all scales. Many important scientific problems are addressed using codes and computational methods that will not initially scale well to the full size of an exascale system, but are nonetheless crucial to the success of BES research programs and user facilities. In addition, runs of moderate size are needed for code and algorithm development, debugging, parameter studies, code validation, and scientific validation of results. These workflows need to run at all ASCR facilities. Making effective use of all the available computational resources requires common middleware — compilers, schedulers, interfaces, libraries, and access mechanisms — across the facilities. Currently, the centers have different access mechanisms and scheduling software, adding to developer cost and requiring additional implementation in workflow automation tools.

Broad adoption of multiscale approaches also increases demands on facilities. At the facility level, techniques for coscheduling different calculations are required: exascale calculations are envisaged to depend on many prior tera- and petascale calculations. These must be run either locally or remotely, with appropriate data transfers scheduled between facilities.

Access to the increasing range of computational scales may benefit from more flexible allocation mechanisms. Currently, the ASCR Leadership Computing Challenge (ALCC) and INCITE award programs have annual periods, while “director’s discretionary” allocations at each center are perceived to be targeted primarily toward ALCC and INCITE readiness, and not to bursts of

scientific activity. Adoption of a “rapid scientific response” mode will enable pressing scientific problems to be addressed and new techniques to be applied on more suitable timescales.

Despite the availability and potential advantages of advanced software runtime systems, domainspecific languages (Fowler 2010), and execution models (e.g., the POSIX Threads library) most quantum simulations are still programmed traditionally using explicit parallelism via MPI, OpenMP, or a combination of both. These newer runtime systems are expected to reduce demand on the application programmer and are better suited to achieving performance portability than legacy approaches such as MPI. However, currently available runtime systems are largely perceived as not being production-level tools. They are also not directly deployed or supported by the computational facilities, and there can be a lag before new architectures are supported, limiting their uptake. Identification and support of a very few runtime systems and standardized interfaces as part of the computing ecosystem will help applications reach higher performance in a more portable manner and enable developers to adopt a more sustainable development strategy.

One of the greatest challenges to application performance in the exascale era is the adoption of increasingly deep and complex memory hierarchies. Effective data placement and use of appropriate memory hierarchy-aware algorithms will enable order-of-magnitude improvements in effectiveness. This sensitivity is already evident in today’s GPU architectures. While some current algorithms have a clear mapping to complex hierarchies, many do not. In addition to developing appropriate language features and libraries that can exploit these features, development of effective and practical tools to characterize the performance of real-world applications is essential. Sufficiently performance-portable and expressive programming models are urgently required; today these are at best developmental research projects.

[BES 3.6] / [BES 3.6.1]

3.6 Computing and Data Challenges @ BES Facilities

3.6.1 Challenges and Opportunities

The BES User Facilities operate more than 240 different instrument types that enable the scientific discoveries of their user communities. BES User Facilities provide photons, neutrons, or electrons and the means to manipulate and detect those particles after interaction with user samples. The heterogeneity of instrumentation contributes to the richness of the capabilities of the BES User Facilities and to the breadth of the user base. It also results in the need for many different data acquisition and analysis methodologies. In the past, and to some extent still today, detectors and end stations generated datasets that were relatively small and readily transferred to users’ portable storage devices and to be taken to their home institutions for analysis. The availability of highperformance networking facilitated this transfer to a user’s home institution. Once the user took possession of the data, the operational responsibility of the BES User Facilities to the user ended.

In recent years, the introduction of rapid and high-performance detectors at some beamlines has led to changes in users’ needs. These detectors routinely produce upwards of 10^2 to 10^3 frames/second, with some to become capable of performing at 10^4 to 10^5 frames per second within the next 5 years — and detector performance will only continue to accelerate. As a result, user datasets have started growing such that users can no longer readily transfer data to their home institutions. Multiple 10- to 100-GB datasets per shift are being produced by some tomography and scattering instruments, and tens of terabytes from the newest electron microscopes. In a three-year “pilot” period ending December 2015 at one BES User Facility, the data from two beamlines have produced more than 244,000 datasets amounting to 1.7 PB and have launched >3.5 million jobs at NERSC. In another example, a single scattering station has occupied OLCF’s entire GPU resource (on Titan). We expect only more of these kinds of demands going forward, and the facilities and their users are either not prepared or are only minimally prepared to address these data-intense environments. Going forward, users cannot be expected to manage and analyze such data volumes

and computations on their own. Experiments have become so complex that users require expert data science and mathematics help to develop suitable analysis methods.

Accelerating detector performance has led to other challenges, such as the need to perform:

- Streaming analysis to enable experimental steering and decision making during the experiment to optimize the scientific outcome.
- Multimodal analysis of concomitant experiments.
- Numerical modeling and simulation.
- Background knowledge from curated archives to drive new scientific discoveries.

Each of these approaches will require significant “on-demand” exascale types of computational resources to deliver the necessary feedback and insights in real time as the experimental process unfolds.

The heterogeneous nature of instruments at BES User facilities means that not one solution will be sufficient to address all of the data challenges. Some instruments may lag in the requirement for data-intense management and analysis. In those cases, “beamside” computation may suffice, particularly if affordable and easily managed petascale “boxes” become available, together with higher-speed networking architectures. Such systems must, however, be “user friendly.” The growing complexity of the analysis process (mixing fast data analysis and numerical modeling) at other instruments will require capabilities beyond the petascale-level capabilities that the ASCR facilities offer. Addressing these challenges in support of users to facilitate and provide for the analysis, management and storage of the data signals a fundamental change in the operation and responsibility of BES User Facilities toward users.

In addition, the computational capabilities of the future will provide a platform for real-time modeling and simulation so that experiments can be augmented and understood as they are in progress. This coupling, along with the need to manage data and the ability to steer and make decisions during the experiment to optimize the scientific outcomes, will require significant “on-demand” exascale types of computational resources to deliver the necessary feedback and insights in real time as the experimental process unfolds.

As users cannot be expected to “go this course alone,” BES User Facilities must partner with the resources of ASCR facilities and research to make the process more efficient. Without such partnerships, in fact, the operational budgets of BES User Facilities will soar, as additional staff will need to be hired who are facile with high-performance computing, data management, and analysis. The commissioning, maintenance, and operation of high-performance computers and storage systems are not trivial and thus are not a simple addition to existing BES User Facilities staff members’ responsibilities. In an attempt to achieve these partnerships, BES User Facilities have begun to pilot, where appropriate, the interaction with ASCR User Facilities, applied math resources, and computer scientists to establish initial working relationships that can produce tangible improvements to the experience of the user base of BES User Facilities. Partnering with ASCR Facilities will also allow BES User Facilities to explore ways to collaborate more closely and efficiently on the development and sharing of analytical tools because each BES Facility would not have to develop its own suite of tools. Cutting-edge computing, data management, and computation are necessities to achieve world-class science. Additional resources from both BES and ASCR will be required, in a partnership, to prepare the BES User Facilities to support the world-class science endeavors of its users.

[BES 3.6.2.1.3]**3.6.2.1.3 Computing Needs and Requirements****Requirements in 5 years:**

- Streaming analysis on high-data-volume beamlines allowing scientists to see high-quality results on timescales sufficient to influence experimental decisions at the beamline or end station.
- A structured engagement allowing computer science and mathematics experts (e.g., a dedicated SciDAC Institute) to work with BES User Facilities to create fully functional and performant analysis pipelines including algorithms, adaptive work flows, streaming data and information visualization, provenance, effective programming pattern, performance assessment, and optimization.
- Automated workflow scheduling systems that enable the effective placement of workflow tasks to minimize and meet demanding response time requirements of running experiments. This capability would include advice on optimal placement and use of compute and networking resources at specific experiments (e.g., the use of petascale systems near the beamline), depending on their analysis pipeline characteristics (e.g., placement of additional computing resources at the instrument to handle the initial data reduction or compression).
- Real-time visualization of results that are viewable at the experimental end station and available/accessible to collaborators at remote locations.

Requirements in 10 years:

- Streaming analysis feeding automated feedback systems that steer DAQ and instrument control systems to facilitate high-quality, time-resolved experiments.
- Streaming analysis and decision support systems that enable scientists to steer their analysis and data taking adaptively and thus optimize scientific outcomes.
- Digital twins integrated with streaming analysis and potentially with DAQ and control systems to steer experiments to temporal, spatial, or parameter regions of maximum scientific interest.

[BES 3.6.2.2.3]**3.6.2.2.3 Computing Needs and Requirements****Requirements in 5 years:**

- Ability to run digital twin experiments routinely, including with comparative analytical capabilities during the experiment. This capability will include multisource streaming analysis in distributed computing environments.
- Access to experimental, computational, and network resources co-scheduled across different ASCR and BES facilities.
- Structured support from computer science and mathematics experts (e.g., a dedicated SciDAC-like Institute) for BES facilities to create fully functional and performant multisource analysis pipelines, including algorithms, adaptive workflows, streaming data and information visualization, provenance, effective programming pattern, performance assessment, and optimization.
- Automated workflow scheduling systems that enable the effective placement of workflow tasks to minimize and meet demanding response time requirements of running experiments. This capability would include advice on optimal placement and use of compute and networking resources at specific experiments, depending on their multisource analysis pipeline

characteristics (e.g., optimized distribution of workloads depending on data source locations and response time requirements).

- Real-time, streaming visualization of multisource results that are viewable at the experimental end station and available/accessible to collaborators at remote locations.

Requirement in 10 years:

- The necessary tools and infrastructure to elevate multimodal data analysis from “one-off” studies to routine and rigorous projects that pull full significance from data. This capability will require a large-scale framework development effort that allows researchers in different experimental domains to develop algorithms and codes that can be combined to fit a single, comprehensive model to all data collected on sets of related samples. Such work must anticipate current and future platforms to allow for the computational demands.

[BES 3.6.2.3]

3.6.2.3 Data Curation

In 1991, the International Union for Crystallography (IUCr) developed a then-state-of-the-art data format for exchange of data and results (Hall et al. 1991). Over the next decade, a data dictionary was developed that carefully defined fields for specification of raw data, metadata, and limited amounts of provenance first for small-molecule, single-crystal diffraction (Brown and McMahon 2002) and later for areas such as macromolecular and powder diffraction crystallography (Hall and McMahon [eds.] 2005), and work continues in areas such as magnetism and superspace structure definitions. During the same period, the IUCr encouraged software developers to support this format and made the use of this format a prerequisite for publication in IUCr journals. The result has been that all crystallographic results, comprising more than 500,000 known structures and five different databases, are available in a single widely implemented format, enabling projects such as the Materials Genome Project. A “CIF” has become a slang synonym for a crystal structure result. Raw data can be imported into analysis programs directly and automatically from journals for alternate hypothesis testing. This level of data sharing and mining is needed to make full use of the data produced by the nation’s user facilities; however, it should be noted that enabling this functionality is more than a simple file format issue: this level of curation and reuse will not be possible without a similar large-scale effort.

[BES 3.6.2.3.1]

3.6.2.3.1 Driving Use Cases

In 2013, the Office of Science and Technology Policy (OSTP) released a directive on Increasing Access to the Results of Federally Funded Scientific Research (OSTP 2013). Its implementation has been linked to providing sufficient supporting information with scholarly publications to enable their validation. To enable compliance, scientists need selective curation — the ability to easily select suitable raw data, analyzed data, metadata, provenance, and tools from their scientific work that led to the insights in the paper in question. As the crystallography case shows, the results of a scientific experiment can be the basis of future research, either as a direct starting point, or as a basis to identify new research directions. To facilitate this reuse, the data have to be sufficiently standardized, documented, discoverable, and immediately usable. The more complex the data collection and analysis process, the more help scientists will need to leverage the results for reuse easily. Furthermore, with large data volumes, fast and targeted discovery methods are vital. New analytical methods, such as the streaming analysis and decision making required to steer complex experiments to more optimized science outcomes, will increasingly rely on fast access to context-relevant data to support scientists in decisions about whether observations are new or how they are different from prior results. Given the time-critical nature of the decisions that the scientists need

to make, fast discovery and extraction of relevant information and delivery are essential. Science relies on scientific reproducibility to validate scientific discoveries, so in a world with exponential growth in scholarly publications, data tools are needed not only to identify research that has produced similar insights but also to facilitate comparison and contrast of those results.

[BES 3.6.2.4.2]

3.6.2.4.2 Gaps

Although sophisticated simulation tools are available, significant gaps and deficiencies are apparent when approaching the tasks described herein. There is a need for global parallel optimization using high-fidelity models, which will help produce robust optimized accelerator designs for present and future projects and with less uncertainty, risk, and guesswork. Real-time, high-fidelity modeling of both storage rings and FELs is needed to anticipate and resolve complex operational issues. Current week-long runtimes for many simulations inhibit using modeling for exploration and insightful design; to be truly useful, runtimes instead need to be on the order of minutes. In addition to reducing runtimes, new and computationally demanding models must be implemented for new regimes where existing models might be inadequate (e.g., 3D CSR).

[BES 3.7]

3.7 Mathematics and Computer Science Transforming BES Science

3.7.1 Challenges and Opportunities

Bridging the gap between BES scientific goals and ASCR computing capabilities will fundamentally rest on transformative physical models, mathematics, and computer science. Mathematics and computer science are how scientific theory, experiment, and computers talk with one another. Obtaining near-term scientific goals and realizing future visions will mean investing in and capitalizing on state-of-the-art evolving mathematics and computer science.

Indeed, this linkage cannot be overemphasized — for it is mathematics that provides the language and blueprint to transform models into equations, approximations, and algorithms that set the stage to take advantage of ASCR’s computing portfolio; and it is computer science that provides the theory, tools, and methods to efficiently execute these blueprints on the most advanced computing architectures.

[BES 3.7.2.2]

3.7.2.2 Experiment

Ever more powerful experimental facilities are creating vast amounts of data — far more than have ever existed before. Experimentation is an integral part of scientific investigation, and BES facilities such as synchrotron radiation light source facilities, neutron scattering facilities, and nanoscience centers generate vast amounts of data. For example, beam science is undergoing a rapid change as facilities probe matter at higher and higher physical resolutions and rapid timescales. These experiments generate massive amounts of data; in the future, they may generate multiple terabytes of data per sample run. The data are often statistical in nature and replete with noise, poor contrast, and signal dropout. Advances in science using these facilities require fundamental advances in the mathematics associated with data science.

Goal 2: *Deliver the mathematical algorithms and unified software environments that allow fast, multimodal analysis of experimental data across different imaging modalities and DOE facilities.*

Fundamental statistical, mathematical, algorithmic, and computational methods are needed to extract information from murky data, interpret experimental results, and provide on-demand

analysis as information is being generated. To make sense of this information, new algorithms that fuse different branches of mathematics are at work. For example, algorithms may combine dimensional reduction, graph techniques, and computational harmonic analysis to perform robust reconstructions from scattering data, and merge partial differential equation methods with machine learning to analyze experimental image data.

At the same time, the landscape of experimental facilities is rapidly changing. In some situations, quick and rough results are desirable while an experiment is under way. In other situations, considerable computation time can be dedicated so as to provide the most accurate reconstruction and analysis possible.

The desire for immediate results from algorithms embedded close to detectors spawns different mathematical questions from those involved in post-processing aided by high-speed networks and extreme-scale computing:

- One end of the spectrum aims at “-demand” computational tools for analysis, data reduction, and feature extraction next to facilities, using embedded advanced algorithms and special-purpose hardware. Here, questions that arise include: What is the minimum/fastest computational model/algorithm that gives (at least some) useful information? Can users quickly determine whether data are useful, are not useful, or are in between? By taking advantage of powerful increases in core hybrid CPUs and general-purpose computing on graphics processing units (GPGPUs), can users quickly perform an analysis in order to steer ongoing experiments to more optimal configurations or output?
- The other end of the spectrum aims at post-processing using reconstruction, inter-comparison, simulation, and visualization using high-performance and extreme-scale computing. Here, different questions arise, including: What is the maximum amount of information that can be measured, processed, organized, and displayed to help understand and shed light on further experiments? Can data be transformed to initialize computational models, with output framed to complement the experiment?

In the near term, identified goals include:

- Building mathematical tools and software environments for real-time streaming analysis of experimental data generated at BES facilities. This desired capability can range from algorithms running on hardware directly embedded with detectors to codes analyzing data efficiently that are shipped using ESnet from facilities to ASCR’s most advanced computing platforms.
- Developing “triage” algorithms that can quickly determine which data are useful, are not useful, and are in between in order to steer the experiment as it is performed and to reduce the amount of data that is sent across networks.
- Devising multimodal algorithms that fuse information from multiple imaging modalities.
- Constructing efficient and robust data and dimension reduction methods, including the use of methods to detect key features in data automatically, perform pattern matching to search across datasets, and employ past information to analyze new results.

In the long term, scientists will best make use of an integrated environment that seamlessly fuses experiments across multiple facilities and compute environments. As an example, in this view, users can sit at one facility and conduct experiments across a range of beamlines, automatically processing samples and fusing results to vastly increase our ability to understand chemicals and materials with great complexity. Keys to this vision include:

- Advanced mathematical algorithms that can extract information from time-resolved, high-resolution data from multiple imaging sources simultaneously.

- Coordinated scheduling of local and remote resources across the full spectrum of the available compute ecosystem.
- Unified software and data curation environments that can allow data to be analyzed across the DOE landscape.

[BES 3.7.2.3]

3.7.2.3 Software Development and Optimization for Extreme-Scale Computers

The extreme-scale systems that will be delivered in the eight- to ten-year time frame will offer unprecedented compute power and challenging parallelism because the machines will contain millions of heterogeneous cores with deep memory hierarchies, but with relatively slow interconnects between nodes. Issues that will significantly complicate software development include the need to explicitly manage data movement, depth and types of memory, and burst buffers and parallel file systems, as well as working with billions of threads, optimizing for power usage, and understanding the behavior of algorithms in the presence of hard and soft faults. For these reasons, refactoring application codes to utilize these architectures effectively will be a massive undertaking. The requirements that were strongly articulated at the meeting by BES scientists were largely focused on improving the ease-of-use of future computers: researchers do not want to spend years of time rewriting their software to utilize the machines, thereby hindering the progress of their scientific endeavors. Thus, our 10-year goal for ASCR computer science support is:

Goal 3: *Build the tools that will make efficient programming of tomorrow's machines as straightforward as programming today's laptops.*

Achieving this goal requires significant development and changes to the entire software stack, from operating system software and run time systems to cross-cutting applied math libraries and application software. Ideally, programming models and languages will emerge that make obtaining efficient performance on the highest-end computers as easy to accomplish as on today's laptops.

This goal is complicated by the fact that computer architectures are becoming increasingly diverse, where code that is optimized to perform well on one computer may not perform well on a different architecture. Tools that allow application software developers to manage this complexity are in high demand. However, the desired realization of this goal has many different forms, ranging from a python-like or domain-specific languages (DSLs) that hide all of the details of the architecture from the application code developer to supplements to the MPI that exposes those details and makes them accessible to the application scientist. Goals in the area of programming models are to provide a range of options, including the following:

- In the short term, it will be necessary to enable the continuity of development of existing codes while ensuring the forward-looking move toward the longer-term goals. This effort will require continued support and development of existing programming models (such as MPI + X where X is OpenMP, OpenACC, or CUDA) and providing tools (such as profilers and debuggers) that will enable developers to analyze their software to prepare for the move to new architectures.
- Development must proceed on community-based standards such as MPI and OpenMP that provide low-level programming models and libraries and that allow effective use of the features of computer architectures. Examples here would include standard methods to deal with new storage and memory hierarchies, which may encompass diverse devices like stacked memory, dynamic random access memory (DRAM), video random access memory (VRAM), and solid state drive (SSD). These developments are already under way, and DOE scientists are playing a key role in their definition. Education and outreach to the application scientists as the standards evolve will ensure that the latest functionalities are as widely used as possible.

- Constructs and tools are needed that work within existing languages, such as C++, to provide higher-level abstractions that will hide the details associated with obtaining performance portability across different computer architectures by decoupling the specification of the science application from how it is mapped onto target platforms. Examples of this capability that are under development today include the Kokkos, Resource-Adaptive Java Agent (RAJA), and Legion tools at Sandia National Laboratories (SNL), Lawrence Livermore National Laboratory (LLNL), and Stanford University and Los Alamos National Laboratory (LANL), respectively. These tools are currently used in a wide variety of math libraries and application codes and are showing significant promise in providing performance portability across different many-core architectures. Incorporation into application codes can further extend the impact of these tools.
- Runtime-centric frameworks are needed that allow for specification of data-dependence (such as graph methods) and asynchronous task-based programming models to significantly improve performance, facilitate portability, increase the level of resiliency, and ultimately improve scientific productivity. Expressing node-level tasks that can be executed asynchronously can lessen the impact of operating system noise, variations within the performance of individual cores, and overall nonuniformity.
- Development must also proceed on workflow tools that enable a suite of interoperable capabilities such as embedded DSLs that emit high-performance tasks for a myriad array of chemical properties, chemical mechanisms, high-order adaptive mesh refinement (AMR) partial differential equation (PDE) stencils, interpolants and operands for a range of turbulence and combustion physics, analytics, UQ algorithms, topological segmentation and tracking algorithms, multivariate statistics, and visualization. Policies and mechanisms need to be developed to manage the execution of end-to-end workflows under strict power and performance constraints while still maximizing the throughput of scientific research.
- Longer term, we can develop DSLs and the associated compilers that provide high-level programmability for particular application areas. Examples of these languages currently exist in materials science for the generation of custom materials (e.g., Matriarch developed at the Massachusetts Institute of Technology [MIT]) but are in limited use. In quantum chemistry, tensor contraction-based DSLs are starting to be used extensively in the community. Research on the use of DSLs to manage performance portability has been explored in relevant applied math and algorithm areas (e.g., structured AMR), but these implementations are still research prototypes and are not yet widely available for use.

Related to this are tools that are being developed by the computer science community to help application scientists improve their productivity when programming for exascale computers. This effort encompasses a broad array of tools, such as those that allow software developers to understand deepening memory hierarchies, data motion and cache use, the trade-offs in performance and power usage, debugging and profiling at scales that contain millions of processes, and optimization of the performance of the software. It is critical that such tools be available on the extreme-scale computing architectures and that they are robust and easy to use for adoption by the application community. This pursuit can be enabled by a co-design type of process where these tools are designed with application specialists and applied mathematicians to develop effective strategies for the broad scientific community. For example, domain-specific scheduling policies may be appropriate for some software while autotuning-driven approaches might be very effective for others. In addition, there must be an effective mechanism for taking the research-level software to hardened software with long-term maintenance.

Even with effective computer science tools to help maximize the use of extreme-scale computers, the architectures are changing in such a way that many of the algorithms currently used by the material science and chemistry communities may not scale or perform well. Thus, existing algorithms may need to be recast, or in some cases radically redesigned, to leverage key

architectural features. In particular, as data motion becomes increasingly expensive compared to floating-point operations, algorithms that have high arithmetic intensity (many floating-point operations for each byte moved from memory) become increasingly attractive. In many cases, these algorithms have been known to exist for many years; however, they were not considered attractive on previous computer architectures. Thus, the community is reconsidering “expensive” but highly accurate methods such as Green’s function methods and AIMD methods and finding that they will be increasingly cost effective on the machines of the future.

Similarly, many algorithms in chemistry and material science scale steeply with the number of particles (e.g., N^3 to N^6), and even significant increases in hardware capability result in minimal increases in the number of particles that can be simulated. Methods that scale linearly or quadratically exist for many application areas; however, until a certain threshold number of particles is reached, these methods have overheads that make them more expensive than their more steeply scaling counterparts. The machines of today and tomorrow are putting us at that cross-over threshold point for many applications. Thus, we are witnessing a revival of linear scaling algorithms with successes for certain classes of applications, such as first-principle electronic structure calculations. Additional research is needed to expand the regimes and problem spaces for which these methods are suitable. Collaborative development of prototypical comprehensive simulation code(s) representing the range of application motifs in BES would be helpful to further flesh out the key requirements for next-generation programming environments and runtimes so these can be adapted to changing architecture designs and new algorithmic approaches.

Finally, it will be important to leverage as much parallelism as possible in the problem formulation. For example, in the case of the GW-BSE approach, it is possible to leverage band parallelism, parallelism over the frequency space, and plane wave basis set parallelism. Parallel-in-time methods also provide parallelism in time as well as in the energy/gradient evaluations.

[BES 3.8.1]

3.8.1 Meeting the Exascale Challenge Is Grounded in Workforce Development

A common theme across the different BES topical areas is that the need for workforce development poses a serious bottleneck to reaching the exascale level of computing — and thus is an essential part of the equation of success in reaching exascale computing. The changes envisioned in computing technologies over the coming decade will require a new generation of computational scientists who are well grounded in their science and engineering disciplines but also knowledgeable about the major computational issues being addressed in computer science and applied mathematics. Likewise, the field also needs computer science and mathematics experts who are knowledgeable about scientific computing and the HPC needs of domain scientists.

While BES researchers have an historic, strong, and recognized synergy between theory, math, algorithms, implementation, and computing, the increasing complexity and dependencies of the exascale era will require these collaborations to be further enhanced. The complexities and multiple layers of hierarchy of next-generation programming environments will require cross-cutting, multidisciplinary teams of domain scientists, applied mathematicians, and computer scientists, who together can cope with the broad range of challenges, from developing the physical and mathematical models, to expressing the scientific workflow, to developing the numerical algorithms, to decomposing the algorithm to the optimal level of task granularity, to expressing fine-grained parallelism in DSLs, and to ensuring that all of the layers of the programming model and runtime have the right abstractions to enable flexibility and performance. Providing the kinds of education and training deliverables we need to prepare these future computational and domain scientists and computational software developers for exascale computing presents a major challenge. In fact, having available and sustaining a sufficiently skilled workforce is considered the greatest risk factor in realizing exascale computational science. Consequently, we anticipate

that significant investments must be made in training a new generation of computational scientists — individuals who are well grounded in their science and engineering disciplines, but also knowledgeable about relevant computer science and applied mathematics issues — to realize the promise of exascale computing.

[BES 3.8.2]

3.8.2 Priority Directions to Promote the Next-Generation Workforce

As stated, staffing and producing the next generation of domain scientists with the requisite training in computer science will be increasingly critical in the near future in pursuit of exascale computing.

Several efforts can be implemented at the DOE/programmatic level. For instance, partnerships between DOE's BES and ASCR that fund computational scientists in the facilities and help develop clear career paths will encourage students and scientists to pursue those career paths. Furthermore, to increase domain scientists' interactions with graduate and undergraduate students, DOE should consider:

- Increasing the number of Computational Science Graduate Fellowship (CSGF) graduates. In fact, expansion of the CSGF program is regarded as one very useful way of pursuing workforce development.
- Creating joint programs between DOE, the National Science Foundation, and universities.
- Having the ASCR facilities, NERSC, and universities share graduate students to foster training and knowledge transfer.

Additional ideas for fostering workforce development, particularly with regard to integrating mathematics and computer science knowledge, include:

- Embedding math/computer science expertise in domain-specific science groups.
- Establishing a center of mathematicians and computer science experts connected to multiple domain science teams.
- Assigning “roaming” architectural experts to visit science groups for 3- to 4-month terms.
- Establishing resident math and computer science experts in the facilities.
- Pursuing funding mechanisms or programs to fill the need for cross-trained graduate students.

HPC training programs similar to the Argonne Leadership Computing Facility's ATPESC (Argonne Training Program for Extreme-Scale Computing) should be expanded, as well as hack-a-thons and coding contests in universities to attract and encourage students. Only by taking actions on all of these fronts can DOE and the universities hope to grow the next generation of domain specialists with the skills needed for achieving success in the exascale computing era.

[BES 3.8.3.1]

3.8.3.1 Collaboration Opportunities

For example, the SciDAC program exemplifies the collaboration that is possible between application developers, applied mathematicians, and computer scientists and engineers in the development of new and better algorithms. These interdisciplinary code development teams are supported (and should continue to be supported) over long time periods to guarantee better code design, increase the lifespan of codes, and give more flexibility to future developments. Expanding this program and developing similar programs to address the training of future scientists (e.g., graduate students and postdocs) would be highly desirable, because the ability of emerging scientists to understand the importance of integrating expertise in a chosen field such as chemistry

or physics (for example) with applied mathematics and computer science will result in needed advances in development of the next-generation workforce.

Although collaboration mechanisms currently exist at scales ranging from the individual PI to national labs and institutions with national reach, collaboration between fields should be particularly encouraged. At the individual PI level, the CSGF program is regarded as highly successful. It receives many more strong applicants than can currently be funded. At the research group level, the SciDAC program is also successful; however, SciDAC's reach has been limited by the number of opportunities and scope of the program. At the national level, additional mechanisms to bring key stakeholders together should be encouraged. One possible mechanism, particularly in the domain of new algorithms, is the CECAM (for the Centre Européen de Calcul Atomique et Moléculaire) workshop model. CECAM is responsible for sponsoring more than 100 workshops, schools, and tutorials each year that typically focus on specific applied topics, with attendance limited to less than 100. NERSC and the ASCR facilities are in a particularly attractive position to lead these efforts with their continuity in expertise in the hardware, runtime environments, mathematics, computer science, and domain science (e.g., catalysts and liaisons). Adapting this format may provide the required mechanism for moving computational research ideas and tools developed in computer science over to the applications domains.

Another mechanism to facilitate information exchange and adoption of new strategies for scientific software design and implementation within the community is to create a shared communication mechanism where domain scientists, mathematicians, and computer scientists could describe their algorithmic challenges and mathematical processes for calculations at scale in a virtual forum. Not only would this facilitate training of the next generation workforce, it would provide continuing education and development opportunities for the current workforce.

[BES 4]

4 PATH FORWARD

For researchers to move forward in addressing the scientific challenges documented in this review, an evolving computing ecosystem must support them. This computing ecosystem includes computational and data systems; scientific applications and software; and the infrastructure for data transfer, sharing, access and analysis — each of which must undergo further investment and development on the path to exascale. Realization of advances in real-time computing, flexible queuing structures, complex workflows, and other policies coupled to the operation of the computing ecosystem must also occur. The coupling of an exascale ecosystem, along with a convergence of theoretical, mathematical, computational, and experimental capabilities, will bring many opportunities for new scientific breakthroughs at an unprecedented scale.

Collaboration between BES and ASCR scientists and facilities staff will help ensure development and deployment of an effective, realistic computing ecosystem that enables revolutionary discoveries in areas described in this report. The computing ecosystem requirements resulting from this review will form the basis to direct future investments of time and resources. These requirements fall into broad categories: methods development; computational environment; data and workflow; and communication and community involvement.

[BES 4.1]

4.1 Methods Development

The advancing complexity of computer hardware requires BES researchers to have more scalable, performant algorithms and applications that are capable of efficient execution on future computing architectures fielded by ASCR facilities. Meeting participants discussed those computing ecosystem aspects that will accelerate or impede their progress in the next 5–10 years. Participants named application codes as well as models and algorithms as key factors, requiring significant methods

development activity. A representative list of the methods development topics discussed by the review participants is as follows (see Section 3 for a more detailed overview of the methods development topics presented by the review participants):

- Strong-electron correlation methods for ground and excited states.
- Theoretical foundation for multiscale coupling of different representations of QM and wellparameterized, validated effective Hamiltonians.
- Inclusion of electronic dynamics such as in TDDFT, DMRG, DMFT, GW, QMC, and emerging methods.
- Full incorporation of relativistic effects.
- Prediction of structures at interfaces.
- Parallel-in-time approaches and other new methods to accelerate the sampling of rare configurations and more realistic systems to generate ensembles with enough statistics for long-timescale sampling of rare events.
- Uncertainty quantification.
- Mathematics that enables order-of-magnitude improvements in speed and accuracy in predictive materials and chemistry modeling.

Other methods development requirements suggest that innovations are needed in order to use the computing resources more effectively and to develop pipelines to and from the resources:

- New strategies for memory management and low-communication/highly parallelizable algorithms to achieve load balancing and effectively use the fast multicore platforms.
- Adaptive algorithms that update “on the fly,” for example, in performing rapid data analysis to guide experiments in real time and changing grids and multiscale methods.
- Fully functional and performant analysis pipelines including algorithms, adaptive work flows, streaming data and information visualization, provenance, effective programming patterns, performance assessment, and optimization.

One of the methods development requirements points to the need to provide a structured engagement allowing computer science and mathematics experts (e.g., a dedicated SciDAC Institute) to work with BES user facilities and researchers.

A close dialogue between BES and ASCR researcher and facilities staff will streamline and promote research and development through the exchange of information about computing ecosystem roadmaps and application requirements and the availability of systems for simulation and testing.

[BES 4.2]

4.2 Computational Environment

Requirements for the access, scheduling, and software ecosystem identify an evolving use-model. The “traditional” HPC model, defined as a large simulation generating data that is then post-processed, is no longer the only primary use-model for many BES projects. Emerging demands, such as for complex workflows and near-real-time computing, are changing the landscape.

New requirements for the computing ecosystem include the following.

- Real-time and near-real-time computing (i.e., nearly instantaneous, with short turn-around times), to provide support for streaming analysis on timescales sufficient to influence experimental decisions, steer DAQ systems, and visualize experiment progress in real-time.

- Proposal/award processes to support the wider array of requirements, including flexible allocations mechanisms that allow for “on-demand” allocations for bursty computational needs and software development.
- A user-friendly development environment, with uniform environments among DOE HPC centers supporting portable, high performance across systems with improved and new runtime systems that mask HPC complexity from application programmers, and training aimed at all levels of HPC developers, including nontraditional HPC users.
- High-performing languages, libraries, and DSLs; and profiling and debugging tools for complex workflows.

[BES 4.3]

4.3 Data

The scale of data generated from both BES simulations and experiments has created an opportunity and a challenge. Even with increased network speeds, datasets from experimental facilities as well as from simulations can no longer be feasibly transferred to a user’s home institution for further analysis. As a result, ASCR and BES facilities must create more data-centric environments with highly effective data analytics tools for their users. Development of such environments and tools will require expertise from domain scientists, data scientists, and applied mathematicians. Continued collaboration will be required to assess proper deployment of the environments as computing resources evolve.

Requirements related to data generation, storage, transport, curation, and exploration include the following:

- Improved tools for data analysis and visualization including machine learning and deep learning techniques for pattern recognition and data quality assessment.
- Tools for real-time analysis of large-scale simulations to obtain, for example, configurations and coordination numbers for metal centers, as well as pattern recognition in real-time to control dynamics in order to steer systems.
- Improved methods and guidelines for data tracking and provenance with searchable and consistent data and metadata from experiment to simulation, analysis, and curation. Push notification for new datasets of interest that have become available, including the ability to link up with the creator(s).
- The ability to carry out “meta” experiments (with other groups’ data as well), including the option to develop new algorithms that can go beyond the original questions that the researchers who collected the data were asking.
- The necessary mathematics, tools, and infrastructure to elevate multimodal data analysis from “one-off” studies to routine and rigorous projects that pull full significance from data. This capability will require a large-scale framework development effort that allows researchers in different experimental domains to develop algorithms and codes that can be combined to fit comprehensive models to all data collected on sets of related samples.
- Improved data access, transfer, storage, and management capabilities that include federated logins across BES experimental and ASCR facilities and high-performance networking from experiment to HPC system.

[BES 4.4]

4.4 Communication and Community Involvement

To foster development of the requisite exascale-level skills and to disseminate this learning widely throughout the community, DOE (with the ASCR facilities) must seek to create or make use of existing initiatives that promote the following:

- Workforce development (education and training).
- Collection and sharing feedback from involvement with standards committees.
- Development of better training materials including best practices, examples, etc.

These activities are ongoing today in multiple institutions; however, efforts to connect them to the larger science community have been attempted on an “ad hoc” basis to date. ASCR facilities can explore new or improved communication channels and activities. In addition, experience has shown some of the best impact from strong collaborations. The previously identified structured collaborative efforts could focus more attention on this important mechanism for community involvement.

[BES page C-36, Dixon]

4. What top three computing ecosystem aspects will accelerate or impede your progress in the next 5-10 years?









Accelerate	Why?
1. Need new approaches to deal with multireference systems (theory)	Multiple metals or even single actinide complexes require multireference techniques and these are not currently available for $> \sim 15$ electrons.
2. F12 correlated methods for all electron relativistic calculations hold promise for reducing the basis set needed to reach the CBS limit.	Large systems with complex ligands will benefit from improved scaling to reach chemical accuracy
3. Spin orbit approaches	Needed for all molecules closed or open shell with heavy atoms ($Z > \sim 40$) to get chemical accuracy.
Impede	Why?
1. Lack of memory and local memory bandwidth	Reliable correlated molecular orbital calculations using current algorithms require substantial memory
2. Processor architectures do not get too simple	Have complicated mathematical algorithms and cannot reprogram for every generation of chip architecture.
3. Fast I/O	Current algorithms for accurate calculations require large, fast, local I/O. Global I/O does not work.

Fault-tolerant architectures are an additional aspect to be considered.

[BES page C-94, Hexemer and Parkinson]

4. Top three computing ecosystem aspects to accelerate or impede progress, 5–10 years?

Accelerate	Why?
1. Development of distributed, scalable software infrastructure incorporating workflow, data management, and resource management.	To enable the BES facility goal of delivering science knowledge (rather than data) requires the development of the Super-Facility, using ASCR resources.
2. Access to guaranteed Network QoS (perhaps with SDN) and the ability to co-schedule network and compute resources with beamline experiments.	In-situ, time-resolved experiments needing real-time feedback have burst needs for network, compute, storage, and other resources.
3. Inter-facility federated ID and single sign-on capabilities.	As BES facilities build partnerships with multiple ASCR facilities, disjoint trust domains and procedures are barriers.
4. Low-overhead and high-bandwidth data I/O capable systems (e.g., burst-buffer or RDMA-like I/O).	Many processing and analysis systems and applications for experimental and observational science are very data I/O intensive.



APPENDIX H

Sections Cited in This Report
from the Exascale Requirements
Review for Biological and
Environmental Research (BER)

This page is intentionally left blank.

SECTIONS CITED IN THIS REPORT FROM THE EXASCALE REQUIREMENTS REVIEW FOR BIOLOGICAL AND ENVIRONMENTAL RESEARCH (BER)

The following sections from the Biological and Environmental Research (BER) Exascale Requirements Review report are cited at least once in this Crosscut report.

[BER ES.1], [BER ES.1 Key Findings], [BER ES.3.1.1], and [BER ES.3.1.4]

Appendix B contains the BER Executive Summary.

[BER 2.1]

2.1 BER Vision

BER's two divisions, the Biological Systems Science Division (BSSD) and the Climate and Environmental Sciences Division (CESD), have research programs and user facilities that encompass laboratory- to field-based research and observation, as well as numerical modeling across the full range of spatial and temporal scales of interest to BER and DOE.

BER's transformative impact on science encompasses mapping the human genome, including via the U.S.-supported international Human Genome Project that DOE began in 1987. Today, BSSD researchers are using the powerful tools of plant and microbial systems biology to pursue fundamental breakthroughs needed to develop sustainable, cost-effective biofuels and bioproducts from renewable biomass resources. Conducting this research poses challenges in the face of extremely rapid changes in biotechnology and high-throughput analysis techniques. Current genome sequence production, particularly for plants and microbes, outpaces the field's ability to efficiently interpret gene function. In addition, high-throughput analytical and instrumental capabilities produce enormous data streams that pose daunting challenges for data management and analysis. New approaches are needed to make more effective scientific use of the enormous volumes of data generated within BER's biological science programs.

In addition, BER has been a core contributor to atmospheric, environmental, and climate science research since the 1950s, beginning with atmospheric circulation studies that were the forerunners of modern Earth system models (ESMs) and by pioneering the implementation of climate codes onto high-performance computers. Today, CESD research contributes to model development and analysis using community-based models, including the Community Earth System Model (CESM), the Energy Exascale Earth System Model (E3SM), and the Global Change Assessment Model (GCAM), as well as numerous system components that are deployed broadly across international modeling communities. These leading U.S. models are used to address the most critical areas of uncertainty in contemporary Earth system science; such as cloud changes and feedbacks, aerosol-cloud interactions, and changes to the most sensitive ecosystems. BER has been a pioneer of ecological and environmental studies in terrestrial ecosystems and subsurface science and seeks to describe the continuum of biological, biogeochemical, and physical processes across multiple temporal and spatial scales that control the flux of environmentally relevant compounds between the terrestrial surface and the atmosphere. BER-supported modeling includes development of scale-adaptive approaches for all Earth system components, including atmosphere, ocean, cryosphere, and land.

For land systems, the research extends across CESD and BSSD, as scientists work to quantify fine-scale hydrologic and biogeochemical processes from surface waters through groundwater in key watersheds. To achieve these types of integrated systems-level modeling, advances are needed in multiscale, multiphysics codes. Therefore, efforts include improving library interoperability among different codes and implementing software improvement practices that will enable the future development of an agile collection of interacting components to form a “software ecosystem.”

[BER 2.1.1]

2.1.1 BSSD Vision

Biological systems science integrates multidisciplinary discovery- and hypothesis-driven science with technology development on plant and microbial systems relevant to national priorities in sustainable energy and innovation in life sciences. As such, BSSD supports systems biology research: the multidisciplinary study of complex interactions specifying the function of entire biological systems from single cells to multicellular organisms (rather than the study of individual isolated components). These systems biology approaches to BSSD research seek to define the functional principles that drive living systems, from plants, microbes, and microbial communities. These principles guide the interpretation of the genetic code into functional proteins, biomolecular complexes, metabolic pathways, and the metabolic/regulatory networks underlying the systems biology of plants, microbes, and microbial communities. Advancing fundamental knowledge of these systems will enable new solutions to long-term national challenges in sustainable energy production, breakthroughs in genome-based biotechnology, the understanding of how microbial activity affects the fate and transport of materials such as nutrients and contaminants in the environment, and development of new approaches to examine the role of biological systems in carbon cycling in the Earth system.

BSSD systems biology research is primarily organized under the DOE Genomic Science Program. A major effort within the current portfolio is to obtain a fundamental understanding of the biology of plants and microbes as a basis for developing cost-effective processes for production of bioenergy and bioproducts from cellulosic biomass and other plant-based energy resources. Broader research efforts in plant and microbial biology within BSSD seek to expand the range of organisms useful for bioenergy purposes; understand relationships between plant and microorganisms relevant to sustainable biomass production; and develop the biotechnology approaches to design new biological systems with beneficial bioenergy or environmental properties. BSSD also uses systems biology approaches to advance DOE missions in environmental process understanding related to climate variability and the movement of contaminants through ecosystems. This research examines biological processes occurring in terrestrial soils, subsurface aquifers, and a variety of other environments relevant to BER. The goal of these efforts is to gain a predictive understanding of factors controlling carbon and nutrient cycling, determine how biological communities respond to changing environmental variables, and integrate micro-scale biological process understanding into ecosystems. DOE systems biology research includes large, team-oriented multidisciplinary efforts (such as the DOE Bioenergy Research Centers); scientific focus areas within the DOE national laboratories; medium-scale collaborative efforts between academic institutions and national laboratories; and focused, single-investigator projects.

BER operates and manages user facilities that advance biological and environmental research for DOE mission solutions. BSSD’s portfolio is supported by the DOE Joint Genome Institute (JGI), a DOE national scientific user facility providing genome sequencing, DNA synthesis, metabolomics, and interpretation capabilities to the research community; the DOE Systems Biology Knowledgebase (KBase), an open-source computational platform for assembly, analysis, and sharing of complex “omics”-based data; and infrastructure facilities for development of new bioimaging, measurement, and characterization technology for visualizing and describing genome-based processes within living cells.

With a long history in microbial- and plant-based genomics research coupled with substantial biotechnological and computational capabilities available within the DOE user facilities, BSSD is well positioned to make significant contributions in bioenergy and biotechnology research. The complex nature of BER research themes and the increasing need to develop new ways to assemble, integrate, and analyze multiscale and multipurpose component codes, as well as enormous data sets from diverse disciplines across multiple scales, require new advances in computational approaches and capabilities. It is from this perspective that exascale computing offers the ability to examine BER-relevant science in new and unprecedented ways.

[BER 2.2.1]

2.2.1 BSSD Objectives and Grand Challenges

Driven by revolutionary advances in genome sequencing technology and new tools and methods for biotechnology, biology is becoming an increasingly quantitative science. The essence of the New Biology for the 21st Century (NRC 2009) has arrived. Biologists now have access to suites of new tools and integrative technologies to explore and probe complex systems in ever-greater detail, from molecular events in individual cells to global biogeochemical cycles, accelerating our fundamental understanding of these systems. Exciting new integrative opportunities with the physical sciences, engineering, computational science, and mathematics offer a wealth of idea-rich exploration. In addition, new technologies are revolutionizing biological research, making it easier, cheaper, and faster to generate greater volumes and types of data. A major challenge for genome-based research is the development of new, more distributed and collaborative approaches to meaningfully analyze and interpret large, diverse data in an effective, reproducible, and shareable manner. In this context, advanced computational systems and methods are key to future research efforts in biology and biotechnology. By integrating genome science with advanced computational and experimental approaches, the BSSD programs seek to take advantage of advanced computational systems at the exascale to propel genome-based science toward the goal of gaining a predictive understanding of living systems, from microbes and plants to their interacting communities. Toward realizing these goals, BSSD's scientific thrust is in the following two areas:

1. *Systems analysis of the collective -omics (e.g., transcriptomics, proteomics, and metabolomics) of plants and microbes.*

BSSD supports programs and scientific user facilities with integrated experimental and computational functions for discovery science and technological innovation. The rapid development of -omics techniques have provided an unprecedented amount of biological data, thereby shifting the bottleneck in scientific productivity from data production to computational analysis, interpretation, and visualization. When integrated, these data will accelerate scientific discovery and become a foundational archive for interdisciplinary data mining, data analysis, visualization, and predictive modeling. Creating detailed characterizations of the genomic, transcriptomic, proteomic, and other -omic states of biological samples and combining these multiple -omic data types in an integrated framework over time and space would enable us to probe the molecular mechanisms of diverse environmental and host-associated microbial communities. However, the high dimensionality of the -omic data generated from systems across scales of space, time, and organizational complexity present significant computational challenges in identifying causal variants and modeling the underlying systems biology. To address these scientific challenges, BER researchers need novel computational tools to direct systems-level investigations to develop:

- New and innovative computational strategies to enhance, scale, and optimize the management and processing of large, complex, and heterogeneous data generated from different scales for effective integration and interpretation of systemwide data.
- Powerful algorithms for modeling and optimization to advance more sophisticated analysis and (re)design of genome-scale metabolic networks.

- A generalized framework for comparative analysis of data obtained from multiple modalities and across experimental conditions and environments in order to build predictive models of the system.

2. *Development of new and advanced methods for characterizing and imaging molecular systems.*

To understand how genetic information translates to function, BSSD is seeking development of new multifunctional, multiscale imaging and measurement technologies to visualize the spatiotemporal expression and function of biomolecules, intracellular structures, and the flux of materials across cellular compartments. BSSD is seeking new ways of combining existing technologies such as the molecular-scale science capabilities within BSSD's Structural Biology component, the division's new effort in Bioimaging Technology development, and technologies available at DOE user facilities (such as the JGI and the Environmental Molecular Sciences Laboratory [EMSL]) to develop approaches to identify, image, track, and measure key processes occurring within plant and microbial cells. These collaborations will enable testing and validation of current hypotheses of cellular function and the generation of new ideas for modified cell functions.

A comprehensive, biologically realistic characterization of a system relevant to bioenergy and the environment integrates not only genomic and physiological data but also three-dimensional (3D) and 4D molecular structural data to gain a deeper understanding of BER-relevant microbial and plant systems.

The wealth of -omics data currently available to researchers has enabled new areas for bioenergy research. However, quantifying and localizing the specific sites where particular enzyme reactions are occurring, or identifying the regulation of material flux into, within, and out of a single cell, are difficult tasks to perform without supporting imaging data. Complex biosystems such as plants or plant/microbe interactions make this task even more challenging. New integrative modeling approaches are needed to understand the flow of material and other dynamic processes in a quantitative manner while also placing systems biology information into the whole cell/biosystem context by using bioimaging data as spatial, temporal, and chemical boundary constraints. These new approaches should create an iterative experimental regime where initial -omics models drive experimentation and experiments further refine virtual cell/biosystem models, leading to more targeted hypotheses and continued experimentation. Constantly updating theory and virtual models with experimental data for finer resolution will enhance the likelihood that one day, we will be able to fully predict and control biosystem behavior. To realize the objective for creation of quantitative and integrative models localizing cellular dynamics and development of a "virtual cell" iterative modeling framework, BSSD seeks to:

- Understand how genetic information translates to function by developing new multifunctional, multiscale imaging and measurement technologies to visualize the spatiotemporal expression and function of biomolecules, intracellular structures, and the flux of materials across cellular compartments.
- Foster the development of new nondestructive, in situ imaging and measurement technologies to visualize the spatial and temporal relationships of key metabolic processes governing phenotypic expression in living biological systems, including plants and microbes of potential interest to BER.
- Build on the iterative experimentation approach of physics-based simulations and multimodal imaging systems biology to inform and validate biodesign principles.

Systems biology studies involving integrated approaches from physical, computational, and experimental sciences will enable a mechanistic understanding of the spatiotemporal expression

of biomolecules and structures within microbial and plant cells, as well as the dynamic nature of cellular metabolism.

[BER 2.3.2]

2.3.2 Mapping CESD Programmatic Objectives to Computing Ecosystems

The CESD science goals and grand challenges rely upon robust computational facilities to provide leading-edge computational performance for the advanced Earth system modeling capabilities, along with substantial mid-level computing for component development and model testing purposes, and finally computers designed for large-data manipulation and analysis capabilities.

CESD modeling supports the development of high-fidelity models representing Earth system changes in order to improve understanding of the significant drivers, feedbacks, and uncertainties within the integrated Earth system and thereby to provide vital information needed for effective energy and connected infrastructure planning. While predictive capability is improving in Earth system models for features such as long-term mean temperature change, other features continue to be major challenges, such as trends in precipitation and extremes in temperature, precipitation, and storms. These challenges require computational modeling of the fully coupled human-Earth system at high resolution to reliably simulate water cycles, biogeochemical cycles, and cryospheric processes that determine the changes in the Earth's energy balance, weather systems, regional precipitation and water resources, extreme events, sea level, and coastal inundation. Uncertainties in future projections persist because of a combination of insufficient process representations, model resolutions, and ensemble members. Consequently, we will need advances in theory and modeling to determine how best to configure climate simulations (e.g., by adding more detailed processes, more ensembles, or higher resolution) for Earth system prediction. For some purposes, we will need a hierarchy of models, from global to process scale, with varying degrees of detail that will provide essential information across scales and sectors. Research and attention are needed to couple processes where important feedbacks occur, and to apply one-way coupling for those needs where feedbacks are less important, such as for particular impacts. Improved frameworks will enable us to integrate models and observations for model testing and initialization, with appropriate attention to scale.

Within CESD, the Earth System Model (ESM) program's Accelerated Climate Model for Energy is a key capability for modeling the coupled Earth system, including significant human activities such as land and water management, to run at highest and at a regionally refined resolution given current computational capabilities. With advances in computation, ACME will continue to increase resolution into the nonhydrostatic state (e.g., 1 km or smaller with regional refinement) in order to capture important high-resolution features, such as mesoscale convective systems, cyclones, frontal systems, ocean-eddy transports, ocean-boundary and mixing effects, polar changes including sea-ice and ice-sheet changes, coastal inundation, and storm surge. Coupled cryospheric changes will be included to accurately model sea-ice change at high-resolution, the coupled dynamic ice-sheet changes that are needed for sea-level simulation, and permafrost degradation. ACME's land simulation will develop to include integrated hydrology and biogeochemistry, vegetation that is dynamic and increasingly based on plant trait methods, and subgrid-scale orography to more accurately simulate and couple atmospheric and land processes.

It is helpful to consider the scientific potential of exascale computing. The current high-resolution (25 km) coupled version of ACME can be run to simulate about 100 years' worth of time using substantial allocations on DOE's current petascale computers. This level of resolution has some scientific value for examining the ability to capture high-resolution and weather-scale phenomena such as fronts, hurricanes, atmospheric rivers, and eddy-transport and boundary-layer phenomena in the ocean. However, to use the ACME for climate-related research purposes, performing longer simulations would achieve climate equilibrium, and a large ensemble of simulations (up

to 100) would be needed to obtain useful climate statistics. Exascale machines could be used to achieve these more extensive simulations with the current model resolution. In addition, Earth system modeling could continue to push the margin of resolution down to 10 km globally, with regional refinement to 1 km; or alternatively, to achieve cloud-scale-resolution using a super-parameterized approach (with a cloud model embedded in the Earth system model grid boxes). This higher resolution would expose yet more detailed phenomena, such as thunder-storm systems and interactions between population centers and weather within single “hero-simulations,” but it still would not be sufficient to provide the statistics that an ensemble would provide.

The potential advances for Earth system modeling will be affected by computer architecture choices. In general, higher memory bandwidth is very helpful for Earth system models, given that most codes utilize operators that have a relatively high ratio of load/stores to flops, so improvements in memory bandwidth within the node will have significant impact. Second, climate modeling performs better on fewer yet more powerful nodes over systems with large node counts but less-capable nodes. At current spatial resolutions, scaling to very large node counts means low subdomain size, where the reduced workload cannot make efficient use of the cores and means there is more sensitivity to the communication overhead. Finally, emphasis on LINPACK performance should be de-emphasized in favor of hardware designs that allow more general code to obtain a significant percentage of the processor’s peak performance. Earth system models often obtain less than 10% of a peak, and this percentage is expected to further decrease under current hardware trends.

To achieve the advanced modeling capability also requires accelerated cycles of model development, testing, calibration, and evaluation using advanced metrics and diagnostic tools. Furthermore, it requires substantial increases in computing power. Based on current trends, all future increases in computing power are expected to come from new architectures, with increased use of multicore and other types of acceleration. Adapting Earth system models to take advantage of these new architectures presents several challenges. For example, all performance gains on new architectures are expected to come from increases in concurrency. Algorithms with serial bottlenecks will have to be replaced with new algorithms and modeling approaches that avoid these bottlenecks. Yet even with a highly parallelizable algorithm, it can be difficult to write parallel software that can take advantage of this concurrency. New programming models are needed that make it possible for the developer to exploit every last bit of parallelism in the algorithms. In addition, programming model abstractions are needed that can help the developer quickly adapt to new hardware and support competing hardware approaches. Finally, the complexity and size of Earth system models, which involve thousands of subcomponents coming from many different modeling teams, make it challenging to adapt new programming models and abstractions. The BER-ASCR Scientific Discovery through Advanced Computing (SciDAC) partnership program is critical for achieving these computational and algorithmic advances and for developing theory around connecting predictability to uncertainty in climate processes.

The overall goal of the Regional and Global Climate Modeling (RGCM) program is to enhance predictive understanding and modeling of climate variability and change by advancing capabilities to design, evaluate, diagnose, and analyze a suite of global and regional Earth system model simulations informed by observations. Analysis of a hierarchy of Earth system models is important for providing a holistic picture of the predictability of the Earth system. The RGCM program is designed to address uncertainties in regional climate projections from the perspective of a predictive understanding of water cycle, clouds, biogeochemical cycles, high-latitude feedbacks, extreme events, and modes of climate variability. High-resolution models and the large ensemble of simulations needed to understand the uncertainties within the Earth system give rise to copious amounts of model output that encompass temporal scales spanning seconds to millions of years (10^0 – 10^{13} s) and spatial scales of microns to tens of thousands of kilometers (10^{-6} – 10^7 m).

Integrating model output and simultaneously synthesizing observational data to enhance an understanding of the Earth system offers new opportunities for scientific discovery. Data mining algorithms, data assimilation techniques, and advanced statistical tools are required to extract knowledge and information from observations and model data. In situ visualization approaches that exploit large-scale, distributed-memory, parallel computational resources offer significant promise. In addition to supercomputer architectures designed for compute-intensive simulations (many cores, fast interconnects, small memory), large-scale climate analysis will require machines with fewer cores, large and fast on-node memory, high-bandwidth input/output (I/O), and fast access to large volumes of storage.

Atmospheric System Research (ASR) supports basic research on atmospheric processes involving clouds, aerosols, precipitation, and radiative transfer by using process-focused models, observations (particularly ARM), and laboratory research. Current processes of interest include aerosol microphysical processes, convective physics, and microphysical and dynamical processes in low clouds in both temperate and high-latitude regimes, including land-atmosphere interactions. Because the native scale of these processes and of their observational reference data is at the subgrid level even for current regional models, the ASR vision is to support the parameterization strategies for large domain model development with observationally constrained, scale-aware process models, which can be realized with direct physical representation on high-resolution temporal and spatial domains, such as large-domain Large Eddy Simulation (LES) models. These process models can then inform advancement of physically realistic global model components for the exascale computational environment.

The vision of the ARM Climate Research Facility, a DOE Office of Science user facility supported by CESD, is to provide a detailed and accurate description of the Earth's atmosphere in diverse climate regimes to resolve the uncertainties in Earth system models. ARM is increasingly coupling high-resolution LES models with its observational capabilities in order to accelerate the understanding of key atmospheric processes. ARM's computational challenges in the 5- to 10-year time frame include: large-domain LES to extend ARM's observational-modeling coupling to mid-latitude, deep convective systems and Arctic regions with complex surface and lateral boundary conditions; incorporation of complex instrument simulators and/or data assimilation in LES; processing and storage of high volumes of observational and model data; and computational techniques for quality control, data mining, and visualization of large-volume observational data sets.

The Subsurface Biogeochemical Research (SBR) program seeks to advance a robust, predictive understanding of how watersheds function as complex hydrobiogeochemical systems and how these systems respond to perturbations caused by changes to water availability and quality, land use and vegetation cover, elemental cycling, contaminant transport, and compounding disturbances. SBR researchers are encouraged to use a systems approach to probe the multiscale structure and functioning of watersheds and to capture this understanding in mechanistic models representing both the complexities of the terrestrial subsurface and ecohydrological interactions with surface water bodies and vegetation. The SBR program efforts include development of genome-enabled biogeochemical models of the multiscale structure and functioning of watersheds. These mechanistic models are based on reactive transport codes, which incorporate metabolic models of microbial processes; molecular-scale understanding of geochemical stability, speciation, and biogeochemical reaction kinetics; and diagnostic signatures of the system response across vast spatial and temporal scales. State-of-science understanding codified in models provides the basis for testing hypotheses, guiding experimental design, integrating scientific knowledge on multiple environmental systems into a common framework, and translating this information to support informed decision making and policies.

A priority for the SBR program is to advance the development of a community-driven collection of multiscale, multiphysics models to facilitate the iterative cycle of model-driven experimentation and observation and thus accelerate scientific discovery. An initial co-sponsored SBR and ASCR project entitled “Interoperable Design of Extreme-Scale Application Software” (IDEAS) is using SBR-supported, terrestrial use-cases to drive improvements in software development practices and library interoperability, underpinning a shift toward a more agile collection of high-quality composable components to ultimately enhance productivity (DOE 2017). This paradigm shift to an agile collection of interacting components, or a “software ecosystem,” acknowledges the need to progress beyond the modularity of traditional multiphysics codes to a higher level of interoperability. More recently, members of the IDEAS project team proposed a more comprehensive software development project entitled xSDK4ECP that was selected by ASCR’s Exascale Challenge Project for support.

High-performance production computing has long been a significant capability offered to the scientific user community by EMSL, a DOE Office of Science user facility supported by CESD. In contrast to the capability computing embodied in ASCR’s exascale systems, EMSL’s computational system and associated codes are oriented more toward capacity computing. However, users of EMSL’s computing capabilities are also able to iterate between simulation, experimentation, and observation to enable greater understanding of biological and environmental systems of interest. From a BER standpoint, EMSL’s production computing capabilities, software codes optimized for molecular to mesoscale modeling, and experimental data archive capabilities enable BER-funded scientists to undertake systems science research ranging from study of molecules to genomes, single cells, microorganisms, microbial communities, atmospheric particles, the rhizosphere, subsurface and terrestrial ecosystems, watersheds, and regions. From a computational standpoint, EMSL’s midscale production computing environment fills a gap in hardware resources available to the scientific community, and therefore provides an important complement to ASCR’s exascale systems.

Both Earth systems modeling and systems biology multiscale modeling will require efficient methods to address the analysis and comparison of extremely large genomic, meta-omics, and multimodal imaging data; furthermore, large Earth system model and complex observation data sets and associated data acquisition, storage, management, analysis, and utilization challenges will need to be addressed.

[BER 3.1]

3.1 Biological Systems Science Division (BSSD)

Biological Systems Science has the ambition to discover, characterize, and predict the complex interplay of biological and abiotic processes that govern ecological dynamics, environmental change, and, ultimately, the health of our biosphere and its constituent organisms. The resulting foundational knowledge base should facilitate our ability to (1) harness biological processes to employ biology as catalysts to transform waste streams and other diverse feedstocks into sustainable and renewable energy and chemical outputs, and (2) improve and prevent environmental degradation by climate effects and anthropogenic inputs into water, soil, and air. An important side benefit of these goals are the discoveries of new branches of life, an understanding of their evolutionary and ecological histories, and the elucidation of the genetic potential encoded in the heritable materials of Earth’s biome including, for example, new antibiotics, new catalysts for important industrial routes to advanced materials and drugs, and materials for defending against plant and animal pathogens.

In contrast to the physical sciences, biology presents key challenges to becoming the predictive science that is necessary to achieve these goals. There are still deep knowledge gaps in mapping the diversity of biological organisms on Earth and understanding the vast panoply of functions encoded

by their genomes. These functions act within diverse and changeable environments, themselves only shallowly characterized. This interdependence of micro- and macro-systems presents hurdles to understanding adaptive dynamics and the evolution of biological communities and the environmental transformations they can render.

An extraordinary upward inflection point in measurement technologies is now exposing an unprecedented array of data, spanning molecular structures to ecosystem productivity and producing unprecedentedly large data streams across a wide array of spatiotemporal scales. These include, but are not limited to, scaling innovations in nucleic acid sequencing; mass-spectrometry methods for protein and metabolite detection and quantification; and multiscale imaging techniques spanning small-angle X-ray and neutron scattering, cryo-electron microscopy tomography, hyperspectral and ultra-resolution optical imaging, and mass-spectroscopic imaging. These new data sources are driving the biological community toward a common goal of producing multiscale models that cross-reference and connect these data to facilitate effective predictions of function from the gene to biosphere level.

There are challenges that are specific to this biological ambition. Data that are relevant to any particular question in biology tend to derive from diverse types of complex measurements. The objects under study — biomolecules, cells, and organisms and their communities — are highly variable (behaviorally and genetically) and have activities that are highly context dependent. In turn, this variability requires complex experimental design to control for it as much as possible. Large-scale analysis and data integration rely on having high-quality and formally controlled labels for experimental conditions and designs, as well as biological identities so that data can be properly cross-referenced and analyzed with the proper statistical methods. Because of the remaining knowledge gaps concerning many of the key players (molecules or organisms and their activities), making inferences about these “unknowns” is a constant activity, alongside training models on the known molecules and processes. Finally, it is not uncommon for single base-pair changes to drastically change the fitness of an organism in a given environment and permit its takeover dependent on spatial, environmental, and population genetics factors. Thus, it is possible for a single-molecule event to propagate to ecosystem scale relatively quickly. Multiscale prediction is thereby a critical goal.

Achieving DOE goals for biological system science demands computational innovation in: (1) sophisticated data streaming, online analysis, storage, and representation; (2) rapid, nearly interactive access to and manipulation of large heterogeneous data sets; (3) algorithms for machine learning that respect the heterogeneity, interconnectedness, and uncertainty in biological data and that can be run and updated continuously by the biological community — where each laboratory will be evolving into petabyte-level data producers over the next ten years; (4) innovations in the discrete mathematics and graph-based algorithms that are central to sequence assembly, phylogenetic analysis, and analysis of large-scale influence networks in biological data; and (5) integration of large-scale biophysical codes ranging from quantum mechanical (QM) methods through molecular dynamics (MD) to stochastic reactive transport codes to facilitate the multiscale predictions required at all levels of biological systems sciences. Together, these recommendations suggest that DOE bioscientists will need facile access to a diversity of biological data systems and computational frameworks, including real-time and interactive access. The computer systems required will include those optimized for images, genomics, and machine-learning workloads, as well as exascale systems for the largest simulation and analytics challenges. Complementing these systems is a need for new algorithms in analysis and simulation, along with community software that can be easily adapted to multiple biological scenarios.

On the following pages, we divide these challenges into four areas with interrelated computational and data requirements that build one on the other. The first is Section 3.1.1, “Multiscale Biophysical

Simulation from Molecules to Cells,” which addresses the challenges that arise as we drive models down to an increasingly atomistic/biophysical understanding of the biological material from single proteins, through their complexes, to the complex spatiotemporal dynamic and material and energy interchange in complex tissues and communities of organisms. The second, “Mapping Sequence to Models” (Section 3.1.2), specifically outlines how we organize information around genetic sequences to aid in making models of biomolecular, organismal, and community function through prediction of encoded macromolecule identity and function and their regulation and variation. The third, “From Microbes to the Environment” (Section 3.1.3), examines the requirements to map processes at the microbial scale all the way up to environmental processes that impact water and soil quality and climate variation. Finally, “Biological Big Data Challenges” (Section 3.1.4), recognizes the issues surrounding the diverse, heterogeneous, and high-volume/velocity issues that biological systems sciences face.

[BER 3.1.1.4]

3.1.1.4 Computing Needs and Requirements

3.1.1.4.1 Molecular Dynamics

Historically, MD codes have been able to use all of the available floating point operations (FLOPs) on whatever hardware has been available. Therefore, a key requirement for exascale MD is high number of FLOPs (in particular, a large number of single-precision FLOPs rather than a lower amount of double-precision FLOPs); memory requirements, in contrast, are modest. Another key requirement for exascale MD is high-bandwidth, low-latency communication between computing nodes, between central processing units (CPUs) and accelerators, and between accelerators. Presently, this capability is conspicuously lacking for GPUs: while GPUs provide tremendous computational power for MD simulations, they come with a burden of high-latency and low-bandwidth communication. The GPU-to-GPU communication bottleneck is unlikely to be solved by the upcoming NVLINK communication technology. In addition, for GPUs, there is projected to be an additional exascale kernel launch latency problem, the solution to which is likely to involve both GPU hardware and software.

Code development for accelerators, both GPUs and Intel Phi, currently requires a high level of sophistication from the software developer. In addition, many approaches used today, such as CUDA or Intel C language intrinsics, are highly hardware-specific, requiring considerable code modification when ported to new hardware. Use of directive-based approaches to software development, such as OpenACC and OpenMP 4, may be beneficial for many parts of the MD code; however, in their current incarnation, they fail to provide the level of performance needed when implementing the most compute-intensive parts of the code. Some code rewriting is inevitable because different algorithms are optimal on different hardware platforms.

On CPU machines, low-overhead threading programming models are needed to improve node-level parallelism. The existing Open Multi-Processing interface still suffers from high scheduling overhead at large thread counts per rank. Also required are advanced vector, single-instruction multiple data (SIMD) instructions/compiler intrinsics with broader vector width to push data parallelism.

3.1.1.4.2 Quantum Chemistry

Hardware requirements for QM codes have some overlap with MD codes, where overall speed is limited largely by latency owing to the extreme need for internode communication. However, QM methods also have unique requirements for large memory and I/O, the latter primarily for storing temporary files when memory is insufficient.

Priorities include these: efficient concurrency management to maximize the parallel performance, auto-tuning and library-oriented design of key computational kernels, network topology-aware execution for networks with hierarchical topologies, and fault-tolerant execution models that can recover from inevitable hardware failure events. The main overarching theme in all of these directions is to shift from monolithic code structure to loosely coupled and nearly independent execution processes.

[BER 3.1.3.4]

3.1.3.4 Computing Needs and Requirements

To address the scientific challenges of multiscale/multiphysics model integration outlined above, several computing requirements were identified, the most pressing of which is the fundamental algorithm and software development necessary to enable superparameterization and model coupling. The scientific disciplines of interest here are relatively immature computationally and otherwise; significant research investment is needed both to fully evaluate the computational requirements and to enable HPC utilization at the extreme scale.

1. *Algorithm development needs.* New and continued algorithm development will be an ongoing need during the next decade. Algorithm development needs were identified that could address the following: dynamic (spatiotemporal) modeling; improved graph algorithms for network inference; statistical machine learning methods and dimension reduction methods for extremely high-dimension, molecular-scale data; integration of heterogeneous data types that span multiple scales of resolution; algorithms for thermodynamically informed biochemical reactions; agent-based simulation algorithms; and algorithms for high-resolution computational fluid dynamics.
2. *Software and library needs.* Software codes exist to address many of these scientific challenges; however, most have not been modified to take advantage of advanced hardware technologies such as accelerators. In a few cases, codes that are less CPU-intensive have not yet been parallelized, and improved data representation will be needed to exploit parallelism of codes. In addition, we identified a need for a common framework to visualize simulation results across the various scales: a “Google Earth” for biological and environmental simulations. This framework will require the development of novel data representations and visual analytical methods. We also identified a need for workflow management to provide connectivity between codes that perform simulations at different scales: atomistic/molecular, single-cell, pore, porous medium, field, regional, and global. Such workflow management will need to include data specifications at the code interfaces. Software libraries will need to provide numerical methods for: mixed-integer linear programming (MILP), nonlinear programming (NLP), coupled PDE-ODE-AEs (algebraic equations), and portable reaction solvers.
3. *Data analysis/accessibility needs.* Data-intensive applications such as graph-based methods and network abstractions will require scalable architectures to address the irregular data footprint, for example, rapid disk access (I/O) and data movement, the ability to perform *in situ* analysis on data (e.g., via in-memory processing), advanced methods for data reduction and compression, and benchmark data sets.

Table 3-1 provides a set of partially complete data with estimated computational needs in 2025. We note that, because many of the algorithms and HPC codes needed to achieve the objectives we have outlined are not yet mature (or, in many cases, do not yet exist), it is difficult to project accurately into the future.

[BER 3.1.4]**3.1.4 Biological Big Data Challenges**

“Big Data” is not a science area in itself, but rather a set of challenges and techniques used on biological data; we give a high-level overview of big data and how this relates to biology in Section 3.1.4.1. The primary drivers of big data in biology come from a vast array of instruments for examining biological data, each having its own data analysis challenges and science implications as described in Section 3.1.4.2. Once the data have been analyzed to remove errors and redundancy and to catalog and perform basis annotation, algorithms are used to cluster, reduce dimensionality, compute graphs of probabilistic dependencies, and generally find models of the data. This latter set of analyses is performed on a more abstract version of the data, typically represented as a matrix or graph, and is therefore independent of the specific input data type or format. These algorithms constitute the crosscutting themes in big data and are described in Section 3.1.4.3. The computational requirements of big data are still an active area of study and debate as described in Section 3.1.4.4.

[BER 3.1.4.2.3]**3.1.4.2.3 Other Sensors and Imaging*****Microscopy***

Microscopies for biology have shifted from “pictures” to “movies” — super-resolution light microscopy, serial crystallography at free electron lasers, and cryo-electron microscopy have all been dramatically advanced by the ability to record images at high frame rates. While firmware data reduction remains a desirable goal, the variability of methods and experiments means that leading-edge experiments will always seek to record complete raw data sets for offline analysis.

Very-high-speed scanning transmission electron microscopy detectors, analogous to X-ray techniques, will enable new imaging modalities by recording an N² pixel diffraction pattern at each scan point rather than simply a single pixel intensity. Ultimately, firmware data processing will reduce the data volume; however, learning those data reduction techniques will require high-end computing platforms. This objective is currently being explored at Lawrence Berkeley National Laboratory at the National Center for Electron Microscopy (NCEM) facility using NERSC for quasi-real-time processing. Special networking capabilities are required, including a 400-GB/sec router that will transport data from NCEM to NERSC, sending individual images to single cores. The goal is to implement analysis techniques wherein data are completely resident in memory and never need to be stored or retrieved from disk.

[BER 3.1.4.4.1]**3.1.4.4.1 Data Storage Needs**

The exponential growth in data from sequencing, MS, cryo-EM, and embedded environmental sensors makes it difficult to accurately predict data and computing requirements. There are already multi-petabyte data sets of genomics data; and based on historical data, we could expect an annual doubling of that data, resulting in exabyte-sized data sets by 2025. The addition of other devices for measuring biological data, both in raw and processed forms, suggests that an exabyte aggregate number is likely to be reached much sooner. As in other scientific domains, some of the observational data are irreplaceable, having been collected from a particular environment or time period that will not exist again naturally. Data sizes may vary significantly as the data are processed; for example, raw sequence data contain multiple reads of the same data, which can be compressed once the data are assembled or aligned. However, because there are many different analysis techniques that evolve over time, researchers may need the original data to allow for reanalysis.

The Joint Genome Institute

DOE's Joint Genome Institute is a raw data generator, as well as a large repository of genomic data, which is accessible to external collaborators through several web portals. The JGI runs several sequencers on a nearly 24/7 basis and has recently added two mass spectrometer machines for metabolomics analysis. Most of the JGI computing resources are co-located with the NERSC computing facility, and the sequencers and MS devices send their data to NERSC in real time. Scientists from around the world submit applications to have their samples sequenced and the resulting data processed at the JGI. Following are JGI-related metrics:

- 140 terabases of genomic sequence are generated on behalf of the BER community (in 2016, the number will be ~124 terabases).
- A total of 17 sequencers generate ~5 TB/day, and eight mass spectrometry systems generate ~100 GB/day.
- JGI features an 8,400-core cluster, 72 nodes that have more than 256 GB of memory, and one 2-TB node.
- There are 7.1 PB of IBM general parallel file system (GPFS) storage and 4 PB of tape storage in a high-performance storage system (HPSS) — the JAMO system has made it possible to reduce the need for larger file systems because most of the data that are reused can be stored on tape and retrieved in a matter of minutes. In addition, 400+ TB of data were downloaded by external users over the past year. Most of this data is restored from the tape system prior to download through the JGI Globus endpoint.

Other Biology Data Facilities

The Sequence Read Archive (SRA) at NCBI, a public archive, has more than 3 petabytes of raw metagenome sequence data, and the volume doubles every 11 months. The Beijing Genomics Institute (BGI) has a peak sequencing capacity of more than 16 TB/day, which is rapidly increasing with the advent of the BGI sequencer (based on the Complete Genomics technology). Numerous large-scale projects frequently operate at this peak capacity. Predicting the future scale of BGI is challenging, particularly now that mass spectrometry services are offered as part of the center's portfolio — certainly, the increase will be exponential in the coming years, and it seems inevitable that it will reach a level of petabytes per day by 2020.

[BER 3.2]

3.2 Climate and Environmental Sciences Division (CESD)

3.2.1 Atmospheric Simulation and Data Assimilation within the Earth System

3.2.1.1 Atmospheric Research/Simulation

3.2.1.1.1 Scientific Challenges and Opportunities

Understanding climate and the implications of energy use continues to be a primary concern of the U.S. Department of Energy, while our increasing computing capability continues to enhance our ability to simulate future climates and their implications. Toward this end, the Atmospheric Research Breakout discussed important areas where computing resources expected to be deployed in the next 5–10 years can open up new areas of possible research and help solve the current limitations of present Earth system models. A common theme was the handling of clouds and representing their impact within atmospheric models. Improving simulation of clouds within the Earth system models, such as ACME, will require integrating observations with a hierarchy of models ranging from direct numerical simulation (DNS) through computational fluid dynamics (CFD) models, large-eddy simulation (LES) models, regional models, and global climate models (GCMs).

Extreme Events

An area where we expect improved understanding enabled by larger computers is simulation of extreme events, such as the probability distribution of precipitation events. Current Earth system models used for decade-to-century-length simulations typically use grid spacings of around 1° , which is too coarse to resolve the small scales where heavy precipitation forms. As resolution increases, the Earth system models are better able to capture heavy precipitation events, which permits capturing the long tail of the probability distribution of precipitation, that is, the infrequent events that have important consequences, such as flooding. An example of this advance is the recent ability to begin capturing aspects of tropical cyclones in Earth system models with 0.25° grid spacing that are just beginning to be used more regularly (Reed et al. 2012). This grid spacing captures tropical cyclones sufficiently that we can now begin to examine their climatologies within the model. The expected decrease of grid spacing below 0.25° will enable better capturing of aspects of the tropical cyclones as well as convective events over both land and ocean. Fundamental improvements to the Earth system models will also further improve the representation of convection, such as development of resolution-aware convection and boundary layer parameterizations and the use of the quasi-3D multiscale modeling framework (Q3D-MMF) (Jung and Arakawa 2014).

Cloud Feedback and Climate Sensitivity

Cloud-climate feedback remains one of the largest uncertainties in Earth system models that affects the magnitude of simulated climate effects in response to external forcing (Sherwood et al. 2014). Clouds strongly modulate the energy balance of the Earth systems. They reflect solar radiation to cool the planet; they trap infrared radiation to warm the planet. The net cloud radiative effect depends on the temperature, altitude, and optical properties of clouds. How clouds respond to climate effects determines whether they will amplify (positive feedback) or mediate (negative feedback) climate effects in response to an external forcing such as greenhouse gases. The challenge in accurately simulating cloud feedbacks is that cloud systems span a large range of scales, of which many of the critical scales cannot be resolved by current models. Recent research has indicated that shallow convective clouds and marine boundary clouds play important roles in determining a model's cloud feedback. Shallow convective clouds have spatial scales of about 1 kilometer, while marine boundary clouds have sharp vertical gradients at their tops in temperature and moisture fields that need to be resolved with vertical resolution of several meters. There is therefore a significant gap between (1) the resolutions of current and near-term Earth system models at tens of kilometers in the horizontal and several hundred meters in the vertical; and (2) the necessary resolutions to accurately simulate the formation, maintenance, and dissipation of clouds that play crucial roles in determining the cloud feedback and the sensitivity of Earth system models.

Aerosol Forcing of Climate Variability and Interaction with Precipitation

The extent and types of aerosols in the atmosphere have changed greatly as a result of anthropogenic land use change, urbanization, and use of fossil fuels. Aerosols directly affect the transfer of solar and infrared radiation and thus the energy budget of the atmosphere. Aerosols also provide the nuclei for water vapor in the atmosphere to condense to liquid particles or freeze to ice particles. They therefore indirectly affect the energy balance of the Earth through their impact on the number and size distributions of cloud particles. Through clouds, they also affect precipitation processes. In the last several decades, the direct and indirect anthropogenic effects of aerosols on radiation at regional scales may have been larger than the greenhouse effect of anthropogenic carbon dioxide. Current Earth system models differ severalfold in simulating the indirect effect of aerosol on radiation (Shindell et al. 2013). The state-of-the-art Earth system models parameterize the aerosol properties and their interactions with cloud particles only by tracking the total mass and

number in a small number of aerosol types. To accurately simulate the direct and indirect effects of aerosols and their impact on precipitation, Earth system models need to calculate the time evolution of the number, size distribution, and chemical and physical properties of the dominant types of aerosols, as well as their size-dependent interactions with cloud and precipitation particles.

Land-Atmosphere Interactions

Interaction between the atmosphere and land is also a critical area for climate research. It affects both cloud characteristics as well as serves as the lower boundary to the atmosphere. Fluxes of energy, water, trace gases, and aerosol (such as dust) all affect climate in important ways. Therefore, correctly representing these processes in Earth system models is a high priority. Increased understanding is needed of the fundamental science governing the processes, which will need to be incorporated into the parameterization of the processes affecting the transfers across the land-atmosphere boundary. Specific examples are the biogenic emissions of trace gases, which are important for the formation of secondary organic aerosol; the impact of heterogeneity in surface characteristics, which affects cloud characteristics; and how urban environments alter the weather, which is important both for how cities alter the climate as well as for implications for the world population where a majority live in cities.

[BER 3.2.1.1.4]

3.2.1.1.4 Computing Needs and Requirements

The climate and atmospheric research community is a mature user of high-performance computing with a long history of being at the forefront of taking advantage of computing to advance understanding of the atmosphere and future climate. Traditionally, Earth system modelers have been able to use the ever-growing computing capabilities effectively, and this usage success is expected to continue during the next 5–10 years. However, current and expected changes in computing hardware pose challenges that will need to be overcome. Of particular importance for climate is the rapidly increasing ratio of calculations to I/O bandwidth. The ability to perform calculations has increased significantly to date without concurrent increases in communication between nodes or to long-term storage. This ability is particularly important for climate, as opposed to some other heavy users of HPC such as computational chemistry, because the evolving time series during the model integration is as important, if not more so, than the final result at the last integration timestep. Integrating diagnostics, also referred to as *in-situ* analysis, within the Earth system simulations will become increasingly important to reduce output demands. Examples include instrumenting the models to output probability distributions of variables in addition to instantaneous or time-averaged values, and more tightly incorporating satellite and other instrument simulators to calculate observation-comparable diagnostics. The danger is that the inclusion of these more detailed diagnostics within the simulation will upset load balance and potentially affect overall performance given that satellite simulators would only be sampling a small portion of the domain while other portions of the domain will need to wait unless this step can be performed asynchronously. There is also the possibility that this approach will result in even greater data output rather than less, as many of the currently outputted details are needed for understanding overall model behavior, both for improving scientific understanding and for diagnosing problems in the simulations.

Optimizing the overall modeling workflow would also greatly benefit climate research. Discussion included quantifying the end-to-end lifecycle of Earth system and atmospheric modeling to incorporate efficiency of researcher time in addition to efficient use of the available resources. This improved efficiency would involve better optimizing usage of computers by allowing more slack in computer usage to reduce queue wait times, as well as providing queues that enable sufficiently quick turn-around time for model development and test purposes. Post-simulation analysis will also become more difficult as model sizes increase, which will require careful consideration of resources

devoted to this purpose that have sufficient I/O bandwidth but that will need less computational power than the machines used for generating the simulations.

Data sharing and archiving are also of critical importance to the climate community. Increasingly, journals (where results are published) and funding agencies (such as DOE) require archiving of results for many years, often beyond the lifetime of the projects that generate the data. This long-term liability is of critical importance for making science open and responsible to the funders that pay for the work. Researchers need facilities where the very large computational data sets can be stored and shared easily with the research community.

[BER 3.2.1.2]

3.2.1.2 Data Assimilation, Model Initialization, and Reanalysis

3.2.1.2.1 Scientific Challenges and Opportunities

Determining the hourly evolution of the Earth system — atmosphere, ocean, land, and ice — with quantified uncertainties from instrumental observations taken over the past two centuries is a key problem that can be advanced in the next 4–9 years with developments in computational capabilities, algorithms, and models and aided further by continued data rescue thanks to efforts such as the Atmospheric Circulation Reconstructions over the Earth (ACRE) (Allan et al. 2011; www.met-acre.org) and intensive observation facilities such as the DOE’s Atmospheric Radiation Measurement (ARM). Knowing the actual evolution of the climate and weather, particularly in their extreme ranges, is critical to understanding and predicting how these extremes may change as the composition of the atmosphere is altered from increasing greenhouse gases. An important opportunity is to assess Earth system models with respect to their ability to represent weather events as well as their fidelity in representing the probability distribution of the Earth system and extremes such as heat waves, cold spells, hurricanes, storm surges, hailstorms, and wind storms for as long a comparison period as possible. This aim is particularly important as variations in extremes, as well as in important climate phenomena, such as the Madden-Julian Oscillation, El Niño Southern Oscillation, and the Atlantic Meridional Overturning Circulation, may occur on decadal to multidecadal timescales, so the commonly used baseline of a 30-year period for model/observation comparison is insufficient (Sardeshmukh et al. 2015).

The most widely used technique for determining this evolution is data assimilation: forming the state of the system, the “analysis,” by optimally combining a model-generated, short-term (e.g., 6-hour) “first guess” with observations and then weighted by the uncertainty in each. The Ensemble Kalman Filter (EnKF) and the 4D-Variational assimilation algorithms have both been employed to provide subdaily atmospheric and land estimates (retrospective analyses or “reanalyses”) spanning more than 100 years using only sparse surface observations (Compo et al. 2011; Poli et al. 2016). Intensively observed areas, such as the ARM sites, use related techniques (Zhang et al. 2001; Xie et al. 2004) to fuse high-resolution observations into complete descriptions of the atmosphere and subsurface variability at the site. Many other techniques have been used to provide atmosphere, ocean, and land estimates focused on the satellite, conventional upper-air, and ocean observing system eras (see, e.g., Reanalyses.org for comprehensive lists and references).

A key challenge in this area is determining the uncertainty of the estimates of the climate and weather states spanning the instrumental record (i.e., from the nineteenth to the twenty-first century). This quantified uncertainty should include uncertainties arising from the assimilated observations, whether surface, subsurface, upper-air, global positioning system (GPS), satellite radiance, or gravity, as well as include uncertainties arising from the nonlinear equations describing the Earth system and uncertainties from errors in the representation of those equations.

Another challenge is to utilize data assimilation to better understand, diagnose, and improve Earth system models by estimating the model error identified from consistent differences between the short-term, first-guess forecast and observations. Differences with forecasts out to several days can also be assessed. The systematic differences are associated with “fast physics” and will point to areas where the model can be improved (Klinker and Sardeshmukh 1992; Phillips et al. 2004; Rodwell and Palmer 2007).

Additional opportunities and challenges are related to the consistency of the state estimates. A key challenge is to avoid spurious drifts of coupled models from the assimilated initial state over a period of integration of several years. In a similar vein, avoiding spurious jumps in reanalysis records that occur as the observing system changes dramatically from the period before satellite observations to that after remains a difficult issue (Compo et al. 2016).

New opportunities are arising from using data assimilation and the rich observational collections provided by intensive observations, such as from the ARM sites, to diagnose physical processes and create more complete four-dimensional data sets for model evaluation. Data assimilation also provides a physically consistent method to use high-resolution observations to constrain or estimate non-observed quantities, such as the global carbon and water cycles.

[BER 3.2.3.3]

3.2.3.3 Cross-Cutting Research Directions

The community should strongly consider a wide range of approaches for the *in-situ*, real-time processing of ensemble Earth system simulation. *In-situ* data analysis will alleviate I/O bottlenecks that are likely to develop with exascale computing as the relative energy cost of moving data off-chip continues to grow. *In-situ* data analysis also creates opportunities to expose additional parallelism and, thereby, more efficiently exploit the large core counts accompanying exascale machines. The community should recognize that much of the analysis capability embedded within ESMs is often not exploited because analysis methods are typically not as scalable as the forward model and, as a result, can degrade simulation throughput. New programming models and system support are needed for internode memory copy beyond that available from MPI. Such technologies would promote the development and deployment of diverse analysis and visualization tools that can digest and enhance the value of ensemble climate simulation data in real time.

The primary computational cost of atmosphere chemistry and marine biogeochemistry is the advection (or transport) of trace constituents by the fluid motion. Physical constraints related to conservation of tracer mass and monotonicity of tracer concentration under the process of advection results in relatively expensive tracer transport algorithms. As a result, valuable information about chemical processes in the atmosphere and ocean is not obtained because transporting these trace constituents can increase the total model cost by 3 to 10 times depending on the number of constituents. While continued research into accelerating traditional models for tracer transport should be supported, additional and more novel lines of research should also be considered, such as separating tracer transport into its own stand-alone executable in order to expose additional parallelism and not hinder the simulated-year-per-day throughput metric.

Over the entire history within ESMs, couplers have been primarily data managers that focus on the mapping and conservation of state and flux variables from one physical component on one grid to another physical component on another grid. Data are passed to and from each model “component” of the coupled system in a sequential, time-lagged manner. The sequential coupling strategy is not rigorously correct from a computational physics perspective (i.e., guarantees of numerical convergence are lacking). This lack of rigor can manifest as instabilities in the simulations that have to be managed with ad hoc coupling strategies (this topic of coupling ESMs is also discussed in Section 3.2.4.3, Cross-Cutting Research Directions, on page 83). The exploration of coupled

system dynamics exhibiting large system stiffness that requires coupled, implicit solvers is all but impossible within the current coupling framework. The Earth system modeling community should seek out coupling approaches that can instantiate mathematically rigorous methods that guarantee numerically consistent and convergent simulations and can be deployed on DOE LCFs with the same computational scaling and efficiency as present-day coupling strategies. Success will require the adoption and tailoring of advanced computational science approaches for data management, dependency graph abstractions, and automated process coupling.

Exascale computing is, more likely than not, bringing with it a hierarchical computing environment. Exploiting current on-the-floor realizations of this hierarchical computing environment has proven to be a challenge for ESM. Looking forward, we are anticipating the deployment of fine-scale simulators on the accelerators with coarse-scale simulators residing on the traditional CPU. While the atmosphere “super-parameterization” is one such example of multiscale simulation, we expect that many more components of ESM could benefit from an embedded, multiscale approach.

More broadly, the community should make substantial investments in the development, testing, and deployment of new programming models. All aspects of the ESM enterprise have opportunities to benefit strongly from these new approaches. New programming models could enable more computationally robust model coupling, better utilization of hierarchical computing, improved task management, and recovery from system faults. Maybe more important than any of these technical aspects, new programming models will act to entrain the most talented computational scientists into the Earth system modeling activity (see Section 3.2.3.4, Computing Needs and Requirements).

[BER 3.2.3.4]

3.2.3.4 Computing Needs and Requirements

The community continues to grapple with the trade-offs between capacity computing and capability computing. Particularly within DOE, we are charged with demonstrating a continual growth in capability. This charge is appropriate, and the international modeling community reaps its benefit. But we need to recognize that the grand challenge simulation that demonstrates capability occurs only after a long sequence of low- and intermediate-scale simulations. Across the exascale computing initiative, all the way from procurement to batch queueing, we need to recognize that capacity computing provides the essential support needed to demonstrate capability. Beyond this supporting role, capacity computing is the workhorse supporting the vast majority of science inquiry.

The long-term health of the Earth system modeling community depends on moving away from a reliance on Fortran paired with MPI+X (message passing interface extension). While the importance of climate variability can and should continue to attract the world’s most talented computational scientists, the reliance on a somewhat outdated language in numerous legacy code bases will act to deter next-generation computer scientists. We all understand and appreciate that we have important science questions that must be answered today with code bases that exist today; however, we also have to recognize that we will have important science questions 10 years from now. Investing in programming models and languages used and developed by the broader computational science community will energize the Earth system modeling effort in the decade to come.

Back end data analysis and management needs should be considered alongside the computing requirements. We should strive to develop analysis systems where the location of the data products is transparent to the scientist. Adopting this approach will require a significant transformation in both computing paradigms and community workflows. While the Earth System Grid and the underlying Globus facility are attempting to move beyond the need to co-locate data set production with data set analyses, the analysis tools are, for the most part, not yet up to the task.

[BER 3.2.4]

3.2.4 Coupled System Integration — Earth System Models

3.2.4.1 Scientific Challenges and Opportunities

Fully coupled ESMs integrate the physical and biogeochemical components of the Earth's climate to capture the many feedbacks in this complex system. Resolving processes at relevant space and time scales and providing decision-relevant information are driving requirements for very high spatial resolution and an increased use of integrated ensembles of simulations that can only be enabled by exascale computing systems.

Sea-level rise due to melting of large ice sheets in Greenland and Antarctica is one of the significant impacts of climate effects. Quantifying the rate of sea-level rise and understanding the coastal impacts require integration across Earth system components, as well as high spatial resolution. Modeling the possibility of the rapid collapse of the Antarctic ice sheet requires the coupling of ocean and land ice with a spatial resolution of ~100 m near the grounding line to capture dynamic processes (this topic of the cryosphere in Antarctica is also discussed in Section 3.2.3.1 on page 75). Similarly, modeling the coastal impacts, flooding and inundation, will require achieving ~1-km levels of resolution and capturing cyclones and other events that contribute to storm surges on top of the mean sea-level rise. Century-scale or longer integrations are required to understand the stability of Greenland and past ice sheet behavior in the paleoclimate record.

Most of the economic impacts of climate effects result from extreme weather events, including severe storms, drought, heat waves, and extreme precipitation events. Changes in the frequency of these events result from changes in the hydrological cycle and global circulation patterns. Generating climate statistics requires integrated ensembles of ESMs to generate probability distributions of such weather events. In addition, high spatial resolution (~1 km) is required to resolve cloud and convective processes. Accounting for the impacts on water use and availability also requires watershed-scale resolution and the inclusion of subsurface hydrology.

Another important direction for ESMs is an increased focus on biogeochemical exchange and issues in atmospheric chemistry and aerosol/cloud interactions. Better simulations of climate sensitivity require an understanding of how land and ocean ecosystems sequester carbon and how vegetation changes in response to the physical climate. Aerosol exchange between land/ocean and the atmosphere influences cloudiness, precipitation, and albedo, while aerosol deposition of dust and black carbon supplies nutrients to the ocean and darkens the albedo of ice surfaces. Biogeochemical simulations bring their own challenges, with the need to transport a large number of tracers and manage reactions among species. In addition, biogeochemical models include long timescales and require long integrations to create equilibrated initial states.

Climate projections with quantified uncertainties remain a persistent challenge to computational science. Error bounds on future climate impacts are needed by decision makers charged with mitigating and adapting to climate variability. However, traditional uncertainty quantification methods typically require large ensembles of simulations to explore the uncertainty space. There is a trade-off between devoting computing resources to (1) a few very high-resolution, high-fidelity simulations that attempt to minimize prediction bias and (2) a larger number of lower-fidelity simulations that can better characterize uncertainty. Understanding the natural variability in the Earth system and decadal predictions with data assimilation may require ensembles over varying initial conditions.

[BER 3.2.4.4]**3.2.4.4 Computing Needs and Requirements**

Computing facilities must be able to support “hero”-class simulations that occupy a significant fraction of the machine, as well as ensembles of simulations under different initial or boundary conditions, parameter settings, etc. Ensembles are embarrassingly parallel and could be run as independent processes, but increasingly, they will be embedded within a single simulation to generate statistics and analysis or potentially guide the ensemble as it progresses.

In addition, a substantial fraction of the workflow in Earth system simulations, even as we prepare grand challenge simulations, requires moderate-size facilities with rapid turnaround. Overnight turnaround is required for testing new developments and for tuning model configurations. Because leadership-class facilities focus their environment on leadership-class simulations, development time is often difficult to obtain, particularly for multi-lab projects like Earth system simulations, where obtaining access to other institutional resources is difficult. Similarly, as the climate community invests in a more substantial testing infrastructure, we are not able to run test suites (including nightly regression tests) on the target leadership class architecture and software environment, as queuing policy does not allow for routine testing. Related to this problem is the need for a stable software stack. Frequent improvements to the software stack on advanced architectures are necessary, but often break application codes and disrupt production schedules. Alternative resource management, containerized environments, and other strategies are needed to support robust testing and stability for application codes.

As exascale systems become more vulnerable to bit errors, fault tolerance mechanisms will need to be developed. It may become necessary to aim for detection of errors, if possible, rather than correction. If errors go uncorrected, the simulation effectively becomes nondeterministic, and new methods will be needed to validate the reproducibility of model results to acceptable tolerances. Continued support and I/O infrastructure for checkpoint/restart will still be needed for the foreseeable future.

Applications must be portable with reasonable effort across the diverse heterogeneous architectures likely to exist at the exascale. While the MPI+X programming model appears to be the most likely target through the next generation of architectures, new portable programming models will likely be needed beyond that time to address issues related to managing memory and fault tolerance and increasing scalability. Work must begin now to develop and explore candidate programming models before exascale machines are deployed.

Climate data management and analytics will be an increasing challenge at the exascale. In-situ analysis will become more necessary given the limited ability to save the output to disk, yet it is difficult to anticipate what kind of analysis users will need to perform. A supported Fortran library of common statistical summary routines may suffice for many scientific users but cannot accommodate the research frontier of big data analytics and therefore risks stagnation. Software ecosystems are available for data analytics, and communities have been trained to use such software in machine learning, statistics, model reduction, etc.; however, this software is implemented in modern programming languages that cannot interface with existing ESMs. New intrusive methods for uncertainty quantification may require access to state variables at the time-step level or to the results of intermediate computations within the call graphs, or even the ability to modify state variables. Climate data must eventually be shared with a broad community of stakeholders; however, new data sets with regional information will likely create some barriers to a typical end user because of storage capability and bandwidth. Server-side analysis is attractive,

yet leadership computing facilities would have difficulty permitting arbitrary user code to be executed by unauthenticated users. Data analytics research could be performed at a smaller scale on more flexible institutional computing resources but cannot take advantage of relevant software ecosystems without hooks into the model. Sponsors and scientific publications are increasingly requiring longer-term archiving of some subset of simulation data that may not be compatible with current computing center policies and/or institutional capabilities.

Computationally focused workforce development is needed for the user community. There is inadequate staff experience with HPC programming, software engineering, and big data. It is difficult to recruit staff against industry competition; training (and time away from deliverables for training) will be needed, as well as software frameworks that abstract architecture specifics. Reward structures for personnel with a strong software focus are also needed as these staff are critical to the success of any HPC code development group, yet they are still judged with the same metrics as scientific staff. Proposals also tend to favor research-oriented work rather than the day-to-day software work required for the development, testing, and performance optimization of application codes.

[BER 3.2.5]

3.2.5 Integrated Assessment Modeling

3.2.5.1 Scientific Challenges and Opportunities

The overarching scientific challenges in the field of integrated assessment are to (1) understand the implications of climate effects on human systems and vice versa, and (2) quantify the uncertainty surrounding human-Earth system interactions.

3.2.5.1.1 Human-Climate Interactions

Although integrated assessment models (IAMs) are designed to capture interactions between human and Earth systems, these models have largely focused on understanding the implications of human activity on the climate system at the global scale (Clarke et al. 2014). As a result, these models operate at relatively coarse spatial and temporal resolutions and have limited inclusion of the impacts of climate variability. However, the questions we are asking more recently require higher resolution and often the inclusion of new processes. In addition to enhancing existing models, IAMs are increasingly coupling to other models, both Earth system models (e.g., Collins et al. 2015) and impacts, adaptation, and vulnerability models (e.g., Kraucunas et al. 2015). These model couplings present a variety of theoretical, operational, and computational challenges as they include multiple scales, multiple processes, multiple disciplines, multiple institutions, and multiple sets of heterogeneous data. In addition, as different questions may require different suites of models, the coupling infrastructure needs to be flexible, modular, and extensible.

3.2.5.1.2 Uncertainty

The systems addressed within coupled human-Earth system models of this type are inherently uncertain. In addition to uncertainty regarding how the environment will react to anthropogenic activities, there is significant uncertainty in the activities themselves. This latter type of uncertainty presents additional challenges in that it is not independent or governed by known physical laws. Humans make choices that affect their environments, and the decision-making criteria are not always well understood. Quantifying this uncertainty is necessary and will require large ensembles testing the effects of parametric, structural, and scenario uncertainty. Different users will have different requirements, affecting both the design of the uncertainty quantification and the presentation of its results. Possible user requirements include (1) predicting a variable of interest, (2) generating a probability density function around that variable, and (3) understanding the implications of our actions on that variable.

[BER 3.2.5.4.5]**3.2.5.4.5 Data**

New methods of developing, storing, and analyzing data are also required, particularly because of the diverse, heterogeneous, and large-scale nature of data — factors that present complex challenges for integrating data with the model development process. Meta-data will become increasingly important for understanding the provenance, credibility, and geographic information related to the data. Artificial intelligence and data mining techniques could aid in processing and assimilating large suites of input and output data. Spatial statistics and scale translation tools are also necessary.

[BER 3.2.6] / [BER 3.2.6.1.1]**3.2.6 Transforming Science through Exascale Capabilities: Model-Data Fusion and Testbeds****3.2.6.1 Large-Scale Heterogeneous Data Management****3.2.6.1.1 Scientific Challenges and Opportunities**

Large-scale heterogeneous data management begins with the conception of private or shared factual information (such as numbers, records, documents, files, etc.) and spans every aspect of the high and low ends of the computational and data ecosystems. In other words, whether execution begins at HPC or ends at the desktop, the management and organization of large-scale heterogeneous data must be tracked and managed for its entire lifecycle. The grand science questions in climate demand a unified data capability that is not possible today, that is, a data ecosystem that includes the concepts of:

- **Critical complex data-generating systems:** high-end supercomputers, clusters, and computer servers to sensitive environmental detectors, lab analyses, and orbiting satellites;
- **Data collection and management:** for organization and easy user discovery and accessibility;
- **Data analytics:** for pattern discovery, structure identification, dimension reduction, image processing, machine learning, and exploratory visualization anywhere throughout the data lifecycle;
- **Data-intensive computing:** for describing applications that are input-/output-bound and enabling large and complex data manipulations both remotely and locally (including in situ analytics); and
- **Decision control:** for knowledge discovery breakthroughs.

The computational and data ecosystems must include a pervasive provenance capture throughout. Data from the “critical complex data-generating systems” are housed and securely managed at many worldwide sites. Local and remote computation is necessary, as the increasing data size and algorithm complexity is leading to more data-intensive and compute-intensive user requests. For data backup with easier data access, the network must be able to move petabytes/exabytes of data between computing/data centers. Finally, analytical modeling of the computational/data ecosystems assists users in making smart choices in managing and using community resources for moving and computing large-scale data.

We recognize that with fast development of large complex systems, issues with resiliency arise. We are also aware that resilience is a systems problem, not an individual component problem, and requires a systems approach. Therefore, we will need to ensure that when a user runs an end to-end workflow, it will engage many different components (i.e., either a component run to completion or the user receives a meaningful error response). Consequently, steps must be taken to maximize resiliency, and we plan to continue climate research efforts in these key areas:

1. **Standards and protocols.** Community-developed software, standards, protocols, and techniques for hardware, network, and software architecture design must be leveraged for data and computational analysis. Examples include the following:
 - a. The Climate Forecast (CF) and visualization output (e.g., portable network graphics [PNG]) data conventions for data archiving, cataloging, analysis, and discovery.
 - b. The Universal Web Processing Services (WPS) application programming interfaces (APIs) to provide well-formed communication points for disparate components.
 - c. The OAuth2 security protocol that creates a universal authentication environment, bringing all security features utilized by each component into one common framework.
 - d. Usage of well-established, open-source software with existing strong community support wherever possible.
 - e. Engagement in community efforts to develop standards and protocols, such as the Open Geospatial Consortium (OGC), Research Data Alliance (RDA), National Institute of Standards and Technology (NIST), etc.
2. **Unit testing.** Unit testing of key individual hardware, network, and software components contained within the computational and data workflow must be increased, including automated tests and nightly builds.
3. **Regression testing.** Full system, multi-component, and regression testing must be created for critical use-cases that are frequently encountered by scientist (and perhaps nonscientist) end users.
4. **Hardening.** The hardening of current hardware, network, and software features frequently used by end-users must be prioritized over the development of new features used by only a few.

Today, the end-to-end computational and data workflow is composed of multiple components working together to create a unique process that must continue to be integrated for DOE's Accelerated Climate Model for Energy project, which is scheduled to run at the three ASCR computing facilities (ALCF, NERSC, and OLCF). Today, key components of the petascale heterogeneous computational and data ecosystem include these:

1. The *Earth System Grid Federation (ESGF)* enterprise system for ACME data storage, cataloging, and sharing. ESGF employs a decentralized, peer-to-peer architecture using modular components and standard protocols, which help increase resiliency. This robust design relieves the system from single points of failure.
2. The *Ultrascale Visualization Climate Data Analysis Toolkit (UV-CDAT)* for visual data exploration and analysis. UV-CDAT is the first successfully designed system to run unrelated analysis and visualization tools and techniques while capturing independent workflows and provenance for fault tolerance and reproducibility. In addition, the scripting and command line interface is utilized by Metrics (i.e., the ACME diagnostics suite) and other tools needed for ACME.
3. The *ACME Metrics and the Exploratory Analysis (EA) Classic Viewer* for model diagnostics generation and Web-based visualization, which leverages UV-CDAT components as an efficient back-end processing engine and Django as a Web application framework.
4. *GridFTP servers.* Although many sites deploy and run GridFTP servers, different sites varied in terms of the effort applied to managing and monitoring those deployments. Globus has addressed the problem of data transfer services, performing a number of tasks that are relevant to ACME data workflow.
5. *Velo and Pegasus* for model run configuration, build, and runtime output capture and storage. Together, they enable the capture of sufficient information so scientists can reproduce previous

ACME runs. Furthermore, these tools extend the provenance format so that it can also capture and link to performance information for specific workflows and model runs to enable in-depth performance analysis. Velo and Pegasus together support the creation, submission, and monitoring of ACME runs on any system available to the ACME project via the ASCR compute facilities. In addition, Velo will be the collection point for all provenance information. The ACME workflow is extended to capture the relevant reproducibility and performance provenance information.

Many of these components are maintained by a broader development community and often contain a rich suite of testing frameworks and methodologies. For example, the ESGF data publishing utility offers a variety of data publishing tests that ensure that configuration of a data resource is being published correctly.

[BER 3.2.6.2.3]

3.2.6.2.3 Cross-Cutting Research Directions

In the world of ever-expanding observing and modeling facilities in terms of complexity, sophistication, and capabilities, the issue of continuous training and development of the workforce is central to realizing an optimum utilization of the scientific and technical tools available. There are two key areas where enhanced learning and training are needed in the area of observational data processing. First, the DOE climate research facilities operate sophisticated, state-of-the-art instrumentation (e.g., millimeter wavelength radar, lidars, radiometers). Similar instruments (especially mm-wavelength radars) are not operated by university-based research groups, and thus, these institutions cannot train graduate students who have expertise with and training on these sophisticated sensors. The second area where training and workforce development are needed is in the area of data analytics and machine learning methods. Graduate student internships and summer workshops are good mechanisms to promote workforce development. The development of a “sabbatical” concept for DOE laboratory employees should also be considered.

We need to better identify the user needs and their requirements for data analytics and visualization. A workshop focused on addressing these issues that brings together observational data producers, users, and data and computations scientists could foster development of a set of common tools and standards. For example, this workgroup could develop standards for data descriptors that extend beyond data quality, and it could also develop standards for instrument network simulators and visualization capabilities that can operate on both model and observational data sets. Funding to support such cross-cutting partnerships between science users, data scientists, and instrument engineers is needed.

[BER 3.2.7.2]

3.2.7.2 Programming Languages and Programming Models

3.2.7.2.1 Scientific Challenges and Opportunities

Within the climate community, simulation is dominated by the nexus of the Fortran/C/C++ languages and the OpenMP/OpenACC programming models. The legacy inertia behind Fortran, coupled with perceptions that Fortran compilers deliver superior performance, ensures its continued use. Conversely, the broad computing community’s use of C++ coupled with its access to novel and emerging architectures and new language constructs has resulted in a gradual shift to C++ as the basis for simulation. Both Fortran and C/C++ are interoperable on almost all computing platforms available today. However, it became clear that there is a strong desire for other productivity-oriented scripting languages in a distributed environment (e.g., Python/pyMPI). Although such demands run counter to the developments in processor architecture, there is a strong reluctance to abandon them.

Concurrent with a language switch is the desire to move away from the flat MPI model, in which there is one MPI process per core and no ability to exploit GPUs. Although this flat MPI model may be functional on the next generation of systems, it will likely underperform on Xeon Phi-based systems (Cori, Theta, Aurora) and will lose a factor of 10 to 50 times on GPU-accelerated systems (Summit). To that end, pragma-based, hybrid programming models are currently being explored in the climate community. These include MPI+OpenMP and MPI+OpenACC. Predominantly, the functionality afforded by MPI 2 is used, whereas on CPUs (including Xeon Phis), the functionality afforded by OpenMP 3.1 is used. Although some aspects of OpenMP 4 are currently being evaluated (SIMD clauses), OpenMP support for GPUs lags behind that afforded by OpenACC 2. Thus, while research efforts use the latest bleeding-edge, compute-unified device architecture (CUDA) to fully exploit GPUs, production computing uses the more general OpenACC with the hope that future versions of OpenMP will provide the requisite functionality and performance. This drive to exploit new architectures via OpenMP and OpenACC is tempered by individual experiences in which substantial effort may be expended with the hope of attaining massive speedups. Unfortunately, the realities of computer architecture resulted in moderate speedups of less than four times previous levels — a clear indication of the need for tighter computer science-computational science collaborations that first set realistic expectations. Similarly, substantial effort can be consumed attempting to manage the data locality challenges inherited in hierarchical memories via OpenACC/OpenMP while newer, vendor-driven architectural and runtime constructs obviate these concerns.

Orthogonal to the pragma-based imperative OpenMP/ACC approaches to exploiting on-node parallelism are the more declarative approaches afforded by domain-specific embedded languages (DSeLs) and C++ templates (e.g., Kokkos). In both cases, researchers write legal C++ code that could be compiled with any compliant compiler. In the DSeL approach, the researcher exploits a source-to-source compiler (e.g., ROSE, CHiLL) to transform and optimize the code cognizant of the underlying domain (e.g., stencils on structured grids) and to generate code for a variety of target architectures (CPUs, GPUs, etc.). The C++ template/Kokkos approach expresses computations as a series of composable parallel constructs (parallel, reductions, scans, etc.) but requires the core Kokkos developers to map these templates onto CUDA, OpenMP, or OpenACC constructs. As C++ evolves, it may subsume these constructs and supplant the need for the pragma-based OpenMP and OpenACC approaches and allow vendor compilers to provide very aggressive optimizations for their respective architectures.

The static domain decomposition of data and computation afforded by MPI can be augmented or replaced with either partitioned global address space (PGAS) models and languages such as UPC, UPC++, and CAF (Coarray Fortran) or task-based models such as Legion. In a PGAS language such as UPC, one can construct a global, shared-memory construct spanning all of the nodes in a supercomputer (petabytes of shared memory). Processes can allocate structures in this memory and access with reads, writes, or atomic operations. Such facilities are particularly useful across a range of BER domains, including genome alignment, adaptive mesh refinement, and potentially with couplers in climate science as they obviate the bulk synchronous and collective-based operations often used in MPI as well as costly demands on hardware vendors and facilities for high memory capacity. Whereas UPC programs are often written as single-program, multiple-data (SPMD), Legion offers tools for constructing distributed, task-based execution in which data and computation can be decoupled, thereby facilitating finer-grained execution than users may naturally write in MPI. Although PGAS and task-based models focus on distributed-memory computations, they remain interoperable with MPI, OpenMP, and OpenACC.

[BER 4]**4 PATH FORWARD**

The wide range of scientific research directions described herein report a common challenge of integrating heterogeneous, distributed, complex, and large amounts of data with the development of underlying models that are needed to carry out descriptive and predictive modeling and simulation. This challenge applies to both the biological systems and environmental science mission areas, and requires support from a computational ecosystem that has a wide range of capabilities for large-scale computing and data analytics and management. Making progress on the research directions described in this document requires making advances in data analytics, assimilation, curation and annotation, complex workflows, and flexible computer access policies to support fast turn-around for development and testing in addition to large-scale production runs. A common theme is the need for integration of experimental and computational capabilities.

The deployment of effective computational infrastructures that enable new scientific advances requires collaboration between BER and ASCR scientists and facilities. The specific requirements for such a computational ecosystem can be categorized in the broad areas of methods development, computational environment, data, and communication and community involvement.

[BER 4.1]**4.1 Methods Development**

Current advances in computing technologies are expected to result in unique exascale computational resources well within the next decade, and will make petascale resources readily available. To make the scientific breakthroughs that these unprecedented, highly complex computer resources enable, it will be necessary to develop accurate physical models and highly scalable algorithms for performant and portable applications that integrate (1) physical measurements with associated uncertainties, models, and parameters from large-scale heterogeneous data sources; and (2) theoretical and mathematical approaches capable of representing multiscale, multiphysics system descriptions. A key capability required is the combination of less-structured, machine-learning approaches with deterministic models based on rigorous physical theory. These generic needs have particular specializations for the two divisions of BER. BSSD has specific challenges in deriving biological function scaling from atomic representations of macromolecules through the networks of thousands of chemical reactions that drive the physiologies of single cells to ecological models of whole biomes. CESD has, in general, a more physically coherent framework for representing the fluid and transport dynamics for water and atmosphere, physical inputs to these dynamics through models of reflectance and absorbance, and activity of various biotic and abiotic participants in land and air geographies. Here the scales may range from porous flow around soil particles up through differential heating due to the albedo of large-scale land formations and atmospheric gas-exchange with soil and plants, to large-scale eddies in atmosphere affecting and affected by cloud formation. The review participants articulated a range of method development requirements that include:

- Methods for conversion of primary measurement to physical quantities for training and comparison to models of function and dynamics, with integrated uncertainty quantification and analysis. Examples include preprocessing and analysis capabilities for large-volume data from visible light, X-ray and electron imaging, super-resolution light microscopy and serial crystallography, mass spectrometry, and liquid and gas chromatography.
- Methods for integration of these diverse, heterogeneous measurements on both natural and simulated systems to infer models and parameters with rigorously characterized uncertainty. Examples of these include new algorithms to accelerate accurate metrology of new experimental methods; methods of deconvolving and quantifying individual biological/chemical contributions to mass-spectral and optical hyperspectral data; preprocessing algorithms for large-volume

data; data analytical methods; functional inference from biochemical, genomic, and structural measurements; methods for large-scale discovery of associations and causal interactions among measured objects and new numerical algorithms for multiscale multiphysics modeling and simulation; and methods that integrate holistic information about interacting systems and algorithms capable of identifying individual events and behaviors within large datasets and the ability to capture “semi-resolved” behaviors.

- Improved core algorithms for particular physical representations of dynamics. Examples include the ability to identify functional dynamics of systems at successively linked system scales; development, implementation, and testing of ensemble methods; hybrid and multiphysics approaches; and hierarchical applications capable of passing relevant parameters and data between successive methods.
- Discipline-specific methods for modeling key biological and environment systems. Examples include methods for modeling cellular and organismal growth, fitness, and ecology; hybrid modeling of communities of organisms at multiple scales from the cellular level, through populations, to full ecologies combining models of different levels of phenomenological, statistical, and biophysical abstraction; methods for multinetwork and whole-cell modeling; multiscale approaches for deriving physical descriptions of whole cells enabling biophysical simulations; and advanced, implicit algorithms for the coupling between Earth system components with improved methods for parameterization and modeling of individual components, such as cloud-resolving models, beyond solely increasing the model resolution.

Core themes are the requirements for methods for rapidly turning large-scale primary measurements into physical quantities as input to and comparison for models of biological and environmental function and dynamics, and statistical approaches for inferring and training models at various scales from heterogeneous data. Each of these needs include theory development as well as algorithm development. Some of these require that existing algorithms be adapted to new architectures. Many also represent a growing need for tight coupling to data systems that support modeling and analysis. It was also noted how important it is that these algorithms become accessible to the largest number of scientists possible by supporting the open “hardening” of these tools into well-supported, distributable, reusable, modular software components that can be used and extended by the community. The use of common software stacks and engineering standards will also enable a larger community to help develop these tools for broad utility and stability.

[BER 4.2]

4.2 Computational Environment and Resources

Leadership-class facilities focus their environment on large production-class simulations and data analysis, but there are other requirements of the computational ecosystem, in particular, to support algorithm and application development and persistent support for access and transport of large data sets. The attendees expressed a need for access to a stable software stack on advanced architectures to enable development in an environment with minimal disruption. Of particular importance is the rapidly increasing ratio of calculations to I/O bandwidth. The ability to perform calculations has increased significantly to date without concurrent increases in communication between nodes or to long-term storage. A common aspect of the scientific approaches discussed is the integration of observational and experimental data in the development of models used in the descriptive and predictive computational modeling and simulation. The requirements for the computational environment include the following:

- A software stack for analysis, machine learning, and visualization of data that allows low-latency access, rapid data migration, and integration across multiple sites and platforms including seamless, fully integrated access between experimental and extreme-scale supercomputing facilities.

- Capabilities to use systems in weakly coupled parallel modes and to support statistical and Markov chain methods, large-scale parameter sweeps, or optimization.
- Interactive access to testbed systems to support preparation of large-scale modeling and exploratory data analytics, with support for common, tested engineering frameworks for code development and execution under different software development approaches and environments; strategies for testing of robustness, stability, portability, resilience, fault tolerance, and adaptability of application codes; and productivity-oriented languages in distributed environments.
- Scheduling tools and policies for optimized usage of computers by allowing more slack in computer usage to reduce queue wait times, as well as providing queues that enable sufficiently quick turn-around time for model development and test purposes, which will improve researcher efficiency.
- Support models and tools for specific community codes that include implementation and optimizations on new architectures; effective use of high-bandwidth, low-latency communication between systems components on exascale computers; and support for memory- and communication-constrained applications.

A few of the key challenges are deep integration with and low-latency access to large-scale, deeply structured physical data for analysis; model training and comparison; provision of a development and testing environment that is representative of how large-scale codes will run in production; provision of a software stack and engineering standards that provide effective community building; extension and use of complex codes and support for the modularity of such codes; provision of resources for both tightly coupled and easily distributable computations on the same architecture; and support for different configurations of memory, communication latency, and disk to support diverse applications from genomics to large Earth system simulation.

[BER 4.3]

4.3 Data

A common discussion topic is the integration of large-scale observational and experimental data into the development of physical models. Increasingly, the BER community is depending on a large portfolio of programs and facilities that are generating huge amounts of data, sometimes continuously and in real time. These data are constantly ingested for analysis by and to confirm predictions of algorithms that support understanding everything from protein function up to full atmospheric dynamics. One of the largest challenges for the next generation of computing architectures is to provide the means of capturing, representing, and providing low-latency access to these exceptionally large, highly structured data sets that inform both statistical and physical models of biological and climate/Earth systems. Such data, as well as the data generated from the modeling and simulation, needs to be made available to the wider scientific community in a way that preserves provenance, metadata, and annotation. This effort requires long-term and large-scale storage of and accessibility to both primary and derived data types from geographically distributed sources. The increasing need for and dependence on ensemble simulations leads to a different set of data challenges related to approaches involving outputting results from every ensemble member and developing the ensemble statistics after the simulations are complete. Requirements for the data-related aspects of the computational environment for data include:

- Mirrored data facilities, which will be needed to ease the difficulty of large-scale data transport from a single site, with synchronization mechanisms and database and data search systems that are optimized for the diverse data types of physical data, functional data, spatially indexed data, and complex relationships among data. Enabling efficient queries and comparisons will require

new capabilities in coherent, well-maintained data representation; ontology; metadata and provenance preservation; and transport formats.

- New exploration, analysis, and visualization tools, which will be needed to enable the integrated analysis and comparison of data from multiple modalities and across experimental conditions and environments, and the assimilation of data at the appropriate spatial scales with the appropriate targeted observations. This capability will address the specific science needs of the Earth system models, including extreme weather events, land-atmosphere interaction, cloud-climate interactions, and aerosol-cloud-precipitation interactions, as well as the biological models, including the integration of biomolecular structure all the way to spatially organized biological communities.
- Complex analysis workflows for in situ analysis; methods for accelerated data compression and dimensionality reduction; and supervised, semi-supervised, and unsupervised methods for statistical data analytics, machine learning, and inference. All are central to virtually all large-scale data analytics in the biological and environmental sciences.
- Data sharing and archiving capabilities to support the increasing requirements of journals and funding agencies, such as DOE, to archive results for many years, often beyond the lifetime of the projects that generate the data. This long-term liability is of critical importance for making science open and responsible to the funders that pay for the work. Researchers need facilities where the very large computational datasets can be stored and shared easily with the research community.
- Methods for quantifying the uncertainty in observational and simulation data. Improving Earth system models requires determining uncertainties in the estimates of the climate and weather states spanning the instrumental record; utilizing data assimilation; and estimating the model error identified from consistent differences between the short-term, first-guess forecast and observations. Similarly, improvement of the biological models requires determining the uncertainties in experimental data, their annotation, and subsequent network inference and model predictions.

[BER 4.4]

4.4 Communication and Community Involvement

The review participants acknowledged the high level of sophistication required from the software developers who will implement and deploy the methods and algorithms on current and future HPC architectures. Continued advances in addressing the scientific challenges depend on:

- Dedicated staff with the expertise across all computing aspects for successful development and implementation of algorithms.
- Computationally focused workforce development of the user community. There is currently inadequate staff experience with HPC programming, software engineering, and big data. It is difficult to recruit staff against industry competition, and extensive training is crucial.
- Reward structures for personnel with a strong software development focus, as these staff members are critical to the success of any HPC code development group, yet they are still judged with the same metrics as scientific staff.
- Proposal mechanisms that not only focus on research-oriented work but also on the day-to-day software work required for the development, testing, and performance optimization of application codes.

[BER page C-63, Brown]

Describe the science challenges that cannot be solved in the 2020–2025 time frame using extant computing ecosystems.

We cannot get to “physics-based,” or even “physics-inspired,” models with current computing power or current measurement modalities. Concomitant advances in computing and multi-scale omics/imaging are needed to transition from predictive/descriptive to predictive/mechanistic models. This transition will mark a new era in biomanufacturing, precision agriculture, and drug discovery.

4. What top three computing ecosystem aspects will accelerate or impede your progress in the next 5–10 years? Why?










Accelerate	Why?
1. Models	Ensemble models will enable inference in currently intractable regimes.
2. Application codes	Embarrassingly parallelizable ensembles will be implemented in Spark.
3. Hardware resources	High-performance and massively parallel computing environments continuously change the scale of what is possible.
Impede	Why?
1. Visualization	Visualization for high-dimensional data is archaic—largely still based on pre-computer approaches surrounding 2D and 3D projections and scatter plots.
2. Workforce development	Few institutes train interdisciplinary computationally sophisticated scientists. Statisticians and computer scientists fail to obtain strong backgrounds in applied sciences that are needed to make substantive contributions to real-world problems.
3. Models and algorithms	Many data structures are massively multimodal: sequencing, metabolomics, and imaging data are frequently collected on the same samples. As in many areas of science, we face the general problem of “weak coregistration.”

[BER page C-159, Urban]

4. Software Applications, Libraries, and Tools: Much of the existing and emergent statistics/machine learning/applied mathematics software ecosystem is implemented in high-level languages (R, Python, Julia, Lua, Go, etc.). As in situ UQ becomes more necessary and useful, it would be extremely helpful to allow these high-level languages to couple to existing climate simulation codes in order to monitor or manipulate system variables as they are incrementally transformed through the call graph. Although these are not traditional HPC languages, and are often slower than Fortran/C++, the analyses needed are typically less computationally involved than numerical simulation, and it is relatively more important to be able to preserve the workflow and toolboxes of UQ personnel.

[BER page D-21, Maranas]

3.9 Software Applications, Libraries, and Tools: In the near future, our needs for proprietary, compute intensive software is expected to be the following: 1) Cplex 2) BARON 3) GAMS 4) Matlab 5) Gurobi. We do notice a considerable interest in the mathematical modeling community in use of the language Julia.



APPENDIX I

Sections Cited in This Report
from the Exascale Requirements
Review for Fusion Energy
Sciences (FES)

This page is intentionally left blank.

SECTIONS CITED IN THIS REPORT FROM THE EXASCALE REQUIREMENTS REVIEW FOR FUSION ENERGY SCIENCES (FES)

The following sections from the Fusion Energy Sciences (FES) Exascale Requirements Review report are cited at least once in this Crosscut report.

[FES ES.4]

Appendix C contains the FES Executive Summary.

[FES 2.2]

2.2 Technical Challenges and Opportunities

Fusion and plasma science researchers have long recognized the opportunities afforded by high-performance computing, establishing the first (unclassified) magnetic fusion energy national supercomputer center at the Lawrence Livermore National Laboratory in the mid-1970s. While plasmas are governed by well-known classical physics embodied in the Maxwell and Boltzmann equations, solution of those equations is among the greatest challenges in physics. The problems of interest exhibit intrinsic nonlinearity, extreme ranges of scales in time and space, extreme anisotropy, and sensitivity to geometric details. Real problems, particularly those that concern the plasma boundary, require additional physics to model atomic physics, neutral and radiation transport, and plasma-material interactions. The early reliance on computation grew from the recognition of these challenges, leading to the understanding that numerical methods would be a necessary complement to analytic theory.

In recent years, advances in theory, numerical algorithms and computer hardware/software have led to dramatic progress allowing meaningful, quantitative comparison between codes and experimental observations across a wide range of problem domains. Notable among many examples are: gyrokinetic modeling of turbulence and neoclassical physics for magnetic fusion energy (MFE) systems and for space and astrophysics applications; full wave calculations of the launching, propagation, and dissipation of short-wavelength radio frequency (RF) waves; nonlinear magnetohydrodynamics (MHD), including self-consistent interaction with super-thermal particles; molecular dynamics models that simulate plasma-material interactions and radiation damage in bulk materials; calculations of magnetic reconnection in plasmas of astrophysical and geophysical interest; and laser-plasma interactions and three-dimensional (3D) hydrodynamics in high-energy-density laboratory plasmas (HEDLPs). Interested readers can find more on these activities in Section 3 of this report and in the recent report on integrated simulations (FES and ASCR 2015).

Also notable, even amid these advances, is that on none of these problems are researchers close to reaching a complete, integrated first-principles solution due partly to the lack of computing power. Progress to date has largely been based on subdividing the problem into separate temporal, physics, or physical domains. It is understood that this separation is only approximate and breaks down in important instances, and thus integration into self-consistent multiphysics and multiscale approaches is critical to making further progress on many important problems. Advances in integrated solutions can be incremental and are already under way. The fusion program has identified a number of “focused integration” or “integrated science applications” that address important challenges with potentially large impact. A sample of significant examples, each of which individually require exascale-class computing to solve, include the following:

- Calculation of plasma profiles — self-consistent with turbulent and neoclassical transport of heat and particles — current drive, and momentum sources: these are crucial for prediction of plasma performance, stability, and sustainment.
- Prediction, avoidance, and mitigation of tokamak disruptions: the tokamak is the most developed magnetic fusion concept and the basis for the ITER device to be built in southern France. However, the tokamak's use as a practical energy source depends on minimizing the occurrence and impact of disruptions — i.e., sudden collapse of the plasma current, which provides the confining magnetic field.
- Simulations of the boundary plasma on both closed and open magnetic field lines, including edge transport barriers, turbulence, neoclassical physics, MHD, and atomic physics: the boundary plasma provides the critical confinement barrier and boundary condition for the plasma core and the interactions with material surfaces including the divertor, vacuum vessel wall, and RF launching structures.
- Plasma-material interactions (PMI) including relevant plasma and material physics and the effects on materials of bombardment from energetic fusion products: these calculations are essential to predict the erosion and evolution of the material surface, long-term retention of the tritium fusion fuel, and fuel recycling and impurity sources that could affect core plasma performance.
- Coupled simulations of 3D hydrodynamics and radiation transport for HEDLPs.
- Extension of magnetic reconnection calculations from small 2D systems to larger 3D systems with kinetic treatments.

Greatly enhanced computational capability is needed to address these challenges. For example, the computation of plasma profiles will require extension of the validity of the turbulence simulation to timescales typical of the transport equilibrium, that is, roughly three orders of magnitude greater than today's capabilities and significantly more again, if electron-gyroradius-scale turbulence must be resolved simultaneously (which recent work suggests that it does). Similar extrapolations are needed for the other problems as well. Of course, perfect simulation is not the goal, but rather a sufficiently faithful simulation to validate the underlying theory and to predict all needed parameters with enough precision for any future design or operational decisions. We still have considerable progress to make before that distinction is important. It is also worth emphasizing that properly executed simulations at almost every level of physics fidelity have been and will be of immense practical value in the development of our field. Computing platforms at varied levels of capability will continue to be essential in the future.

For MFE, PMI, and low-temperature plasma research, the logical endpoint of integrated simulation is the creation of “whole device models” (WDMs), which self-consistently include components that calculate all elements of an existing or proposed machine. Such a model would need to perform all of the calculations outlined above and more, modeling, for example, the system's power supplies, the electromagnets that they drive along with eddy currents in the device's structure, the response of the plasma, and the feedback control system. For a fusion-burning device, nuclear physics and neutron transport need to be included. It seems that simulating WDMs at full fidelity for all detailed physics may not be practical, even on exascale platforms, and may require beyond-exascale platforms. However, achievable and reasonable goals on exascale computers are WDM with important high-fidelity multiphysics and WDM with a mix of high-fidelity and reduced-physics components. The latter WDM is a collection of model hierarchies, selected for particular applications, to enable utilization of smaller-scale computers than the former WDM requires. The formulation and solution of the strongly and weakly coupled multiphysics problems on extreme-scale computers in the former WDM approach, and the factorization of the problem, the appropriate

choice of modules, and development of the required full-physics and reduced-model hierarchies in the latter WDM approach are all ambitious research problems.

Additional computer power offered by exascale facilities could be transformational if it can be applied to our problem set efficiently. Specific approaches to design and build exascale computers are being actively debated and tested, and the outcome of that debate may not be known for years; however, certain architectural features common to various approaches are likely to emerge and will present challenges to our applications. Challenges recognized today include these:

- The central role of power consumption in driving the new generation of machines will require that application codes perform less data movement, use more localized computation, incorporate more on-memory physics analysis, and have lower I/O requirements.
- The large increase in the number of processors with little change in processor speed will push applications toward levels of concurrency not possible for many of today's algorithms.
- More complex hardware with hierarchical and heterogeneous memory structure and with a mixture of processor types will require careful tuning of software to the hardware to realize the full capabilities of the platform.
- More frequent, fundamental changes, likely between generations of this class of computers, will require more timely evolution of software.
- Higher error and failure rates in hardware will need to be accommodated in applications and system software.

Close collaboration among the stakeholders will be crucial to overcoming these and other challenges — many of which will perhaps only become apparent as actual computing platforms are put into use since important features of the exascale hardware and software are unknown moving targets. Technologies developed by the applied math and computer science (AM/CS) communities will be essential for FES domain scientists as they develop new codes or adapt older ones. New programming models are needed that will abstract and encapsulate some of the complexity that will be presented by the exascale environment. A key will be to insulate application developers as much as practical from the inevitable changes in underlying implementations. Accompanying the needs for extreme-scale platforms and methods is a set of technologies and platforms that enrich the broader computational ecosystem. The I/O challenge will put a premium on development of in-situ and in-transit methods for data analytics and visualization. I/O challenges on current platforms already limit what can be saved for post-run analysis, leaving valuable physics on the table. Similarly, there will be a need for better methods of data management, metadata capture, and organization that span on- and off-HPC processing. Techniques for uncertainty quantification (UQ) that are practical in the exascale environment need to be developed and deployed. Finally, there is a need to develop a model for sustaining the exascale-oriented tools, libraries, and capabilities after they are developed. Widespread adoption of common tools can dramatically reduce costs and enhance productivity for individual research groups; however, this adoption will only come about if those groups are confident that the tools will be maintained and that their investment in adoption will be protected.

Although the road ahead will not be easy, the history of FES partnerships with the advanced computing facilities and collaborations with the AM/CS community, which have been crucial elements for advances made in the past, should provide a model for progress in the future.

[FES 3.1.1.4]**3.1.1.4 Computing Needs and Requirements**

Multiscale core-region turbulence simulations of experiments by a continuum gyrokinetic code that have been completed to date required approximately 15 M CPU hours each on the NERSC Edison system and approximately 37 days for completion (using approximately 17 K cores). Clearly, access to larger and more capable computers and development of algorithms that can scale to processor counts in the 10^5 – 10^7 range are likely needed to use multiscale simulation for profile prediction and to reduce the time to solution to reasonable levels. Multiscale edge simulations (even without resolving the electron gyroradius-scale turbulence) by a particle-in-cell gyrokinetic code require runs lasting a few days on the entire heterogeneous 27 PF (peak) Titan system at the Oak Ridge Leadership Computing Facility to complete a one-case physics study, consuming more than 30 M CPU-GPU hours using a scalable gyrokinetic code, XGC1. Increased physics fidelity and simulation resolution will inevitably increase requirements for data analysis and storage. Current physics output files can reach 0.5 TB per multiscale simulation for continuum codes and much more than 100 TB for particle codes. Expected increases in the simulation dimensions and inclusion of additional physics would result in at least an order-of-magnitude increase in storage needs. On-the-fly, on-HPC data analysis and reduction are necessary elements for these fusion codes. Computational requirements for particle codes will likely exceed the requirement for continuum codes. These estimates underscore the need for exascale resources for this challenge, as both approaches could conceivably require in the range of 10 B core-hours. However, it is reasonable to assume that improvements to algorithms over the next decade will likely reduce the requirements of such simulations, perhaps by an order of magnitude or more.

Edge plasma simulation is enormously challenging. It involves multiscale and multiphysics effects that are inherently kinetic and take place far from thermal equilibrium. Simulations with fundamental physics require an exascale computing ecosystem. Data from a 24-hour clock time run can reach exa-bytes assuming an exascale HPC using particle-cell methods. Dedicating this amount of storage may not be possible with future file storage systems. In-situ, on-HPC data analysis, visualization, and reduction are needed, in addition to a physics-/math-based coarse graining of the data. We will also need on-the-fly data analysis and reduction in the filesystem to make data archiving possible. Tools that support portability and accessibility among various extreme-scale HPCs will be important for productive workload sharing of edge physics research, depending upon the algorithms that dominate different edge physics topics.

[FES 3.1.2.4.1]**3.1.2.4.1 Data / I/O**

Fusion simulations can generate enormous amounts of data; for example, a simulation for MHD science can use 10^7 – 10^9 grid points, 5 to 10 variables, and 10^3 time steps, which would amount to ~100 GB in RAM and ~100 TB on disk. As computers become more powerful, increasing amounts of FLOPS can be processed; however, the amounts of data for the anticipated exascale MHD simulation challenges will not fit in memory so the storage systems need to keep pace with computing improvements — otherwise, the fast machines of the future will be wasting time waiting for the storage systems to deliver the data. The trend in HPC is use of multicore architectures, which increases the concurrent load that can be sent to the storage system. On the other hand, the availability of extra computing power can be used to optimize the I/O path and make I/O operations faster. Parallel file systems distributing the data of a single object or block of data across multiple storage nodes will probably continue to be used in the HPC environment in the future, with the degree of parallelization increasing to meet the increasing pace of computing.

[FES 3.1.2.4.3]**3.1.2.4.3 Awards Program Criteria**

One aspect that precludes progress by implicit MHD codes is the emphasis by some computational awards programs on algorithmic scaling (e.g., scaling per time step or solver iteration). The ultimate metric may instead be resources used relative to the science goals achieved. If scaling data is required as a proof of concept for resource utilization, time-to-solution scaling on the proposed problem should be chosen, not algorithmic scaling. Algorithmic scaling, which is useful internally to a project to characterize code-kernel performance, does not provide a full picture of time to solution. For example, the time-step size typically decreases when explicit algorithms scale weakly with a fixed domain size in order to avoid numerical instability; but this decrease is not reflected in a computational cost per time-step plot. These arguments are not against the use of explicit (or other) algorithms, which are well suited for certain classes of problems; rather, we argue that the current system does not permit an apples-to-apples comparison of the capabilities of different codes.

[FES 3.1.3]**3.1.3 RF Heating and Current Drive****3.1.3.1 Scientific Challenges and Opportunities**

The success of magnetically confined nuclear fusion as an economically attractive source of power largely depends on the success of ITER, the next step device currently under construction in France. ITER in turn relies on the successful operation of three plasma heating technologies (Figure 3-6). Of those, two are based on external application of radio-frequency (RF) power. These are in the electron- (GHz) and ion-cyclotron (MHz) range of frequencies and are here referred to as electron-cyclotron heating (ECH) and ion-cyclotron heating (ICH), respectively. The primary science driver in this area is the robust, reliable, and efficient operation of these systems to enable a successful ITER mission, and in the long term, steady-state operation and control of fusion-based reactors (e.g., DEMOnstration Fusion Power Plant [Stork 2009]). However, it has been observed on present devices that the operation of ICH correlates with the production of impurity ions from increased interactions between the plasma and its confining structures (i.e., plasma-material interactions or PMI), which can have deleterious effects like collapsing the plasma. As such, the scientific challenge here is to obtain a fundamental understanding of the coupling of externally applied RF power to fusion plasmas and to use that understanding to mitigate the PMI issue, thereby making available the required robust RF technology that a successful fusion mission requires. In a practical implementation, the fast timescales and extreme environment make diagnosing RF-related phenomena experimentally incredibly difficult, and as such, much of the understanding in this area relies on computer simulation. It is clear that this trend will continue, and that exascale computing resources will be one of the key tools in meeting this challenge.

The physics basis for how the application of ICH power drives plasma waves and enhances the electric potential that exists between the plasma and any confining material structure (the sheath) is thought to be understood, as is the basis for how materials respond to the bombardment of ions accelerated by that sheath potential. However, implementing these understandings in predictive and reliable computational models that have the required fidelity and dimensionality to be directly validated with experiment is only now becoming possible. In the application of RF power, experimental observations (Wilson and Bonoli 2014) have made it clear that it is the details (both geometric and in the physics model) that are important in determining whether that power will heat the plasma to fusion or whether it will burn a hole in a wall tile and collapse the plasma. As such, the present and future state of the art focuses on building reliable simulations that couple the required pieces, all at the required fidelity. In the 5- to 10-year timeframe, exascale computing resources present an opportunity to utilize a linear, kinetic plasma-wave solver that resolves not

only the tens of cubic meters of the core fusion plasma at the fastest (nano- through microsecond) timescale but also the millimeter-scale features and nonlinear field response near the launching antenna structures and confining material walls (Figures 3-7 and 3-8), together with models of the material response to plasma bombardments. Indeed the transport of sputtered impurities through the edge plasma that ultimately affects core fusion performance is a multi-time-scale and multiphysics process that couples the RF, PMI, SOL/divertor, pedestal, and core problems. At the material surfaces, many atoms are quickly ionized and are on gyro orbits that intersect the wall. These are promptly redeposited, but the material created as a result is often loosely bound and has different mechanical and thermal properties. For the ions that escape that fate, subsequent transport is highly dependent on local conditions and the 3D geometry of the machine. The fraction of sputtered atoms that end up in the core plasma is a strong function of source location through processes that are poorly understood. Once in the core, the impurities are subject to turbulent and collisional transport processes that we are only just beginning to understand. Understanding this multiphysics process will allow us to design strategies to mitigate the interaction with material surfaces, while maximizing the heating efficiency and reliability.

In parallel to the above challenge is the impact that the application of RF power has on, and how itself is affected by, plasma turbulence, MHD instabilities, and energetic particle populations, that is, the timescales that exist between the RF and transport scales. For the control of MHD instabilities, RF actuators have long been recognized as tools that will be essential for realizing a steady-state tokamak. How RF power interacts with turbulence and energetic particles is less clear. However, what is clear is that the proper design of reactor-grade, steady-state tokamaks involves coping with a complex interplay of the effects of transport, external current drive and heating profiles, MHD stability, and control of edge density and temperature pedestals and scrape-off plasma parameters. While great strides have been made in developing the modeling capability for most critical areas, very little progress has been made in modeling the whole device: thus, the exascale computing resources represent an opportunity to integrate the advances that have been made in transport, core and edge MHD, RF current drive, and SOL simulations in order to determine optimal reactor configurations and operating scenarios.

A stated top-level goal for DOE's Office of Fusion Energy Sciences (FESAC 2014) is the use of massively parallel computing for validated predictive simulation for magnetically confined fusion plasmas. This capability should ultimately enable, and minimize the risk in, future fusion energy development stages. A subset of this goal is the integration of independently developed computational tools that make up the DOE portfolio of legacy and state-of-the-art simulation codes. For many years, RF source modules have been employed as components within integrated simulation (e.g., within TRANSP [Hawryluk 1980]), and the RF SciDAC program has produced both high-fidelity and reduced models for many aspects of simulating the application of RF power, with the goal of creating a predictive and robust tool that will bring the coupled antenna-to-core system to within reach. This opportunity and the challenge of integration are dealt with more thoroughly in the whole-device-modeling section of this report (Section 3.1.4).

In principle, the RF challenge (and subsequently the entirety of the plasma physics occurring at longer timescales in fusion plasmas) may be met by solving the full 6D + time Vlasov-Maxwell coupled system (with an appropriate collision operator). However, with a naive, explicit time advance approach, this would require time steps of the order 10⁻¹¹ seconds in a system where the resulting impurity transport occurs at milliseconds on a mesh that has to resolve the Debye length throughout the device, meaning many trillions of spatial points. Without suitable application of some (yet-to-be-created) multiscale time integration scheme or the like, such an approach likely exceeds even the exascale, and this still does not incorporate the auxiliary equation set for the material response to the incident plasma bombardment, or generated neutron flux, or interaction of these with any of the myriad of engineering components required in the operation of any fusion

reactor. As such, we therefore expect a continued reliance on scale separation combined with integration within the 5- (and perhaps even 10-) year timeframe, although we continue to keep our eye on possible avenues of computer science and applied math that may enable this 6D approach to be incorporated into a comprehensive RF predictive capability.

[FES 3.1.3.4.5]

3.1.3.4.5 Computing and Programming Environment

- Continued or expanded access to fast job turnaround at moderate concurrency to support rigorous validation and uncertainty quantification of models (probably as a requirement of any contributions to a community whole device model). The use case of running a particular code many times for varied input parameters is likely to become more common.
- Training, specifically student training to build the next generation of leadership computing-capable science experts.
- Common programming model/environment across facilities. While we recognize the difficulty in providing this environment, we imagine that at a minimum, a set of best practices for programming in scientific codes could be made available that result in (1) minimal time invested when porting from one machine to another, (2) minimal time required for a computer scientist to learn a new code, and (3) minimal wasted efforts due to codes falling to legacy status. At more advanced levels, this common environment or approach may be in the form of offering advanced software engineering tools that perform code re-engineering or provide suggestions for doing so.

[FES 3.1.4]

3.1.4 Whole Device Modeling

3.1.4.1 Scientific Challenges and Opportunities

Several factors — including the high cost of building future experiments and prototype fusion facilities, such as the ITER device now under construction, the complexity of the multiscale multiphysics in fusion plasmas, and the advances expected in extreme-scale computing over the coming decade — provide strong motivation for developing integrated simulation capabilities to predict key physical processes occurring in tandem in these devices. Such capabilities can ultimately enable predictive simulation of an entire fusion device, thus minimizing risk to the device and guaranteeing its successful operation. WDMs are required to assess reactor performance in order to minimize risk and qualify operating scenarios for next-step burning plasma experiments, as well as time-dependent or single-time-slice interpretive analysis of experimental discharges.

Figure 3-9, adapted from the Integrated Simulation Workshop report (FES and ASCR 2015), provides a high-level view of the WDM. Highlighted in the figure are the levels of complexity — or physics hierarchy — constituting a WDM that span reduced models to extreme-scale kinetic physics models. The former can be analytic or physics based, should typically be computationally fast, and can run on a reduced number of processors for fast in-situ jobs in direct connection with experimental runs. The latter extreme-scale models provide a deeper physics understanding, need to be first-principles-based codes, and typically require capability computing. Into this category fall gyrokinetic codes coupled with multispecies fluid/MHD, RF, and materials codes for fast timescale physics. In between the two are the so-called advanced reduced models. The fitting parameters in these models are typically derived from ensemble results of extreme-scale calculations over a finite range of plasma parameters and might be applicable to only a limited operational space. Advanced reduced models are not as fast as reduced models, but they can still be used in time-dependent simulations for higher-fidelity reduced calculations. A WDM plan should provide the flexibility to choose among levels of physics hierarchy depending on the needs, which can span from in-situ experimental planning and analysis (simulations requiring only a few minutes) to fundamental physics understanding and prediction (simulations requiring a few hours to a few days).

The choice of multiple components will facilitate the verification and validation of individual physics models and the verification of integrated physics. The VVUQ of a reduced-model WDM will rely mostly on the **availability of a large-scale computer ecosystem** owing to the need to carry out large ensemble calculations. VVUQ of a high-fidelity, extreme-scale WDM still remains an open research topic in the ASCR community. Multiphysics processes are coupled together in a nonlinear manner. Thus, studies on the individual components may not present much reflection on the integration result.

Special attention should be directed to the load balancing among each component, alternative coupling schemes, component interchangeability, regression analysis, and exceptions handling. Collaboration with ASCR is critical. Scalability to the future exascale computers must also be a significant part of WDM research activities.

A WDM should provide the interface between simulations and experiments. The WDM should be connected to an experimental database for two purposes. At one end of the process, the experimental data, once they are calibrated and processed, are an input to simulations; at the other end, they are used for validation of simulation outputs. Data should be prepared using a common format so that interfacing with codes is facilitated, as well as exchange of data among experimentalists and modelers. Along this line, some topical groups of the ITPA (International Tokamak Physics Activity framework for ITER) are considering adopting the ITER Data Structure (IDS) as a standard, as well as the Integrated Modeling and Analysis Suite (IMAS) as a framework to support all of the physics modeling needs. Access to a centralized experimental database would facilitate cross-checks among experiments, as well as submission of runs for Uncertainty Quantification studies. All of this interaction would require storage capabilities, large-scale computing and centralized data management, processing, and visualization.

The output from simulations should be reduced using centralized software capabilities to ensure that the metrics needed for V&V are defined consistently.

Assessing uncertainties in the physics models needs to be incorporated in the WDM codes so that the simulation results include the confidence intervals of the predictions. In predictive simulations, UQ tools can be used to evaluate the probability of events occurring (such as disruptions and edge-localized modes), for the computation of confidence intervals of predicted quantities, and for the optimization of plasma performance. There is very little work being performed in this direction at this moment, and action should be taken to ensure that interfaces for this scientific objective are ready in ten years.

The following have been identified in the white papers as scientific challenges that represent attractive scientific opportunities for the collaborative FES/ASCR community to tackle in the next five or ten years, and whose progress would benefit from the availability of increased computing capabilities.

- Predictive WDM core-edge simulations that are based on the kinetic models: First-principles-based kinetic codes that can treat the core and the edge self consistently are extremely valuable for validation against experiments and for verification of reduced models. The validation activity needs to be increased during this time period. Despite the significant progress in experimental diagnostics in the plasma edge region, the error bars for experimental data in this region remain among the largest in tokamaks.
- Efforts to bridge the gap between short and long timescales: A multiscale time-integration method needs to be developed in order to “telescope” the kinetic simulation to experimental timescale. Plasma turbulence correlation time, which determines the transport coefficients, is milliseconds, whereas the resulting plasma profile evolution can be over seconds in a large size tokamak.

[FES 3.3.2.4]

3.3.2.4 Computing Needs and Requirements

Large-scale particle-in-cell simulations generate extreme amounts of numerical output; a three-dimensional simulation of a $(100\text{-}\mu\text{m})^3$ plasma volume with 400 cells per micrometer and 10^4 particles per cell, where each particle is described by 10 double-precision numbers, generates $\sim(100 \times 400)^3 \times 10,000 \times 100 = 6 \times 10^{19}$ bytes of particle data per time step output. To work with these amounts of data, it is essential to have robust, scalable, and easy-to-use I/O tools, like HDF5 or ADIOS; optimizing I/O on extreme computational scale is “still too much of a black art.” It is expected that for detailed high-frequency data analysis, the required I/O performance levels will greatly surpass the expected availability, necessitating efficient data reduction and in-situ analysis/ visualization capabilities. In this regard, common, versatile, high-quality, and easy-to-use tools for post-processing and in-situ visualization and analysis are needed. Efforts toward standardizing and optimizing parallel I/O for PIC codes are currently under way (e.g., OpenPMD at <https://github.com/openPMD>), which will facilitate data exchanges between codes, data processing, and visualization software. For some codes, the support of Python as front-end to number-crunching modules is essential. The hardware and associated software of exascale supercomputers will involve highly specialized programming skills, requiring proper continuous training as they evolve over time, as well as close partnership between applications developers, users, computer scientists, and applied mathematicians. For maximum benefits to the community, the suites of tools will benefit from more intuitive interfaces, wide dissemination, extensive documentation, and resources for user support.

[FES 3.4] [FES 3.4.1.1]

3.4 Verification and Validation**3.4.1 Needs Driven by Experiments****3.4.1.1 Scientific Challenges and Opportunities**

There are four main HPC drivers arising directly from fusion experiments:

- 1. Experimental Operations Planning.** Experiments are usually motivated by or connected to tests of particular theoretical or modeling questions, but only a fraction of experimental proposals are currently qualified ahead of time by extensive HPC modeling. However, this sort of modeling can be particularly useful for operation in unfamiliar or challenging parts of parameter space. With improvements anticipated in capabilities and fidelity of WDM as an element in experimental planning, activity in this area is likely to increase and should lead to more successful operation and more efficient use of expensive run time. Looking further into the future, it is anticipated that proposals for run time on ITER or devices of that class would only be approved if accompanied by extensive discharge modeling (Greenwald et al. 2015).
- 2. Run-Time Experiment Steering.** While remote HPC will likely not be used in real-time control loops in the present-day short-pulse tokamak experiments, its use for between-shot operations steering has already been demonstrated (e.g., White et al. 2013, where between-shot linear gyrokinetic modeling was employed using a local 500-core cluster at the Massachusetts Institute of Technology). This approach entails accelerated workflows that can provide preliminary analysis of experiments in minutes rather than days or weeks. The result is that experimental operations can be adjusted throughout a run day to more precisely meet the experimental goals and to increase overall productivity of the facility. Early experiments in queuing for near real-time applications are currently being tested at NERSC by the General Atomics group during its current experimental campaign. The trade-off for time spent in data transfer can be favorable for computationally intense calculations. In a steady or near-steady operation of fusion reactors, the technique can be extended to the real-time control of the experiments. Technical challenges remain, as discussed below.

- 3. Code Validation.** A large fraction of run time is devoted to tests of particular theoretical ideas, often backed by simulation results, or for explicit comparisons with code predictions. The increasing rigor of validation activities will likely drive much larger computational requirements — especially for sensitivity analysis.
- 4. Interpretation of Experimental Data.** Comparisons between simulations and experimental data can generate new hypotheses about the physical phenomena being observed. This activity is less formal than VVUQ but can be an essential element of scientific discovery.

[FES 3.4.1.2]

3.4.1.2 Priority Research Directions

Each of the drivers listed above has implications for research directions:

- 1. Experimental Operations Planning.** Fast turn-around and large numbers of runs characterize this application. Research will be required to create the reduced or surrogate models that are well tested against full-physics calculations.
- 2. Run-Time Experiment Steering.** This capability requires exceptionally fast turn-around for submission, queueing, execution, and post-processing of HPC tasks. This application requires more complicated data workflows — for example, reading data stored externally to the HPC system — that are currently not utilized or possible due to firewall related issues. Research and development are likely required to arrive at widely accepted solutions for federated authentication and authorization, credential forwarding, etc.
- 3. Code Validation and Uncertainty Quantification.** These requirements are covered in Section 3.4.2 of this report.
- 4. Interpretation of Experimental Data.** This application would include both capacity and capability-type computing. Research needs would generally be aligned with those identified for the science topical areas (Sections 3.1–3.3), whole device modeling (Section 3.1.4) and the V&V discussion in Section 3.4.2.

[FES 3.4.2.2]

3.4.2.2 Priority Research Directions

We identify four priority research directions in the VVUQ area:

- 1. Develop/identify improved methodologies for code verification, especially for coupled/integrated physics models.** Software quality assurance and engineering techniques are often — but not always — applied. Some activity to identify and implement “best practices” could be useful. As noted above, despite the best efforts and intents of developers, verification typically depends on convergence studies and code-to-code benchmarking. The method of manufactured solutions (MMS) may be practical in our domain, and could be a target of opportunity for further research. This method functions by choosing, a priori, a manufactured solution to a modified set of equations that consists of the original equations plus an analytic source term. The task for developers is to make an intelligent choice for the manufactured solution and then to solve the inverse problem — find the source function that is required to make the manufactured solution correct. This method has been successfully applied to plasma fluid turbulence and computational fluid dynamics, including problems involving multi-physics couplings, suggesting that there may be a broad set of problems in the plasma physics domain that are amenable to this approach.
- 2. Assess and extend existing methods (intrusive and nonintrusive) for UQ/ Sensitivity Analysis (UQ/SA) on our codes.** As noted above, the major source of uncertainty in many plasma simulations is the forward propagation of errors in input parameters. Ideally, a probability distribution representing the uncertainties in all input variables would be known,

and an efficient method would be available to transform these distributions through the code to distributions in the output parameters. A simple and common approach is to try to estimate this level of uncertainty by executing an ensemble of code runs, each with slightly modified inputs. However, for problems requiring extreme-scale computational resources, this sampling-based method can be expensive and limits the range and numbers of parameters that can be tested. For a problem with a large number (high dimension) of input parameters, brute force approaches may not be possible with available computational resources. There is a good deal of ongoing research on sensitivity analysis methodologies aimed at addressing this problem with the goal of improving their efficiency and statistical rigor. The methods can be classified as intrusive — that is, they require some lesser or greater degree of modification of the underlying code — or nonintrusive. Generally, there is a trade-off between intrusiveness, which puts a burden on code developers, and computational intensity. In some cases, the intrusive methods may not be applicable or practical, so the optimal approaches need to be determined on a case-by-case basis.

The efficiency of random sampling-based methods can be improved through several approaches — deterministic sampling (importance sampling, pseudo-random) or non-sampling methods (e.g., pseudo-random, polynomial chaos, probabilistic collocation, etc.). While providing improvements, all typically suffer from the “curse of dimensionality”: that is, they work efficiently for only a small number of input parameters. It is an open research question as to which approach would work best for each of our calculations.

Rather than sampling via large ensemble of code runs, SA methods have been proposed that can perform the calculation (i.e., derivatives of outputs with respect to inputs) as part of the main simulation. These are more intrusive but may offer the prospect of greater efficiency — especially for higher-dimensional spaces. Additional computational overhead is typically a small multiple of a base case, so the win becomes bigger when the dimension of the input space exceeds a few. These methods are currently the subject of intense research, but in our problem space, there are few applications so far. It is an open question as to which of these methods are applicable to our problems, what their limitations are, and what their costs for implementation will be. The most intrusive methods require some significant rewriting of the underlying code.

For example, by solving an adjoint equation along with the originals; a single adjoint can provide the derivative of a single output with respect to all inputs. However, solvers for the original system of equation may have problems with the adjoint equation. This method has been developed and tested for CFD problems in recent years, and a research team is currently trying it out on a fluid plasma turbulence problem. Its general application needs to be demonstrated. Somewhat less intrusive is automatic differentiation. This method calculates local (linear) dependence of inputs to outputs by a symbolic differentiation through applying the chain rule to each elementary operation in the code. The approach is somewhat analogous to code compilation. The method might be extended to higher-order, nonlinear dependences — but at a cost. The applicability to our codes needs to be demonstrated. One limitation is that the method requires source code for the entire calculation — no opaque libraries or subroutines are allowed. Implementation of these approaches may be particularly challenging in the exascale era as codes may increasingly rely on independently optimized external libraries.

Another possible method is to explore the sensitivities of a model through the use of surrogate models or computations with reduced physics or resolution. These can be used to identify the most important inputs, which can then be studied with the full models or in some cases may themselves produce reasonable estimates of sensitivity — even if their absolute accuracy is unacceptable. (That is, for a quantity of interest Q and input variable x , they may produce an acceptable estimate of dQ/dx even if they cannot

produce a useful value of Q .) All of the SA methods described above might benefit from “pre-screening” in this manner. An interesting example was shown at the Exascale Requirements Review meeting where the sensitivity of lower hybrid current drive (LHCD) models to assumed density and temperature profiles was studied. A very systematic exploration of this sensitivity was carried out using a combined ray tracing/Fokker-Planck model (see Figure 3-16). This approach may stand on its own, or help direct studies with the full wave/Fokker Planck model targeting more interesting or important parts of the parameter space. In any particular case, this approach may or may not be applicable and so must be studied for each problem. Studies of the sensitivities of the multiscale turbulent transport code XGC1 are being attempted using a similar method. This approach would also be extremely relevant in the context of integrated and whole-device modeling where a variety of physics models at various fidelity levels will be coupled together. In such a context, understanding whether individual reduced or “advanced reduced” models exhibit the same sensitivities as the “first-principles” models they seek to describe will be important in its own right, as well as understanding how the sensitivities of multiple coupled models interact with each other.

- 3. Develop and adopt tools to enhance data management, including automated capture of provenance and other metadata and improved data browsing of very large and complex data sets.** The data management challenge for VVUQ extends to processes occurring both on and off HPC hardware and will typically involve an extended collaboration that includes modelers and experimentalists. The scope of VVUQ activities includes raw experimental data and its reduction and analysis, preparation of inputs for code runs, post-processing of code outputs, and comparisons between experimental and simulation results including statistical analysis. Meaningful validation requires careful documentation of assumptions and parameters at each step. Overall, a much more systematic approach to metadata and provenance capture is required.

The large datasets generated by simulations and experiments offer additional challenges. We need tools capable of more effective data browsing and exploration of these data — probably driven by more systematic provision of metadata and automated feature or event detection. The ability to explore 5D and 6D data effectively may depend on tags generated during analysis — akin to the process of geo-tagging. The overall challenges in data management are well covered in detail in the report from the recent Workshop on Integrated Simulation (FES and ASCR 2015). A particular challenge in the context of PRD2 is developing workflows for tracking ensembles of simulation results (particularly first-principles ones) used for UQ and sensitivity analysis.

- 4. Develop and deploy new tools that enhance analytics and visualization for computations under emerging computing architectures.** While the new architectures present a challenge to compute, it is already the case that analytics and visualization are lagging. I/O challenges restrict the amount of data that can be stored, preventing the full exploitation of the simulations performed. In the future, the new computing architectures may make this situation worse if the present technologies are used. Post visualization of the huge quantities of data produced, including 5D and 6D arrays, is a daunting challenge. One solution is to perform more of the data analysis and visualization tasks during the computation itself while the data are still in memory (in situ) or being transferred between processors (in-transit). In-situ and in-transit methods are particularly well suited to processor-rich, I/O-poor architectures. These methods are the subject of current research but are not yet widely deployed in production. Finding the best approaches for our domain and deploying them will take close collaboration between applied computer scientists and computational physicists. Ideally, toolkits would be made available that would ease the adoption of these approaches. Because full-resolution data are only available

during the HPC run, decisions about analysis must be made ahead of time — thus, it will also be a challenge to develop strategies that maximize the productivity of each run. To that end, in-situ analysis could also be used to detect interesting events or features in the simulations, triggering particular analyses or higher-resolution I/O around the features for off-line analysis and visualization. For code validation, a particularly important subject of in-situ analysis may be the generation of synthetic diagnostic data. The report from the recent Workshop on Integrated Simulation covers the overall challenges for in-situ analytics well (FES and ASCR 2015).

[FES 3.4.2.4]

3.4.2.4 Computing Needs and Requirements

VVUQ is a strong driver for computing needs. The sensitivity analysis described above is an essential activity — simulation results without an assessment of sensitivity to input parameters have severely limited utility. Quantities of interest typically depend on many parameters, and spanning that space can be computationally expensive — even using some of the techniques described above. This finding suggests that UQ drives 10–100 simulations for every “base case” for which a researcher needs to apply UQ. Extrapolating from the most demanding of fusion simulations, the computation of turbulent-driven transport, implies a requirement of more than 200 million core-hours for each base case of interest using today’s codes. Note that each of these is only for a snapshot in time at a particular location in the plasma. The aggregate impact is hard to estimate but will certainly be significant. One might expect that the amount of computational time dedicated to VVUQ will be at least as large as other activities combined.

It is also important to note that simulation cases run for validation will often be among the most demanding. That is, they typically require using the most realistic physics assumptions. As an example, a turbulent transport calculation might require including both electron and ion gyrokinetic dynamics with the real mass ratio, finite β (plasma kinetic energy/magnetic field energy) effects, and operation near critical gradients (i.e., at marginal stability). Even if agreement is achieved with simpler assumptions, it is critical to understand whether this agreement is fortuitous or whether it is maintained in the full physics calculation.

[FES 4]

4 PATH FORWARD

For researchers to move forward in addressing the scientific challenges documented in this review, an evolving computing ecosystem must support them. This computing ecosystem includes computational and data capabilities and capacity; scientific applications and software; and the infrastructure for data transfer, sharing, access, and analysis — each of which must undergo further investment and development on the path to exascale. New advances required for large-scale verification and validation must be realized. The coupling of an exascale ecosystem, along with a convergence of theoretical, mathematical, computational, and experimental capabilities, will bring many opportunities for new scientific breakthroughs at an unprecedented scale.

Collaboration between FES and ASCR scientists and facilities staff will help ensure development and deployment of an effective, realistic computing ecosystem that enables revolutionary discoveries in areas described in this report. The recent FES and ASCR (2015) Report of the Workshop on Integrated Simulations for Magnetic Fusion Energy Sciences highlights collaboration on similar topics. The computing ecosystem requirements resulting from this review will form the basis to direct future investments of time and resources. These requirements fall into broad categories: methods development; computational environment; data and workflow; and communication and community involvement.

[FES 4.1]**4.1 Methods Development**

The advancing complexity of computer hardware requires FES researchers to have more scalable, performant algorithms and applications that are capable of efficient execution on future computing architectures fielded by ASCR facilities. Meeting participants discussed those computing ecosystem aspects that will accelerate or impede their progress in the next 5–10 years. Participants named application codes and verification and validation techniques, as well as models and algorithms, as key factors requiring significant methods development activity. A representative list of the methods development topics discussed by the review participants is as follows (see Section 3 for a more detailed overview of the methods development topics presented by the review participants):

- Threaded and accelerated versions of the direct and sparse matrix algebra in combination with semi-implicit time advance methods (3.1.1).
- Ameliorating the memory and logic intensity of particle/mesh mapping (3.1.1).
- Efficient algorithms for Maxwell’s equations (3.1.1).
- Improved capabilities for verification, validation, and uncertainty quantification (3.1.2, 3.1.3, 3.1.4).
 - New mechanisms for managing, discovery, method exploration (brute force, adjoint-like, automatic differentiation, etc.).
- Improved multiple-timescale, multiple-physics coupling (3.1.2).
- More scalable and performance linear algebra solvers (3.1.2).
- Domain-specific solvers (3.1.2).
- Improved scaling and resolution in RF algorithms (3.1.3).
- Hybrid algorithms capable of kinetic and magneto-hydrodynamic features (3.3.2).
 - J Improved particle solvers (3.3.2).
- High-dimensional PDE solvers (3.3.2).
- Exascale-ready adaptive mesh algorithms (3.3.2).

In almost all discussions for new models and methods, integrating applied mathematicians and computer scientists with computational scientists is crucial for success. Programs like SciDAC are key examples of success.

A close dialogue between FES and ASCR researchers and facilities staff will streamline and promote research and development through the exchange of information about computing ecosystem roadmaps and application requirements and the availability of systems for simulation and testing.

[FES 4.2]**4.2 Computational Environment**

Requirements for the access, scheduling, and software ecosystem identify an evolving use-model. The “traditional” HPC model, defined as running a large simulation that generates data that are then post processed, is no longer the only primary use-model for many FES projects. Emerging demands, such as for complex workflows and near-real-time computing, are changing the landscape.

New requirements for the computing ecosystem include the following:

- Support a user-friendly development environment, with uniform environments among DOE HPC centers supporting portable, high performance across systems with improved and new runtime systems that mask HPC complexity from application programmers, and training aimed at all levels of HPC developers, including nontraditional HPC users (3.1.1, 3.1.3, etc.).
- Reconsider batch queuing methods and priorities. Queue wait times can limit the ability to push through mid-scale/capacity computing science needs (3.1.1, 3.2.1).
- Promote accessibility (3.1.1).
- Keep balance between storage, memory, and FLOPS (put the problem in memory) (3.1.2).
- Help support better software engineering efforts (programming environments, models, software stacks, etc.) (3.1.3).
- Identify/develop programming models and languages that can increase productivity and still provide performance (3.1.3). There will be a need to abstract and encapsulate some of the complexity that will be present in exascale environments. A key to success in the past has been — and will again be — to insulate application developers as much as practical from the inevitable changes in underlying implementations.
- Reduce the need for porting and maintaining code, perhaps through containers or software-as-service models (3.1.3).
- Sustain the capabilities — tools, libraries, etc. — that are developed in the process of moving to exascale computing; explore models for doing this by other agencies (e.g., National Nuclear Security Administration).

[FES 4.3]

4.3 Data

The scale of data generated from FES simulations and the requirements needed for verification and validation have created an opportunity and a challenge. ASCR and FES facilities must create more data-centric environments with highly effective data analytics tools for their users. Development of such environments and tools will require expertise from domain scientists, data scientists, and applied mathematicians. Continued collaboration will be required to assess proper deployment of the environments as computing resources evolve.

Requirements related to data generation, storage, transport, curation, and exploration include the following:

- Because simulations for particle codes already generate more than ~100 TB of physics data, expect at least an order-of-magnitude growth (3.1.1).
- Develop advanced workflow managers (3.1.3).
- Support access to simulation and experimental databases at runtime for input, analysis, and verification (3.1.4).
- Support in-situ analysis (3.1.4).
- Improve capture and organization of metadata, including provenance.
- Develop community standards and federated databases.
- Develop data curation methods.

[FES 4.4]

4.4 Communication and Community Involvement

To foster development of the requisite exascale-level skills and to disseminate this learning widely throughout the community, DOE (with the ASCR facilities) must seek to create or make use of existing initiatives that promote the following:

- Proposal/award processes to support the wider array of requirements, including flexible allocations mechanisms and metrics based on science goals.
- Expanded involvement of ASCR applied math and computer science (3.1.3).
- Deeper involvement in open source applications and libraries.
- Workforce development (education and training).

These activities are ongoing today in multiple institutions; however, efforts to connect them to the larger science community have been attempted on an “ad hoc” basis to date. ASCR facilities can explore new or improved communication channels and activities. In addition, experience has shown some of the best impact from strong collaborations. The previously identified structured collaborative efforts could focus more attention on this important mechanism for community involvement.

[FES page C-9, Ernst]

4. What top three computing ecosystem aspects will accelerate or impede your progress in the next 5–10 years? Why?

Accelerate	Why?
1. Hardware resources	During this time frame, the Xeon Phi is coming into its own in systems of unprecedented scale (9,300-node Cori II at NERSC and 3,500-node Aurora at Argonne). The Knights Landing and Knights Hill Xeon Phi multicore architecture not only has enough on-chip memory per core to serve as the main memory in many applications, at 4.4 times the bandwidth of off-chip memory, but off-chip memory can also be accessed at the DDR4 bandwidth of 90 GB/s, a dramatic improvement over the previous generation Knights Corner co-processor PCIe connection. The next system at OLCF will be IBM Power 9 + NVIDIA GPU based with NVLINK, which also allows the GPU core to access system memory at a similar 80 GB/s bandwidth. The ability to access system memory at CPU speeds, though not alone sufficient to keep all cores busy, should simplify programming and greatly improve performance.
Impede	Why?
1. Process for allocation of computational resources	For new proposals, sufficient computational resources to meet the proposed milestones are not guaranteed. NERSC ERCAP awards appear to be based on prior use, and significant time increments for new projects are handled through the ALCC, INCITE, and NICE programs. Applicants must compete across disciplines and recompute each year. These proposals are lengthy and peer-reviewed. If legally possible, providing additional computational resources as part of new grant awards would leverage one review process and reduce overhead for researchers, DOE, and reviewers, while helping to ensure sufficient resources for funded research.
2. Application optimization/development support	Support resources for application development and optimization on the new architectures are limited. NERSC often provides good webinars and holds hackathons. However, having an expert working closely on code improvement appears to come mainly with Scientific Application Partnerships (FASTMATH), sometimes possible through SciDAC grants or ad hoc collaborations. Exascale computers will have one-fifth the memory per core.

[FES page C-21, Umansky]

What top three computing ecosystem aspects will accelerate or impede your progress in the next 5–10 years? Why?

Accelerate	Why?
1. Application codes	In the edge plasma field, there are a number of existing application codes that could likely produce important advances if more work could be done on applications.
2. Data workflow	Developing better ways of communicating data should improve the throughput.
3 Hardware resources	Increasing the resolution of edge plasma simulations will open the possibility of including a larger range of spatial and temporal scales.
Impede	Why?
1. Workforce development	Insufficient application workforce would impede progress.
2. Internal/external libraries/frameworks	Application codes used in edge plasma community use standard mathematical libraries and frameworks (SUNDIALS, PETSc, SLEPc, Hypre, Chombo, etc.). If these libraries cannot successfully transition to new architectures coming in next 10 years, then the application codes would not be able to take advantage of advances in computer hardware.

[FES pages C-24, 25; Jenkins et al.]

Computational Needs: Coupling Extended MHD Simulations and RF Wave Codes

T.G. Jenkins,* S.E. Kruger,* J.R. King,* E.D. Held[†]

*Tech-X Corporation, 5621 Arapahoe Avenue Ste. A, Boulder, CO 80303

[†]Utah State University, Logan, UT 84322

Current Science Drivers

The experimental use of radiofrequency (RF) waves to drive localized plasma currents that suppress deleterious tokamak instabilities (e.g., neoclassical tearing modes [NTMs]) has been remarkably successful [1]. Sophisticated control algorithms for the application, timing, and steering of injected RF power have been empirically developed; these have significantly reduced the likelihood of NTM-induced disruptions that could damage experimental hardware in existing experiments [2]. Avoidance of such disruptions will be crucial in the ITER device, because its stored energy is projected to exceed that of any present-day device by at least an order of magnitude. Numerical simulation can augment physics understanding of the coupled RF/extended-MHD interaction that is critical to mode stabilization and can also explore broader issues such as control strategies, RF power optimization, and so on. Parameter regimes for which the plasma disruptivity is not empirically known can be explored in numerical experiments without risk to the device.

The recent derivation of a self-consistent theoretical [3–5] framework in which the RF/extended-MHD interaction can be explored facilitates such predictive numerical analysis.⁵ Loose coupling between the RF and MHD codes is sufficient for such analysis, permitting the use of vastly different data structures and representations (e.g., finite element, spectral, and ray-tracing characteristics) in these two classes of code. Robust numerical techniques [6] have been developed to map RF ray tracing data onto more conventional extended-MHD representations such as finite elements. A python-based simulation framework [7], developed to manage the interactions of loosely coupled physics components, facilitates the RF and MHD data manipulation and exchange.

⁵ Initial theoretical and computational work—carried out by the SciDAC Center for Simulation of RF Wave Interactions with MHD (SWIM)—has continued on a limited basis with support from the SciDAC Center for Extended MHD Modeling (CEMM). Reduced MHD simulations that use heuristic models for the RF-induced currents (methods used primarily by European research groups) have also begun to explore some of the basic physics imparted by RF.

Science Challenges for the 2020–2025 Time frame

Research and progress in coupled RF/MHD simulation requires large-scale computing resources. Physics issues of interest include (a) quantification of the driven RF current efficiency in various operating regimes; (b) the influence of source width and position on island stabilization; (c) the detailed physics of RF effects on closure, and the role of this physics in Fisch-Boozer or Ohkawa stabilization mechanisms; (d) the relationship of detailed 3D models to 1D modified Rutherford equations; (e) development of optimal NTM control strategies for a fixed RF input power; and (f) optimal responses to mode locking scenarios that arise as NTMs interact with resonant magnetic perturbations (RMPs). The small size of the resonant region in which the RF modifies plasma dynamics necessitates very high resolution (possibly sub-millimeter), while the tearing mode size is on the order of the size of the device. Thus, substantial computing efforts (hundreds of runs using tens of thousands of cores) at or near the capacity of existing resources are required.










Exacerbating the problem is the need for higher-fidelity physics models. Closure computations, which in their most general form require solutions of 5D drift-kinetic equations throughout the computational domain [8], impose additional computing and storage requirements; the phase space resolutions required to guarantee numerical convergence may not be attainable using present resources. Tighter coupling requirements between the RF and MHD aspects of the problem may also be imposed by neoclassical or closure physics. Existing computing ecosystems have already modeled rudimentary RF/MHD interaction at marginal resolution. However, the increased capability that larger-scale systems afford will enable more detailed models, including neoclassical effects and full closure computations, to be fruitfully compared with experiments.

Top Computing Ecosystem Aspects to Accelerate or Impede Progress

Accelerate	Why?
1. Dedicated consulting and financial support for code refactoring issues raised by new computing ecosystems.	Minimizes scientific productivity losses as the computing platforms supporting the scientific studies evolve, and ensures the optimal use of new computing platforms.
2. RF ray-tracing code development to make optimal use of GPU architectures.	GPU-enabled RF computations are faster, enabling more tightly coupled RF/MHD modeling scenarios to be carried out efficiently.
Impede	Why?
1. Code refactoring requirements imposed by fundamental computing ecosystem changes.	Diverts time and effort from the physics studies to re-establish currently extant code capabilities.
2. Systems that prohibit the use of Python on the back-end nodes (although we only need the master node to run Python).	The problem is loosely coupled and the feedback systems require rapid prototyping. Python is perfectly suited to this problem.

References

- [1] H. Zohm et al., Nucl. Fusion 39, 577 (1999); G. Gantenbein et al., Phys. Rev. Lett. 85, 1242 (2000).
- [2] E. Kolemen et al., Nucl. Fusion 54, 073020 (2014).
- [3] C. C. Hegna and J. D. Callen, Phys. Plasmas 16, 112501 (2009).
- [4] J. J. Ramos, Phys. Plasmas 17, 082502 (2010); Phys. Plasmas 18, 102506 (2011).
- [5] T. G. Jenkins and S. E. Kruger, Phys. Plasmas 19, 122508 (2012).
- [6] T. G. Jenkins and E. D. Held, J. Comp. Phys. 297, 427 (2015).
- [7] <http://sourceforge.net/projects/ipsframework/>.
- [8] E. D. Held et al., Phys. Plasmas 22, 032511 (2015).



APPENDIX J

Sections Cited in This Report
from the Exascale Requirements
Review for High Energy Physics
(HEP)

This page is intentionally left blank.

SECTIONS CITED IN THIS REPORT FROM THE EXASCALE REQUIREMENTS REVIEW FOR HIGH ENERGY PHYSICS (HEP)

The following sections from the High Energy Physics (HEP) Exascale Requirements Review report are cited at least once in this Crosscut report.

[HEP ES]

Appendix D contains the HEP Executive Summary.

[HEP 1.2]

1.2 Meeting Structure, Report Preparation, and Report Organization

The review began with a series of talks discussing the HEP science drivers, focusing on the scientific goals over the next decade and how exascale computing would play a role in achieving them. Key topics included a deeper understanding of the properties of the Higgs boson and its implications for the fundamental laws of nature via collider experiments (HEP's Energy Frontier); the construction of a new scientific facility at Fermilab that will host an international effort to understand the neutrino sector with unprecedented precision (Intensity Frontier); and large-scale sky surveys to investigate the nature of dark energy and dark matter, the neutrino mass sum, and the origin of primordial fluctuations (Cosmic Frontier).

Following the scientific overview, the computational use cases were divided into two areas: compute-intensive, referring to the use of HPC resources in accelerator modeling, lattice quantum chromodynamics (QCD), and computational cosmology; and data-intensive, referring to the use of HPC, as well as other computational resources (e.g., cloud computing), to tackle computational tasks related to large-scale data streams arising from experiments, observations, and HPC simulations. The use of HPC resources is well established in the compute-intensive arena, whereas the data-intensive set of use cases is an emerging application for HPC systems. The evolution of the above computational streams, and their interaction, will function as an important driver for a future exascale environment that encompasses elements of computing, data transfer, and data storage. For this reason, the review participants not only discussed the computing requirements, but also articulated how the entire computational fabric (e.g., networking, data movement and storage, cybersecurity) needs to evolve in order to best meet HEP requirements. All of the HEP facilities plan enhancements that will significantly increase the data volume, velocity, and data complexity, as well as lead to an increased demand for much-improved modeling and simulation of scientific processes. A common concern related to these plans is whether the current scientific computing and data infrastructure will be able to handle the impending demand for computational resources.

Review participants later assembled into breakout sessions to discuss the key issues associated with the demands imposed on compute-intensive and data-intensive computing resources. These discussions amplified the points made in a series of white papers and case studies (Appendices C and D) prepared by the HEP community in advance of this review.

The review provided a rare opportunity for experts from all of the ASCR facilities and ASCR research and development (R&D) programs and the HEP scientists to interact as a community and learn about each other's expertise, the challenges faced, and the exciting opportunities made possible by the exascale computing environment. The review generated the following specific findings, which are detailed in the body of this report.

- Compute-intensive applications in accelerator modeling, lattice QCD, and computational cosmology continue to remain a critical need for HEP, and their requirements are expected to escalate. These groups have the scientific needs and the expertise to use next-generation HPC systems at extreme scale.
- HEP is expanding its range of HPC expertise to data-intensive use cases. The expansion is driven by large data streams from simulations and from experiments. The HEP experimental software stack, however, is currently not well-suited to run on HPC systems and will need to be significantly modified to use these resources more effectively.
- A strong partnership between ASCR and HEP is necessary to optimize systems design for the exascale environment, taking into account the range of requirements set by the compute-intensive, as well as the data-intensive, applications.
- Long-range planning of computational resources is of paramount importance to HEP programs. Stability of ASCR systems and services (which are typically allocated annually) must align with HEP timescales to allow for proper planning and to optimize design and use of ASCR resources for HEP science campaigns.
- In many ways, the current ASCR and HEP computing facilities are complementary: ASCR resources are more focused toward compute-intensive applications, while HEP facilities are more oriented toward data storage, high-throughput computing (HTC) (the grid), and complex scientific workflows. As the two interact more closely in a future driven by mission needs, it is important that they evolve in a coordinated and mutually beneficial fashion.
- Effective uniformity of the computing environment across the ASCR facilities would be very helpful to the user community and important to simplifying engagement with HEP experimental workflows. A federated authentication/cybersecurity model would be very desirable.
- As part of establishing useful partnerships, it was considered desirable to have ASCR staff partially embedded in HEP science teams.
- To enhance opportunities for establishing long-term partnerships across the ASCR and HEP programs, a pilot ASCR/HEP institute was suggested to foster long-term collaborations that develop capabilities benefiting both programs.

As a broader message, participants agreed that a mission-oriented ASCR/HEP partnership for longrange activities should be initiated.

Prior to the meeting, a number of white papers and case studies were collected that describe, in broad terms, the current computational activities across the HEP science spectrum (white papers) and then detail specific examples (case studies) to provide concrete data in a number of representative applications. The case studies cover both compute-intensive and data-intensive examples and vary in size from “mid-scale” to “extreme-scale” applications (by this terminology, we mean the distinction between small [mid-scale] and large [extreme scale] computational footprints on facility-sized computing platforms). Both the white papers and the case studies project needs in their areas on the 2020–2025 timescale.

Following the review, the organizing and steering committees prepared this detailed report, which describes the HEP community’s vision and grand challenges (Section 2), HEP computing needs and requirements (Section 3), and the path forward for HEP and ASCR. The appendices contain meeting materials, including a list of the HEP Exascale Requirements Review organizers and participants (Appendix A), the meeting agenda (Appendix B), and the white papers and case studies prepared in advance of the review and presented in the format submitted by the community (Appendices C and D).

[HEP 2.2]**2.2 Computing Ecosystem**

Computing has always been — and continues to be — an integral and essential part of all activities in high energy physics, which is a data- and compute-intensive science. As the field of particle physics advances, computing must continue to evolve in order to satisfy the ever-increasing appetite of its researchers. For example, the volume of physics data from LHC experiments already stresses both the computing infrastructure and related computational expertise, and LHC operations in the next decade will likely result in order-of-magnitude increases in data volume and analysis complexity. The data needs associated with experiments exploring the cosmos will greatly expand as vast new surveys and high-throughput instruments come online. Cosmological computing is making significant progress in connecting fundamental physics with the structure and evolution of the Universe at the necessary level of detail. Lattice QCD provides the most precise values of heavy quark masses and the strong coupling constant, important for using Higgs boson decays at the LHC to test the Standard Model and probe new physics. Precision perturbative QCD and electroweak calculations are also beginning to use large-scale computational resources. The use of HPC is advancing full 3-D simulations of nearly all types of accelerators. The aim is to enable “virtual prototyping” of accelerator components on a much larger scale than is currently possible, potentially alleviating the need for costly hardware prototyping. All of these use cases will see significant enhancements and improvements in the future.

The HEP community fully recognizes the significance of exascale computing and is counting on its availability as an essential resource for the next set of major HEP projects early in the next decade. Broadly speaking, there are two types of HEP use cases relevant to the exascale computing environment. The first are compute-intensive, referring to the use of HPC systems in HEP applications such as lattice QCD, computational cosmology, and accelerator modeling. The second are data-intensive, referring to the use of HPC and other computational resources (e.g., analysis clusters, cloud resources, HTC systems) to manage and analyze data streams arising from experiments, observations, and HPC simulations. As the quality of simulation and modeling improves, and the associated datasets become more complex, a number of applications are emerging that merge the two categories by being simultaneously compute- and data-intensive.

Historically, the HEP theoretical community has been a long-standing customer of ASCR HPC resources, a relationship that has been further nurtured by partnerships initiated and continued by the Scientific Discovery through Advanced Computing (SciDAC) program. The situation with HEP experiments has been somewhat different; for the most part, experiments in all three HEP Frontiers have been responsible for handling their own computational needs, which have been mostly satisfied using distributed HTC (“the grid”).

The data from HEP experiments is growing at a rapid rate in terms of volume and throughput. It is now doubtful that experimental HEP will, as a field, own all of the computing resources it needs moving forward. Provisioning will be, at best, capable of meeting steady-state demands — the HEP community will need to find creative ways to meet peak demands. At the same time, to extract the best possible results from the observations, joint analysis with sophisticated theoretical predictions is needed; such analysis comes with its own set of large-scale data and computational needs.

The emerging exascale computing landscape offers many opportunities for HEP both in terms of compute-intensive HPC use cases, as well as the data-intensive use cases — for example, using HPC machinery (or what it evolves into) to perform data reconstruction and simulation tasks that have historically fallen within the domain of HTC systems in Energy Frontier experiments. The data-intensive community is increasingly becoming aware of the value of the HPC resources and expertise available. These use cases bring new challenges in the form of data movement

and storage, data persistence, and integrating these technologies into the scientific workflows of the experiments.

A number of white papers and case studies are presented in Appendices C and D. The white papers provide broad coverage of HEP science and describe the science cases and modes of computational usage. The case studies present individual examples that provide more details about computational methodologies and requirements.

The aim of this Exascale Requirements Review is to understand the requirements that must be satisfied for ASCR facilities to be able to meet HEP needs through the 2020–2025 timeframe. The approach taken in this review is broad, covering the entire range of HEP activities and taking into account not just ASCR computing systems, but the broader environment of data storage and transfer, software evolution, and, the desired characteristics of the exascale environment in the face of HEP’s data-intensive needs, a relatively new requirement. In this broader context, the specific examples discussed above are connected to a number of important issues that are central to the review. These are:

- **Simulations:** The ability to effectively perform simulations is paramount for successful HEP science. As experimental and observational datasets grow — with ever-higher-resolution detectors — the demand for finer-grained, higher-precision simulation will continue to increase, requiring an increase in computational resources and the overall efficiency with which they are used.
- **Architectures:** In order to satisfy increasing computational demands and keep abreast of technological changes, the field needs to make better use of next-generation computing architectures. With increased concurrency at the nodal level and more complex memory and communication hierarchies, the complexity of software and systems will continue to increase, and such systems need to be better exploited and managed.
- **Data Access:** Distributed storage access and network reliability are essential for data-intensive distributed computing, a hallmark of HEP experimental practice and a growing issue for other applications. Emerging network capabilities and data access technologies improve researchers’ ability to optimally use resources independent of location. Treating networks as a resource that needs to be managed and planned for is an important area of future ASCR and HEP interaction. This will require close collaboration between HEP scientists and networking experts from the Energy Sciences Network (ESnet).
- **Training:** As a field, HEP, across all frontiers and applications, must continue to develop and maintain expertise and re-engineer frameworks, libraries, and physics codes to adapt to the emerging hardware landscape, as mentioned above. There is a large code base that needs to be re-factored and re-engineered, and there is a shortage of trained experts to do it. Training the next generation of HEP computational scientists and software developers requires urgent attention.

[HEP 3.1]

3.1 HPC Applications

The category of compute-intensive tasks in HEP belongs primarily to the HPC realm. The principal use cases of HPC within the HEP community include groups working on accelerator modeling, cosmology simulations (Figure 4), and lattice QCD; all have been long-time partners with the ASCR community. These users possess considerable computational sophistication and expertise; they leverage significant state-of-the-art computing resources to solve, at any one time, a single, complex, computationally demanding problem.

Very large datasets can be generated by these HPC applications — potentially much larger than those from any experiment. Analyzing these datasets can prove to be as difficult a problem as running the original simulation. Thus, data-intensive tasks currently go hand in hand with many HPC applications that were originally viewed as being only compute-intensive.

[HEP 3.2]

3.2 HTC Applications

In contrast to HPC applications, the data-intensive use cases in HEP experiments exploit HTC systems to take advantage of the inherent parallelism of event-centric datasets. HEP teams have built large computing grids based on commodity hardware in which each computing node handles a computing problem from start to end. A typical example of an HTC application is event simulation (Figure 5) and reconstruction in Energy and Intensity Frontier experiments.

The HEP community that has primarily focused on HTC applications is actively considering the use of HPC resources. Sherpa, for example, has been ported to HPC systems, and ATLAS and CMS have ongoing efforts that use HPC resources for event simulation. It is fair to state, however, that the community is at an early phase of learning how to leverage HPC resources efficiently to address their computing challenges.

[HEP 3.3]

3.3 Future Architectures and Portability

A common problem that both the compute- and data-intensive communities within HEP face is the possible proliferation of “swim lanes” in future computational architectures and the difficulty with writing portable code for these systems. Currently, and in the next generation of large ASCR systems (“pre-exascale”) (Figure 6), there are only two types of computational architectures available: CPU/GPU (accelerated) and many-core (non-accelerated). While HPC users can imagine running codes on these two types of architectures — and even this is limited to only a few teams — data-intensive users have a much more difficult planning decision to make. Disruptive changes cannot be made often to the HEP experiment software stack, and even then, only with considerable difficulty. This means that it is very likely that future evolution of this software will follow conservative trends, which for now, appear to lead down the many-core path. Although we cannot predict the detailed future of the exascale environment with precision, from what is currently known, this strategy would appear to make sense. The above argument suggests that a parallel effort in developing portable programming models for the exascale would be particularly beneficial for data-intensive applications. These are not typically characterized by a few highly-tuned kernels, but by a number of chained subprograms that may not be individually tuned for performance (nor is there usually an attempt to apply global optimization). Therefore, in this case, portability may not necessarily be accompanied by an unavoidable loss in performance, as is the case for the vast majority of HPC applications.

[HEP 4.1]

4.1 Challenges Ahead for HEP

In this section, we cover specific technical, organizational, and other challenges facing the HEP community as it moves forward to take advantage of the next generation of computational systems and data storage and data transfer technologies. Exascale users will face a number of difficulties, such as stability in hardware and software, system resilience issues, evolving programming models, complex system hierarchies, etc. — some of which are merely extensions of current problems and some of which will be new to the exascale environment. Here, we focus on the issues that are particularly relevant to HEP applications.

[HEP 4.1.1]**4.1.1 Compute-Intensive Applications**

The computational challenges ahead for the HEP community have been broadly outlined in the white papers. The papers cover, in principle, most of HEP's HPC-oriented computational practice. The compute-intensive areas are covered in the following white papers: *C.1 Exascale Accelerator Simulation*, *C.2 Computational Cosmology at the Exascale*, *C.3 Lattice QCD*, and *C.4 HPC in HEP Theory*. These white papers discuss current science topics and computational approaches, as well as future plans.

From the point of view of computational needs, given the science targets, the raw demand in terms of core-hours (from all of the white papers, individually and together) seems to be very reasonable and completely consistent with the expected availability of computational resources at the Leadership Computing Facilities (LCFs) and National Energy Research Scientific Computing Center (NERSC) on the 2020–2025 timescale; the HEP requirements, in terms of core-hours, amount to roughly 10% of the total that is likely to be available. It is important to note, however, that the different applications make demands on HPC systems in different ways: computational cosmology and lattice QCD have the highest computational density, and computational cosmology codes are also memory-limited (as are a few examples in accelerator simulation). To run the problems at the required scales, even next-generation systems are on the smaller side of the requirements. The HEP theory use case is more of a mid-scale computing example and makes relatively light demands on the HPC system network.

The HEP compute-intensive community is abreast of the current computational state of the art, although new algorithm development continues for a variety of tasks. Other work includes performance optimization and adoption of better software engineering practices. The evolution to next-generation systems is less clear, however. Certain teams within the lattice QCD and computational cosmology groups already have extensive experience with GPUs and many-core hardware, but these two options should be viewed as a disruptive set of technologies for most applications. The SciDAC program and other opportunities, such as the Early Science Projects at the LCFs and the NERSC Exascale Science Applications Program (NESAP), will help accelerate the adoption of new programming paradigms.

Participants agreed that what is needed in this area is a set of common tools across architectures — efficient parallel primitives — that are separately optimized for each. These tools should possess flexible data layouts so that codes can adapt to different memory hierarchies. Examples include operators such as shift, scan, reduce, transform, scatter-gather, geometry remapping, etc. Moreover, it should be possible to easily replace primitives with hand-tuned code, when needed. It would be useful to compile a list of desired primitives and define appropriate interfaces. It is important to note here that there is no desire for heavy-duty frameworks and libraries that would be hard to develop, maintain, and modify. The emphasis is primarily on lightweight tools, especially because the future computing environment is expected to evolve considerably. A possible suggestion is that a few application codes should function as testbeds for developing the envisioned toolkits. In addition, emphasis should be placed on the benefits of having better network fabrics on the HPC systems and on increased memory/node. At a lower level, there is some concern directed at the problems of node-level memory management. A common interface for managing data placement and movement across different levels of the memory hierarchy would be very desirable.

As noted in two of the white papers (C.2 and C.3), data storage and transfer will become significant issues. In fact, it can be argued that for the compute-intensive field of computational cosmology, the point where current needs are not being met has already been crossed.

Finally, as mentioned in Section 3.1, compute-intensive applications can generate very large amounts of data — potentially significantly larger than the data volume from the experiments — that then must be processed in different ways. This issue is of particular concern to the lattice and computational cosmology groups, who either have their own resources for off-line analysis and storage (lattice) or are also leveraging other approaches, such as in-situ analysis and use of software containers and virtualization (cosmology). This general area is currently characterized by a number of ad hoc approaches; achieving production-level status in the next 2 years is an important goal, so as to keep in sync with the arrival of next-generation systems at the LCFs and NERSC.

[HEP 4.1.2]

4.1.2 Data-Intensive Applications

Up to this point, the data-intensive HEP community has been very successful in terms of its ability to leverage computing technology to carry out its science. Yet, there are significant challenges ahead that must be overcome before HEP groups can effectively exploit large-scale HPC resources for many of their applications. The relevant white papers here are *C.5 Energy Frontier Experiments*, *C.6 HEP Experiments: Data Movement and Storage*, *C.7 Cosmic Frontier Experiments*, and *C.8 Intensity Frontier Experiments*. Currently, the Cosmic Frontier and Intensity Frontier experiments generate relatively modest amounts of data, and the primary use cases are the two LHC experiments. In the future, the other Frontiers will be characterized by significantly larger data volumes, but they are unlikely to stress the available resources in a way comparable to the LHC needs. Except for a few cases in the Cosmic Frontier, all of HEP’s data-intensive computing uses the grid model; forecasts of future computing needs versus what the HEP grid can provide are making it very clear (see white papers C.5 and C.7) that access to alternative computational resources will be necessary.

With some very rare exceptions, it is not an easy task to run the HEP experiment software stack “out of the box” in an efficient way on HPC systems. The current code base is not well-suited to run on HPC machines (designed for HTC and not for HPC resources, lack of vectorization, threading, unsuitable code organization, heavy I/O with small files), and it will need to be significantly modified and appropriately refactored to use these resources efficiently.

This is a significant problem, and the HEP workforce as it stands does not possess the skill set to tackle it on the required timescale. Training — through summer schools, boot camps, and mentors — will therefore become increasingly important. ASCR expertise can play a significant role in helping to modernize the HEP workforce in this area. An effort is now underway to develop a set of “mini-apps” [24] through the HEP-FCE — for both data- and compute-intensive applications — to help HEP researchers understand the new computing environment and ensure that HEP codes can scale appropriately with reasonable performance levels. The mini-apps can also be used as testbeds for prototype exascale hardware to ensure that HEP applications and workflows will perform adequately.

More generally, however, the problem is that it is simply not possible to train thousands of HEP researchers to become HPC experts, nor is it desirable. The goals of the training program would be (1) to develop a core of expertise within HEP, which would allow structured refactoring of the current code over time; and (2) to construct usable frameworks that would allow thousands of scientists to run large-scale analysis or other data-intensive tasks without having to become HPC experts.

The complexity of the current HEP experiment software base (many millions of lines), a substantial fraction of which needs to be available at the level of the compute nodes, and the fact that it changes on relatively short timescales, are significant problems for HPC systems. Fortunately, these problems can be substantially mitigated by the use of software containers, such as Docker. Although HPC systems currently do not offer this capability, the hope is that they soon will.

One of the significant issues facing data-intensive applications in the exascale environment involves incorporating HPC machines into the experiments' workflows. One difficulty is that Energy Frontier applications require real-time remote access to resources such as databases while the associated jobs are running. Because of the fine-grained nature of the computational tasks, the associated I/O bandwidth requirements are a potentially significant issue and need to be properly characterized, not only for the systems being used, but also at the facility scale. A second difficulty is that HPC systems are scheduled for maximal utilization in a way that is problematic for largescale complex workflows, especially those that require true real-time access. The problem is that a given workflow may require widely varying resources as it runs, in ways that may not be easily predictable beforehand. This mode of operation is essentially orthogonal to current batch scheduling implementations on HPC systems. Moreover, truly addressing the elasticity requirement implies a substantial degree of overprovisioning, which is hard to imagine implementing at current HPC facilities. Note that having containers available will not solve the problem of elasticity, although it will help by potentially providing the ability to quickly spin up and spin down machine partitions dedicated to a particular workflow.

A significant barrier at the moment is that each of the ASCR facilities is configured differently. There is little uniformity with respect to scheduling jobs, cybersecurity, staging data, etc. It would be significantly easier if all HPC systems could be accessed via a uniform fabric. "Edge Servers" needed to stage data outside of the facility firewall will play an important role in data-intensive applications, including potentially addressing access control issues on HPC systems. A uniform implementation of access protocols would make interfacing to the facilities much simpler. Lastly, uniformity in cybersecurity protocols across the facilities would be very desirable. A single, federated model that works at all three centers would be ideal.

The volume of data produced by experiments is rising rapidly. The current LHC output is on the order of 10 PB/year, and in the HL-LHC era (2025), this number is expected to rise to 150 PB/year. No experiment in the Cosmic or Intensity Frontier will come close to this data rate or total volume (exabyte [EB] scales). For the most part, one can expect that if the volume is sufficiently large, the projects will budget for and purchase the required storage. Below a certain threshold, the ASCR HPC centers provide disk storage in parallel file systems and much larger archival tape storage managed using HPSS (High-Performance Storage System). The timescale over which this storage continues to exist is subject to success or failure in the ASCR Leadership Computing Challenge (ALCC) and Innovative and Novel Computational Impact on Theory and Experiment (INCITE) process (and to a much lesser extent at NERSC). In any case, if the HPC systems are used as a data-intensive computing resource, much more attention will have to be paid to storage resources in terms of capacity, latency, and bandwidth (as is also the case for compute-intensive applications that generate large amounts of data, as described in Sections 2.2 and 4.1.1). In particular, high-performance, reliable, and predictable data transfer across multiple storage endpoints would be an essential requirement.

[HEP 4.2]

4.2 Potential ASCR/HEP Paradigms

In the context of a future exascale environment — or even in the near term — a number of needs for the HEP computing program have been identified thus far, as described in Section 3. These needs include programming on new architectures and associated algorithm development, data storage and transfer, refactoring of a potentially significant fraction of the code base (over ten million lines of code), co-evolution of ASCR and HEP facilities, and training of HEP personnel. To address these needs, the HEP community is very interested in identifying and exploring new ways to partner with ASCR. The exascale environment opens up exciting new scientific opportunities; the knowledge and expertise that ASCR has are invaluable to the HEP community in its pursuit of new scientific opportunities.

One difficulty is that current partnership opportunities are limited under today's organizational structure. Outside of the SciDAC programs and informal contacts, there are limited established mechanisms to develop long-term partnerships between the two communities. This is an area that calls for serious attention.

A possibility discussed at the Exascale Requirements Review is to enhance or redirect partnerships where some entity — like a joint ASCR/HEP institute (and similar entities for other sciences if appropriate) — would provide multi-faceted partnerships with ASCR research and facilities divisions. These new entities would have broad-ranging computational expertise and be dedicated to the scientific success of the community they are supporting, forming valuable additions and expansions to the successful SciDAC Program. As an important example, the fact that a complex HEP software suite needs to run “everywhere” requires an approach whereby staff belonging to a joint ASCR/HEP group of experts (“embedded power teams”) work together to solve the associated portability problems within the larger computational context of a given project. In this mode, scientists across different domains would get to know each other, establish a common language, and have enough long-term commitments to tackle difficult problems. In other words, the broader computational support model should involve working with large collaborations, as well as with small teams.

Another issue involves long-range planning to ensure availability of resources; such planning is essential for HEP experiments because of their extended lifetimes of 5 years or more (Section 4.5). As outlined in the white paper C.5 Energy Frontier Experiments, “it is possible that by 2025, most processing resources will be supplied through dynamic infrastructures that could be accessed opportunistically or through commercial providers.” If such a scenario is to become reality, close interaction will be needed not only between the ASCR computing facilities and HEP counterparts to support programmatic needs, but also between HEP and ESnet to make sure that networking resources are adequately provisioned. Joint design of “Edge Servers” (Section 5.1.2) and co-evolution of HEP facilities (Section 4.5.) would also be facilitated by such an interaction. To summarize, a mission-oriented ASCR/HEP partnership for long-range activities needs to be initiated.

[HEP 4.3]

4.3 Summary of HEP Requirements

The purpose of this section is to summarize the quantitative estimates of computing, storage, and networking that are presented in the white papers. The white papers addressed (1) the individual compute-intensive HEP applications (accelerator modeling, computational cosmology, lattice QCD, and HEP theory); (2) data-intensive computing requirements from experiments at the three frontiers; and (3) a separate white paper on Energy Frontier experiment data movement and storage and on the evolution of HEP computational facilities. The authors of the white papers were asked to collate information from all the major use cases in their domains, and in particular, to provide an assessment of the science activities and computational approaches on the 2020–2025 timescale, as well as estimates for the computing, data, and services needs over the same period. The case studies presented in Appendix D are meant to provide more in-depth examples of specific use cases, which can range from medium-scale computing to full-machine exascale simulations.

It is not easy to assess requirements roughly a decade into the future for a diverse set of scientific applications because new breakthroughs and unexpected obstacles are, by definition, difficult to predict. The difficulty is compounded by the fact that two architecture changes are expected over this timeframe: the change from multi-core to many-core (GPUs included in this classification) at all the major ASCR computing facilities (pre-exascale) followed by another architecture change at the exascale/beyond-exascale level, which could be significantly disruptive. In addition, computational needs in different areas do not have equal priority. Much depends on which

computations are considered more relevant at a certain time for HEP science and which can be pushed into the future. Computational needs driven by ongoing projects and those planned or under construction have been taken into account in the white papers. The future, however, may turn out differently if the associated timelines change in ways that lead to lack of coordination with the evolution of the computational resources.

Given these caveats, it is nevertheless encouraging to note that the compute-intensive HEP applications have kept pace with hardware evolution — including key pathfinder roles in several cases (e.g., lattice QCD for the IBM Blue Gene systems and computational cosmology for Roadrunner, the world’s first petaflop system). In addition, as noted in the white papers, the science drivers show no sign of letting up in their hunger for computational resources. The situation in the case of experiments is more difficult to assess. In the case of the Cosmic Frontier, it is becoming clear that the future will lie in using ASCR facilities (NERSC is the host for DESI and also the likely host for LSST-Dark Energy Science Collaboration [DESC] computing; the LCFs will provide additional resources) as the dominant source of computational resources. HEP Energy Frontier experiments have historically not used HPC sites for computing, although this is rapidly changing. Intensity Frontier experiments have not been major consumers of computational resources, but this situation is also changing, even on the current timescale. How quickly all of these areas will be able to take full advantage of ASCR resources depends on the pace with which the relevant components of the HEP production and collaboration software are refactored, as well as how the ASCR facilities evolve — in their turn — to address the HEP use cases.

We now consider the computational, data, and networking requirements that can be extracted from the white papers. Unless explicitly stated, the numbers in this section (in core-hours for a standard 2015 X86 core) refer only to the computational requirements associated with HEP use of ASCR facilities. The disk storage requirements are given as an aggregate number, but it is conceivable that this storage can be split between ASCR and HEP facilities, depending on the particular use case. Wide area network (WAN) bandwidth requirements are given in a few cases in which we expect that the requirement will potentially stretch ESnet capabilities. Local bandwidth within the facilities is assumed to be at least as good as the wide area requirement.

Accelerator modeling can be divided into the following: (1) electromagnetics and beam dynamics simulations for current and near-future technology machines and (2) dedicated simulations to help in the design of future accelerators. Currently, both electromagnetics and beam dynamics consume on the order of 10M core-hours annually each. These numbers are expected to scale up to 10–100 billion (G) core-hours by 2025, depending on the use cases being run. Large-scale simulations for future accelerators (machines that may be built on the 2030+ timescale) focus on plasma-based acceleration schemes. While these simulations are currently at the 10M-core-hours level, they can scale up to an annual requirement of 1G–100G core-hours (or more) by 2025, but there is significant uncertainty regarding the upper value. Storage (or networking) has historically not been a major issue for accelerator modeling, and this is unlikely to change in the future.

Computational cosmology will have to support a large number of Cosmic Frontier projects, some of which are quickly reaching the scale of HEP experiments in terms of collaboration size. Current annual simulation usage is at the 100M–1G core-hours scale and is expected to increase to 100G–1,000G core-hours by 2025. In addition, large-scale cosmology simulations are already memory-limited, and they are likely to saturate the system memory of machines in the exascale era. Storage requirements are likely to be large; they are already at the level of 10PB of disk storage, and they are likely to easily exceed 100PB by 2025. Furthermore, because large-scale distributed analysis will be needed by the collaborations, there will be significant networking requirements. Currently, a pilot project with ESnet is aiming to establish a 1PB/week production transfer rate for

moving simulation data. By 2025, the burst requirements will be approximately 300Gb/s, which is roughly the same scale as that required by Energy Frontier experiments.

Lattice QCD has a long history of efficient use of supercomputing resources, and this trend will continue into the foreseeable future. Current annual usage is in the 1G~core-hour class; this number is expected to increase to 100G–1,000G core-hours by 2025. Memory requirements (unlike for computational cosmology) are nominal. Disk storage, which has historically not been a major requirement, will only increase slowly, from ~1PB currently to ~10PB by 2025. Theory requirements (event generation, perturbative QCD) are at 1M–10M core-hours currently; these will likely increase to 100M–1G core-hours by 2025. Memory and storage needs for this effort will likely remain at easily satisfiable levels.

Cosmic Frontier experiments are currently running at the 10M–100M core-hours scale on HPC resources; on the 2025 timescale, this is likely to increase to 1G–10G core-hours. Disk storage requirements are currently at the ~PB scale and are likely to increase to 10–100PB by 2025. Network requirements are unlikely to stress future capabilities at the same level as Energy Frontier experiments.

Energy Frontier experiments have begun using HPC systems relatively recently, primarily for event simulation tasks. However, annual usage has already reached the 100M core-hour level on ASCR resources. This usage level should be compared with the total U.S. contribution to LHC computing, which is on the order of 100M–1G core-hours annually on HTC systems. By 2025, the requirement on HPC resources could reach 10G–100G core-hours, with extensive storage needs, exceeding 100PB of disk space (the total global storage requirement will reach the exabyte scale). The network requirements will also be high, at the level of 300 Gb/s (more in a continuous mode, rather than burst operation).

Intensity Frontier experiments are at the ~10M core-hour level of annual usage. This number is expected to increase to 100M–1G core-hours annually. Storage requirements are at the ~PB level currently and are expected to increase to 10–100PB by 2025.

The information discussed in this section is summarized in Table 1. The sum of HEP requirements, although difficult to pin down precisely because of the uncertainties discussed, is projected to be ~10% of the expected ASCR facility resources at the ALCF, NERSC, and OLCF by 2025.

[HEP C.2, Almgren et al.]

C.2 Computational Cosmology at the Exascale

Authors: A. Almgren, K. Heitmann (leads), N. Frontiere, S. Habib, Z. Lukic, P. Nugent, B. O’Shea

C.2.1 Current Science Activities and Computational Approaches

Large-scale simulations play key roles in cosmology today, including: 1) exploring fundamental physics and probes thereof, e.g., dark energy, dark matter, primordial fluctuations, and neutrino masses, 2) providing precision predictions for a range of cosmological models, important for data analysis, 3) generating detailed sky maps in different wavebands to test and validate analysis pipelines, 4) understanding astrophysical systematics, e.g., baryonic effects, and 5) providing covariance estimates.

Cosmology simulation codes fit into two main categories: gravity-only (“N-body”) solvers and “hydrodynamics” codes that also include gas physics and associated subgrid modeling (e.g., cooling, astrophysical feedback, star formation). The tasks above require four kinds of simulations using a mix of these capabilities, depending on the specific application: 1) gravity-only

and hydrodynamics simulations over a range of cosmological models to address the first and second tasks, 2) very high resolution large volume gravity-only simulations (for, e.g., large galaxy catalogs) and medium resolution large volume hydrodynamics simulations (for, e.g., thermal Sunyaev-Zel'dovich maps and Lyman-alpha investigations) to address the third task, 3) very high resolution hydrodynamics simulations including treatment of feedback effects to address the fourth task, and 4) a very large number (well beyond thousands) of moderately accurate gravity-only simulations to address the fifth task.

Approaches used today for the gravity-only solvers include particle-mesh plus short-range solvers (particle-particle or tree methods) [1, 2], pure tree methods [3, 4], and pure grid methods. Codes that include hydrodynamics coupled with an N-body representation of dark matter include grid-based hydro methods, typically using Adaptive Mesh Refinement (AMR) [5–8], Smooth Particle Hydrodynamics (SPH) [9], and Moving Mesh Methods [10].

Analysis of the data generated in the simulations is fundamental to addressing the research goals. Currently most of this occurs in post-processing, but the large amount of data from future simulations and increased computational expense due to more complex analyses will result in more reliance on in situ approaches.

C.2.2 Evolution of Science Activities and Computational Approaches on the 2020/2025 Timescale

Figure 1 is an overview of the cosmological surveys that will dictate our science activities until 2025. Future activities will be similar to those today, but the simulations will have to keep up with observational improvements.




Figure 1: Timelines for cosmological surveys (blue: CMB, black: optical, ground-based, red: optical/NIR satellite) and supercomputing resources as well as simulation requirements. DOE HEP plays a major role in the ground-based optical surveys and in some of the CMB surveys.

Requirements are most stringent for optical surveys since they probe much smaller scales than CMB surveys. With increasing depth, fainter galaxies at larger distances need to be resolved, requiring larger simulations (larger volumes and more particles) with more detailed physics implemented in the hydrodynamics codes.

For gravity-only simulations, the ultimate goal is to have enough mass resolution to resolve dark matter halos that host dwarf galaxies in a cosmological volume. These simulations will be needed for surveys such as LSST. In the next decade, we want to cover volumes of several Gpc and achieve a mass resolution of 10^6 – $10^7 M_\odot$. This means that we have to simulate up to a hundred trillion to a quadrillion particles, leading to memory requirements of ~ 4 – 40 PB. In addition, the requirement to capture small-scale structure in large-volume simulations will demand finer force resolution, leading to very compute-intensive runs. Although the general approach to gravity-only simulations will likely not change much in the next decade, two main challenges exist: 1) global load balance and efficient communication to enable use of large parts of supercomputers (e.g., FFTs), 2) local load-balancing to follow the formation of complex sub-structure. Despite these challenges, no other major roadblocks are anticipated in the future.

In the case of grid-based hydrodynamics, scientific targets include modeling the Lyman-alpha forest at scales relevant to baryon acoustic oscillation measurements (with box-sizes of $\ll 1\text{Gpc}$) while maintaining resolution to resolve density fluctuations responsible for the forest ($\ll 10\text{kpc}$). This leads to memory requirement in the range of 4-64 PB. Another challenge in modeling the forest arises from the fact that with future precision requirements the ionizing background can no longer be treated as uniform; radiation transport – and probably multigroup radiation transport – is going to become the norm. This will be computationally costly and will present enormous scaling challenges, depending on the method used. Similar requirements arise in other areas of cosmological hydrodynamic studies, e.g., in the study of clusters of galaxies, or the evolution of galaxies. The resolution needed to build physically reliable subgrid models for star formation and feedback, as well as AGNs, is much more stringent and of the order 100pc, bringing again the total memory requirements into the PB range.

Improvements in grid-based hydrodynamics codes will focus both on on-node performance and load balancing. We will need to make effective use of all the cores on the new many-core nodes subject to low-memory per core and on-chip NUMA effects. This will require new approaches for working on domain-decomposed or block-structured AMR blocks of data using finer granularity. “Logical tiling”, used to control the working size of the block of data being operated on, can improve performance due both to improving the use of cache and allowing threading over blocks rather than loops. “Regional tiling” alters the layout of the data on a node by optimizing for the on-node NUMA effects. Both of these strategies fit well into the AMR paradigm in which the cost of metadata is minimized by keeping the size of individual AMR grids large.

SPH codes provide a different approach to the hydrodynamics problem. They have the advantage of computational efficiency compared to AMR codes but have suffered from problems such as lack of mixing in the past. Significant progress has been made recently with regard to a number of the accuracy concerns and new SPH techniques show great promise. In terms of implementations on future architectures, these improved SPH methods will be an attractive option. More work is needed in assessing the accuracy issues and also in the implementation of these new approaches on future architectures.

In addition to the performance-improving paths identified above, a major development component will be in sub-grid modeling. Here, uncertainties are currently very large and ensembles of runs will have to be carried out to understand the implications of different modeling assumptions in detail.

Finally, as the amount of data generated by large simulations increases, we will move from primarily using post-processing for diagnostics to a mode of in-situ data analysis. This addresses the data storage issue but requires additional run-time optimization of the diagnostic routines as they will either compete with the simulation code for resources or require additional data movement to cores that will not compete with the simulation itself.

C.2.3 Compute, Data, and Services Needs on the 2020/2025 Timescale

In general, all cosmological simulations are memory-limited. The memory requirements on next-generation supercomputers will be in the tens of PB for gravity-only simulations. Each time step would produce tens of PB of data, usually around 100 snapshots are needed for a complete analysis. This amount of data will probably go beyond the available resources in 2020/2025, therefore efficient in-situ analysis frameworks will be essential. In addition to a handful of very large simulations, we will also need to carry out suites of simulations of medium size.

For the hydrodynamic simulations, the memory requirements per particle or grid element are much higher than for the gravity-only solver because of the additional fields being evolved. For both the gravity-only and hydrodynamics runs, the size of the largest runs will be dictated by the total

memory available on the supercomputer. As mentioned above, ensembles of runs will be carried out to explore the effects of different sub-grid models.

The current usage for cosmological simulations is roughly 400M core-hours at the LCFs and NERSC. Simulation requirements for surveys are still being processed by the community but are expected to be very substantial. The demand will scale faster than the available compute because of the need for multiple runs. The hydrodynamics runs will add a multiplicative factor of 20 beyond the N-body requirements. Finally, while some of the science can be accomplished using run-time diagnostics that do not require the storage of the solution, a large community would like to use the results for a variety of additional science tasks. This means that the storage of the output from at least several key runs of different types is highly desirable. Storing and moving such large amounts of data is very difficult and it is not clear 1) how to efficiently serve the data, and 2) how to provide analysis capabilities that can deal with very large datasets. The community has started to address these questions, but with increasing amounts of data a more rigorous and coordinated plan over the next decade is needed to address both of these issues.

- [1] V. Springel, Mon. Not. Roy. Astron. Soc. 364, 1105 (2005)
- [2] S. Habib, et al., arXiv:1410.2805, New Astronomy, in press
- [3] M.S. Warren, arXiv:1310.4502, Proceedings SC'13
- [4] M.D. Dikaiakos and J. Stadel, ICS Conference Proceedings (1996)
- [5] G.L. Bryan et al., Astrophys. J. Supp., 211, 19 (2014)
- [6] A. Almgren et al., Astrophys. J., 765, 39 (2013)
- [7] R. Teyssier, A&A, 385, 337 (2002)
- [8] A.V. Kravtsov, A.A. Klypin, and Y. Hoffman, Astrophys. J. 571, 2 (2002)
- [9] V. Springel, Mon. Not. Roy. Astron. Soc. 401, 791 (2010)
- [10] P.F. Hopkins, Mon. Not. Roy. Astron. Soc. 450, 53 (2015)

[HEP C.6, Wurthwein et al.]

C.6 HEP Experiments: Data Movement and Storage

Authors: F. Wurthwein, G. Oleynik, B. Bockelman

C.6.1 Current Science Activities and Computational Approaches

The landscape of experimental HEP is strikingly diverse. In addition to the large LHC experiments, there are a number of Intensity and Cosmic Frontier experiments at all sizes. While the LHC will continue to be the largest data producer in 2020-25, experiments like DUNE and LSST present their own challenges, and so do smaller experiments.

The ATLAS and CMS experiments produced a few tens of PB of data during LHC's Run 1, 2010-12. Roughly a third of this is from the detector, the rest is simulation. The ~2:1 relationship is likely to be stable; the output rate from the trigger system is thus a rough guide to the increase in data volume over time. Both experiments started with a data taking rate of ~150Hz that increased to ~300-600Hz at the end of Run 1. For Run 2 (2015-2017) the initial rate is 1kHz and is expected to reach ~10kHz by 2025. In the ten years of HL-LHC running (roughly 2025-35) each experiment will transition from O(100) petabytes to O(1) exabyte of data.

Two copies of the RAW data are archived to tape. One is conceptualized as a "backup" copy at CERN, while the other is distributed as an active copy across the Tier-1 centers worldwide for each experiment. Only a single copy of the processed data and simulations is archived. The RAW data from the detector is understood to be the most precious, while all else can be reproduced, if corrupted or lost. Tape is the preferred archival technology because of its durability and lower cost.

It is useful to divide the path from raw data to science results into two steps. First, each collaboration centrally produces official datasets from raw data and simulation. Second, small groups of collaborators produce private custom datasets from (subsets of) the official data, and analyze them to derive publishable results. The first step is consuming ever increasing amounts of CPU power, to produce an output format optimized for maximal flexibility to support the full diversity of the LHC program, allowing for continued improvements in physics object definitions and selections. The second step typically begins with each small group producing “slimmed” and “filtered” custom datasets, resulting in smaller event sizes, fewer events, and substantially faster to process data formats; further analysis tends to be I/O limited. The custom datasets comprised ~4-400TB per individual group during Run 1, and are expected to grow at a similar rate as the official data.

LHC computing was a huge success during Run 1. The speed and robustness with which results were produced exceeds most prior large scale experimental HEP programs, despite the significant increase in scale in data volumes and computing needs. Thus, any proposed change in how data is stored, transferred, processed, and managed, should result in cost savings without loss of speed and robustness in producing physics results. It is worth noting that the LHC program globally involves O(10k) people. At a cost of \$100k per FTE, this amounts to \$1 Billion in global personnel costs, dwarfing the annual LHC computing budgets worldwide. Therefore, a “cost-effective” solution must maximize the effectiveness of the global human capital, or it will invariably lead to loss of efficiency in producing science results.

C.6.2 Evolution of Science Activities and Computational Approaches on the 2020/2025 Timescale

When considering the evolution of scientific and computational activities, we distinguish “technical” and “sociological” opportunities with the aim of identifying more cost-effective solutions as just discussed above.

Among the technical drivers, we identify three high level concepts. First, the trend towards vectorization and parallelization driven by a larger number of simpler cores on commodity/HPC hardware. Second, the advent of Big Data technologies, and third the advent of highly elastic computing. The three combined with the divergence of CPU and I/O needs for the two steps discussed above are likely to drive the need for integrated workflow management across a diverse set of resource types.

In 2025, physics and detector simulation, and raw data processing may be done on three different hardware platforms. Disk buffers in front of tape archives may be minimal in size, as tape retrieval is tightly integrated into the data processing workflow. Such a workflow might be scheduling disk buffers at a remote processing center in addition to disk buffers in front of the tape archive, and the wide area network between the two disk buffers. By 2025, the majority of disk space in ATLAS and CMS may thus be in the analysis systems. These systems may be heavily I/O optimized using disk scheduling ideas from Hadoop/MapReduce, in combination with an I/O layer optimized for partial file reads that is already standard today in ROOT.

Finally, all of the above must be highly elastic. ATLAS and CMS take months to produce releases validated for large-scale simulation and data processing. Once the release is validated, the time for simulation and processing would be shrunk as much as possible, leading to large spikes in desired resource consumption. Since commercial Cloud providers already operate distributed exascale systems, it is natural that ATLAS and CMS will want to use such systems by 2025.

The most important opportunity for cost savings is in placing the divide between the (centrally produced) official data, and custom data produced by the scientists. It may be beneficial to centrally produce data formats that are much less CPU intensive to process, trading flexibility against

reduced size per event. To make up for the lost flexibility, such formats might be produced more often, leading to more versions in use at a given time. Whether this will lead to more human inefficiency than gains in computational efficiency needs to be explored.

The common theme that emerges is that future computing approaches will need to be much more agile, dynamic, and elastic in order to support a more diverse set of hardware resources at scales that frequently change throughout the year.

C.6.3 Compute, Data, and Services Needs on the 2020/2025 Timescale

Data storage and movement services needed to meet HEP needs in the 2020 timescale already exist for the LHC and other HEP experiments, but there are challenges in scaling up by a factor of 10-18 in I/O bandwidth and data volume [3], and to apply these services to collaborations, which though smaller scale in terms of data volume, can benefit from existing infrastructure and architectures. In addition, services are in the process of becoming more agile, but that process is still far from complete. Common facility Services include:

- Long term archival of data (tape storage)
- Dataset services (e.g., data location aware staging service)
- Federated storage access, local posix and WAN data access protocols
- Network infrastructure
- High throughput, low latency access storage for analysis computation
- High throughput, low latency storage for production computation
- Tape backed low latency disk cache
- Global catalog (or mappings to local catalogs)
- Management and monitoring infrastructure

Opportunities and challenges in these service areas are:

- Delivering data is still a major challenge; storage systems do not perform well with random access patterns from multiple users. Services and storage architectures need to convert inefficient patterns into efficient requests to the underlying storage hardware. An example of this is SAM [1], which understands the physical location of files on tape and attempts to optimally stage experiment defined datasets from tape to disk. A related challenge will be to effectively utilize the anticipated 2 GB/s bandwidth of tape drives.
- Smaller HEP collaborations are often limited, not by resource restrictions, but in organizing data to efficiently deliver it to computational workflows. Providing tools and architectures to aid with this could be a great benefit.
- The “storage element” (storage organized as a POSIX-like filesystem with an access interface such as GridFTP) tends to be a too-low-level abstraction, especially when multiple storage systems are involved. The human overhead of maintaining filesystem consistency is high.
- The largest data management systems (CMS’s PhEDEx, ATLAS’s Rucio) have failed to gain widescale adoption outside their respective experiments. A successful reuse example is SAMGrid, adopted by several HEP experiments. This area has had a poor track record of moving computer science innovations from R&D to production. Future experiments would benefit if data management services could be generalized and made production ready.
- The field has rallied around ROOT as a common I/O layer, allowing investments in a single software package to benefit the entire community. However, the core ROOT IO community is

too small; we have the opportunity for significant improvements but not the scale of effort to achieve them.

- Standardized cloud based storage interfaces (S3, CDMI, WebDAV) have not been taken advantage of. Work is needed to assess if they can meet production processing requirements.
- Smaller scale HEP collaborations not directly affiliated with National Labs are often on their own for providing an active data archive and the expertise to manage it. Work is underway to provide storage services to these experiments on a case by case basis. A cohesive set of services for storing and retrieving data at a set of National Labs would be a significant benefit.

Facility based long-term tape storage will be the most affordable option in the 2020 timeframe, less than \$20/TB, and capacity scales better than needed for the HL-LHC run. Tape form factors will likely be the same and libraries will likely still be sized at around 10,000 slots. Tape drive bandwidth is expected to increase by about a factor of $8\times$ to 2 GB/s [2, 4]. Providing data to these drives at full rate will be a challenge. The takeaway for storage is that tape storage and network expenditures will likely be lower, while CPU, disk and tape drive costs will likely be higher than current expenditures. Tape libraries will likely need to be refreshed prior to HL-LHC luminosity running. For more information, see Refs. [3, 4].

- [1] A. Norman, “Data Handling best Practices”, 2015 FIFE Workshop, <https://indico.fnal.gov/getFile.py/access?contribId=35&resId=0&materialId=slides&confId=9737>
- [2] INSIC’s 2012-2022 International Magnetic Tape Storage Roadmap: <http://www.insic.org/news/2012Roadmap/12index.html>
- [3] G. Stewart, CHEP 2015, “Evolution of Computing and Software at LHC: from Run 2 to HL-LHC”, <http://indico.cern.ch/event/304944/session/15/contribution/551/material/slides/0.pdf>
- [4] B. Panzer-Steindel, “Technology, Market and Cost Trends”, 2015 WLCG Collaboration Workshop, <http://indico.cern.ch/event/345619/session/1/contribution/10/material/slides/1.pdf>

[HEP C.7, Borgland et al.]

C.7 Cosmic Frontier Experiments

Authors: A. Borgland, N. Padmanabhan (leads), S. Bailey, D. Bard, J. Borrill, P. Nugent

C.7.1 Current Science Activities and Computational Approaches

The current experiments supported (directly and indirectly) by DOE HEP facilities are cosmic microwave background experiments and low redshift dark energy experiments. The CMB experiments include both space-based experiments (Planck) as well as ground-based experiments; the key science goals are inflation (including a detection of the gravitational waves from inflation), dark matter and dark energy. The dark energy experiments aim to use imaging and spectroscopic surveys to constrain the expansion and growth history of the Universe through weak gravitational lensing and observations of the galaxy distribution (including correlation functions, number counts etc). Exemplars of this are the recently completed Baryon Oscillation Spectroscopic Survey (BOSS), the extended BOSS (eBOSS) survey (ongoing) and the Dark Energy Survey.

Exact analyses of these data are computationally intractable, and therefore one must often rely on Monte Carlo methods to characterize (a) the instrument responses and biases, (b) observational effects from incomplete data/cuts, (c) statistical uncertainties and (d) astrophysical systematics. The generation and analysis of these mock catalogs is often the limiting step in every analysis. For example, to reach 1% statistical uncertainties, the Planck group at NERSC considered 10^4 realizations, each with $O(10^3)$ maps. Low redshift galaxy surveys rely on large numbers of N-body simulations to model the nonlinear formation of structure.

Traditional astronomical data sets are often 10s of terabytes, but are often broken down into very large numbers $O(10^6)$ files, making them unwieldy for storage. Tools to efficiently access these data are either often missing or have yet to see widespread usage. Traditional HPC models are often poorly suited for the analysis of such data sets. An associated challenge is the distribution of these data (and related simulations) amongst large and geographically diverse collaborations.

A third class of experiments supported by DOE HEP facilities are dark matter experiments. These can be divided into three classes – collider production (covered in the Energy Frontier), indirect detection (Fermi Gamma-Ray Space Telescope, not part of this timeline) and direct detection. Direct detection experiments use WIMP-nucleon elastic scattering to put constraints on the WIMP mass. The current Generation-2 program has two main experiments: LUX/LZ (Xenon) and SuperCDMS SNOLAB (Ge, Si) both of which will be located underground to shield them from cosmic rays. They are expected to start operating around 2018-2019.

The analysis of direct detection data closely follows the particle physics model in that there are particle reactions in a detector with associated detector information read out. Because of the low expected WIMP signal rate, a thorough understanding of backgrounds is the critical part of direct detection experiments. Monte Carlo simulations along with dedicated calibration events are used to estimate backgrounds from the detector and associated shielding. Up until now, the data sets from these experiments have been small and computing has not been a priority.

C.7.2 Evolution of Science Activities and Computational Approaches on the 2020/2025 Timescale

The next decade will see an order of magnitude increase (or larger) in data volume and science reach from the next generation of experiments. Each of the three major areas described above have next generation experiments in the planning/construction phase – the CMB experimental community is working towards a Stage IV CMB experiment (CMB-S4) in the early 2020's; the BOSS and eBOSS surveys will be succeeded by the Dark Energy Spectroscopic Instrument (DESI) (~2019-2024), and the Large Synoptic Survey Telescope (LSST) will be the premier imaging survey throughout the next decade (~2020-2030). The science reach of these surveys will require significant increases in computational/storage needs, both for the analysis of the raw data and its cosmological interpretation. We discuss the needs of these surveys individually below. (An important assumption in all of this is that we will be able to maintain computational efficiency on the next generations of hardware, including the expected heterogeneous processing architectures, bandwidth to memory and storage, etc.)

The LSST survey will survey half the sky in multiple bands to unprecedented depths; in additions, it will provide a time-domain view of the sky by surveying it every few days. The primary goals of the HEP LSST effort are to probe the nature of dark matter and dark energy through weak gravitational lensing and galaxy clustering. The weak gravitational lensing signal is very subtle and easily swamped by systematic effects from the atmosphere and detector, as well as imperfections in the analysis algorithms. To quantify/mitigate these systematic effects, the LSST project is undertaking a very detailed program to simulate all aspects of this measurement. The estimated compute cost for this process is $\sim 10^7$ compute hours and 100 TB of storage. The other dominant cost for the LSST analysis are the simulations necessary for quantifying the uncertainties in the measurements. The LSST DESC collaboration estimates requiring ~ 2500 simulations, each with a cost of $\sim 1\text{M}$ CPU hours, for a total of $O(10^9)$ CPU hours and 100 TB of storage. The other analysis tasks are expected to be subdominant to this.




Figure 1: Exponential growth over the past and coming 20 years of CMB data volumes gathered by ground-based, balloon-borne and satellite missions, and supercomputer peak performance using the NERSC flagship system as a proxy.

The DESI survey aims to obtain redshifts to ~ 25 M galaxies over $\sim 14,000$ sq. deg. of the sky. This represents an order of magnitude increase over current galaxy surveys. The estimated computational and storage requirements for the survey are ~ 100 M CPU hours and $O(1)$ PB of data storage. As with the other projects described here, the other dominant cost will be the simulation requirements for error quantification (although these should be less demanding than the LSST case).

The CMB-S4 experiments aim, amongst other goals, at constraining both the CMB polarization caused by gravitational waves from the inflationary epoch as well as constraining the total neutrino mass at unprecedented precision. These will provide astrophysical windows into scales impossible (or very hard) to reach by traditional particle-physics experiments. Reaching these goals will require a 1000-fold increase in the data volume compared to the Planck satellite, and will require improvements in analysis algorithms as well as their implementations to make best use of the architectures of the next generation of supercomputers. We can estimate both storage and compute needs by scaling from experience with the Planck satellite. The Planck satellite used 10^8 CPU-hours for its analysis, a CMB-S4 experiment will require 10^{11} CPU hours. Planck used $O(250)$ TB in 2014, but this constrained us to storing $O(10^6)$ maps instead of the full $O(10^7)$; CMB-S4 will need $O(10)$ PB in 2020.

An important set of considerations for these next generations of surveys are the large and geographically diverse collaborations and serving different views of the data to them. Furthermore, the traditional data models and file systems do not scale well to these data sets, since they often result in $O(10^9)$ files making accesses/queries time consuming. The collaborations will therefore benefit from tooling to simplify these tasks. Furthermore, different analysis tasks often have very different scales of parallelism, making them hard to chain together under traditional MPI/OpenMP models. Finally, issues of openness and reproducibility are becoming very important considerations for these collaborations.

Data volumes and processing needs for the G2 direct detection experiments, while larger than current experiments, are expected to be small. LUX/LZ expects 1.2 PB of data/year, with 250 cores needed for prompt processing and 3k cores needed to reprocess one year of data in three months. Including Monte Carlo simulations the LUX/LZ plan includes support for 10 PB of disk storage and 10k cores during operations.

SuperCDMS, which is concentrating on lower WIMP masses, will have data volumes of the order of 100 TB/year, with a very small number of cores needed for prompt processing. The plan is to adopt a very distributed computing model to be able to access as much CPU as possible at

collaboration sites to minimize the time needed for data reprocessing needs and Monte Carlo production. The overall CPU needs will be similar to LUX/LZ.

C.7.3 Compute, Data, and Services Needs on the 2020/2025 Timescale

The previous section discusses the needs of the three major cosmic frontier surveys individually. We summarize these below:

- Compute: $O(10^{11})$ CPU hours. This is dominated by the anticipated requirements of the CMB-S4 experiment, with the other two experiments, estimating about two orders of magnitude less computation. Therefore, this should be able to accommodate an expanded compute requirement from these surveys.
- Data: $O(10)$ PB. This is again dominated by CMB-S4, with the other surveys requiring an order of magnitude less data.
- Serving the raw and reduced data to large, geographically diverse collaborations.
- The ability to support a broad range of analysis/programming paradigms including (but not limited to) traditional HPC usage, massively (but trivially) parallel analysis tasks, to exploratory analyses.
- Developing the infrastructure to enable reproducibility/openness.

[HEP C.8.2, Viren and Schram]

C.8.2 Evolution of Science Activities and Computational Approaches on the 2020/2025 Timescale

Neutrino Experiments: The next generation Deep Underground Neutrino Experiment (DUNE) employs a 40kton liquid argon time-proportional chamber (LArTPC) far detector with 1.5M channels reading wires spaced every 5mm and acquiring waveforms at 2MHz for about 5ms. This produces approximately 10-20 GB per “event”. In order to be sensitive to supernova bursts the readout must be capable to sustaining 10s of seconds of data collection. Such full-stream readout can produce 100s of Exabyte per year. However, most of the ADC samples will be noise and can be discarded by using “zero suppression” technique in which low-threshold portions of the waveform are discarded. This can reduce the raw data rates to the TB/year.

The DUNE LArTPC is incredibly fine-grained compared to other neutrino detectors (except bubble chambers). Traditional reconstruction techniques will, at best, scale linearly with the number of channels and may scale as worse as N^3 and novel techniques exploiting the unique characteristics of LArTPC are envisioned. At the very least, it is expected that production processing must exploit Grid resources.

Belle II Experiment: The Belle II distributed computing system must handle ~ 85 PB data volume for each year when the SuperKEKB accelerator is operating at design luminosity. The Belle II computing model includes several tasks such as raw data processing, Monte Carlo event production, physics analysis, and data archiving. Belle II has adopted the DIRAC framework for their Grid system. DIRAC provides both a workload and data management system along with other systems, such as data creation/manipulation workflow system, metadata catalog system, etc.

The Belle II software/computing workflow has numerous elements. At the core is the standalone Belle II software framework (BASF2) and external dependencies (Geant4, ROOT, etc.). The code is currently distributed on the grid using CVMFS servers. Grid sites are deployed with commodity machines (x86 processors that run Linux) and require queuing software (HTCondor, SLURM, etc.) and Grid middleware (gridftp, voms, etc.). Currently Belle II jobs are submitted to the Grid as single core jobs, however, Belle II is developing/testing a multicore solution. Multicore jobs will

reduce processing time and RAM per core. This will allow Belle II jobs to more efficiently use opportunistic resources such as Amazon EC2 and HPC clusters. Belle II has started to test jobs on HPC resources and identified some challenges when submitting jobs as backfill; these challenges are expected to be partially resolved with multicore jobs. However, most of the Belle II code and external library do not take advantage of the HPC hardware and are not compiled optimally. Moreover, only one binary is currently distributed on the Grid.

The Wide Area Network requirements for Belle II are similar to that of the LHC experiments and the needs are expected to be satisfied by the NRENs.

Belle II is currently developing extensions to the DIRAC framework to enhance the distributed computing model. Central to this effort are the “Fabrication System” and “Data Management System.”

[HEP D.1.3, Ge et al.]

D.1.3 Current and Future HPC Needs

Computational Hours: Currently ACE3P uses 2.5M CPU hours on NERSC computers.

A computational challenging problem will be to model dark current effects in the entire linac of an accelerator such as the superconducting linac in PIP-II and its upgrade to high power operation. The problem size will be 20-30 times larger than that of current simulation. It is anticipated that a growth of more than an order of magnitude to 50M CPU hours is required in the next decade.

Parallelism: The codes use MPI for parallelism with the average number of cores in the order of 5,000. Parallelism on multi-core nodes focuses on the development of hybrid linear solvers that are scalable in memory. The average number of cores will increase by an order of magnitude to 50,000 due to the increase in problem size.

Memory: ACE3P simulation in frequency domain requires large per core memory and hence benefits from compute nodes with large memory. Currently, the simulation uses 64 GB of memory on a compute node and the aggregate memory can reach up to 2 TB for electromagnetic simulation. Future aggregate memory will increase by an order of magnitude to 40 TB for multi-physics simulation.

Scratch Data and I/O: A typical run in the time domain generates 1-2 TB of data including field and particle snapshots and checkpoint files. A total of 50 TB scratch space is required for ACE3P users to perform their simulations concurrently. The current I/O bandwidth is estimated to be 20 GB/sec. Future requirements will increase the size of output datasets to 20 TB and the I/O bandwidth to 80 GB/sec to maintain reasonable I/O percentage of the runtime, which is about 20%.

Long-term and Shared Online Data: Several of the production runs are shared for the collaboration. It is estimated 5 TB storage for long-term data is required, which will increase to 50 TB in the next decades.

Archival Data Storage: About 50 production runs need to be archived. The estimated current space is 100 TB and the future storage will increase to 800 TB.

Workflows: The data generated from simulations on LCFs are transferred back to local computing resources for analysis, and hence maintaining and enhancing adequate data bandwidth from the remote facility are essential to the scientific process. For the next decade, the use of remote processing and visualization of data will alleviate the demand for high bandwidth of data transfer.

Many-Core and/or GPU Readiness: ACE3P’s current parallel implementation uses MPI. The plan to build a hybrid programming paradigm with OpenMP is under way, for example, for particle

tracking. In addition, ACE3P will benefit from improvement of third-party linear algebra libraries on multi-core architectures.

Software Applications, Libraries, and Tools: N/A

HPC Services: N/A

Additional Needs: N/A

[HEP D.3.2, Finkel et al.]

D.3.2 Computational and Data Strategies

Approach: HACC’s N-body methods use tracer particles to track the phase-space distribution of matter in an expanding universe. Force-splitting methods are used to calculate long-range forces using MPI distributed-memory methods and short-range forces using shared-memory methods. Individual simulations must run longer than individual job runtimes (and mean-time-to-failure) requiring high performance IO for checkpoint/restart files. Particle methods should scale well for the next generations of large-scale HPC systems.

Codes and Algorithms: The HACC framework runs on all types of HPC systems, including power-efficient multi/many-core systems and accelerated systems. Long-range forces are computed by a high-order spectral particle-mesh method. Portable, custom-written MPI code is used to perform a distributed memory FFT to solve for long-range forces. Shortrange force calculations are adapted to the underlying architecture. Tree methods are used to generate particle interaction lists on multi-core and many-core CPUs, while an OpenCL code employs a simpler fixed spatial-scale data structure for calculations on GPUs. Customwritten (non-collective) MPI-IO code with internal checksums is used to checkpoint/restart and output files. We will add a new higher-order smoothed-particle hydrodynamics (CRKSPH) method to HACC in order to scale baryonic physics calculations to future HPC systems.

[HEP D.4.3, Gnedin]

D.4.3 Current and Future HPC Needs

Computational Hours: Currently we use between 50 and 100 million hours per year. “Ideal” reionization simulations mentioned above will require of the order of 500 million hours each, and a statistical ensemble of at least 5 of them will be required, so we are talking about 2.5-3 billion CPU hours for the 2020-2025 period.

Parallelism: The current implementation of the ART code scales to about 10,000 nodes and 20 cores/node for the largest currently feasible simulation sizes. Going beyond this scaling will not be possible with the current implementation both for numerical and algorithmic reasons. We are currently starting work on the next generation of the ART code that should scale on exascale platforms (i.e., reach greater than million core scaling), with the expectation that the new code will be production ready around 2020.

Memory: All of the current and future simulations are CPU limited and the memory requirements are not critical – i.e., they are highly suitable for the future machines with low memory-to-peak-performance ratio. The total memory requirements for the “ideal” simulation described above will be of the order of 1 PB. The per-node memory requirement will depend on the particular implementation of the next version of the code, but in no case should be below 16 GB.

Scratch Data and I/O: Because of the persistent value of these simulations, a sensible number of snapshots will need to be stored. The exact storage requirement will depend on the degree of

compression available to us; something in the range of 10-30 PB seems to be reasonable. The IO bandwidth will be crucial, however, and will need to exceed the IO performance of the BlueGene/Q by at least a factor of 100.

Long-term and Shared Online Data: At present we need about 300TB of active data storage. In the period 2020-2025 this requirement will grow by a factor of 10-20.

Archival Data Storage: At present we have about 1PB of archival storage used for simulation outputs. In the period 2020-2025 this requirement will grow by a factor of 10-20.










Workflows: We don't need any workflows and do not plan to use any in the future. Many-Core and/or GPU Readiness: The current version of the ART code is not ready for exascale and inhomogeneous architectures. As we discussed above, work on the next generation, exascale AMR code has already started, and the new code is expected to be operational by about 2020.

Software Applications, Libraries, and Tools: N/A

HPC Services: N/A

Additional Needs: N/A

This page is intentionally left blank.



APPENDIX K

Sections Cited in This Report
from the Exascale Requirements
Review for Nuclear Physics (NP)

This page is intentionally left blank.

SECTIONS CITED IN THIS REPORT FROM THE EXASCALE REQUIREMENTS REVIEW FOR NUCLEAR PHYSICS (NP)

The following sections from the Nuclear Physics (NP) Exascale Requirements Review report are cited at least once in this Crosscut report.

[NP Abstract], [NP ES.1], and [NP ES.3.2]

Appendix E contains the NP Executive Summary.

[NP 3.1]

3.1 Nuclear Astrophysics

3.1.1 Scientific Challenges and Opportunities

This is a golden era for nuclear astrophysics, a subject that offers tantalizing new opportunities for both nuclear physics and HPC, with the promise of synergistic development of both disciplines. Nuclear astrophysics stands at the intersection of major experimental and theoretical efforts in the areas of nuclear structure and reactions (e.g., FRIB [FRIB undated; FRIB 2012]) and fundamental symmetries (e.g., $0\nu\beta\beta$ decay), and rapidly developing capabilities in astronomical observation — from new ground- and space-based observatories covering the entire electromagnetic band to the advent of gravitational wave astronomy (e.g., LIGO). It is clear then that nuclear astrophysics theory efforts are fundamentally multi-physics in nature, and that large-scale numerical simulations are at the heart of this science. In fact, HPC capabilities are opening up new ways to explore outstanding issues in nuclear physics.

Key objectives in the recent 2015 Long Range Plan for Nuclear Science (DOE and NSF 2015) include exploring the origin of the elements (e.g., ascertaining the site[s] where the r-process elements such as uranium, gold, and iodine are synthesized); the physics of ultra-dense neutron-rich nuclear matter; and the nature of neutrinos and their interactions in nuclei and dense matter. Understanding nucleosynthesis means understanding how compact objects work — from core-collapse supernovae (CCSNe), to Type Ia thermonuclear supernovae, to binary neutron star (NSNS) mergers, to explosive phenomena such as x-ray bursts (XRBs). One immediate objective is to leverage the considerable success we have had in modeling neutrino transport and high-temperature (up to tens of million electron volts [MeV]) dense-matter physics in collapsing and exploding compact objects to explore, at a comparable level of fidelity, new venues (e.g., NSNS mergers). Simulations of these mergers could allow LIGO data to provide insight into ultra-dense nuclear EoS physics and determine whether these events are significant contributors to the inventory of r-process heavy nuclei in the Galaxy.

Many simulations in these areas will share some common characteristics: they will be multi-physics and likely multi-dimensional, and they will share a common strategy in which extracting science from the computations will require extensive parameter sensitivity studies. In some cases, unique and challenging new kinds of physics will be modeled — for example, neutrino flavor evolution with neutrino scattering (neutrino flavor/spin quantum kinetics). Progress has already been made in this area, and future efforts will involve a “layered” approach, exploring the efficacy of different approximations. Another example is modeling the electromagnetic signals that might accompany NSNS mergers or the r-process nucleosynthesis of CCSNe, a key part of “multi-messenger” probes of these events.

The multi-physics nature of the HPC techniques employed in this field nicely aligns with the training of adaptable nuclear and computational physicists. Indeed, it is increasingly common to find universities and national laboratories creating joint physics/computing efforts. These efforts attract outstanding graduate students, postdoctoral researchers, and faculty/staff.

[NP 3.1.4]

3.1.4 Computing, Data, and Networking Needs and Requirements

Effective use of exascale systems for nuclear astrophysics simulation will involve requirements spanning the ecosystem — from the wide-area network connections provided by ESNet to the individual processing cores on each node. Many of these requirements are driven by the need to perform ensembles of simulations to achieve important scientific results. Dozens to hundreds of individual simulations of the stellar explosive events described in Section 3.1.2 (Priority Research Directions) will be required. The intent is to fully explore the space of possible progenitors, initial system conditions, and necessary parameterizations in each scenario. In addition, the impact of new or newly implemented (and more complete) physics at each epoch along the road to exascale will require a re-evaluation of earlier work.

Because much of nuclear astrophysical simulation is multiphysics in nature, possibilities for collaboration with ASCR-supported researchers in algorithm development for multiphysics abound. Applied mathematicians and computer scientists with interests in, for example, methods to move away from operator splitting for coupling physical models, would be excellent candidates for collaboration. Work in other areas — ranging from dense and sparse linear solvers to high-order techniques for transport to advanced AMR techniques — is also needed.

In this context, important features of the exascale ecosystem required for nuclear astrophysics calculations include the following: wide-area networking, disk and archival resources, burst buffers, internode networks, memory hierarchy, and processor and programming models.

The number of reduced-data products that could be useful to a broad audience and could be made publicly available from nuclear astrophysics simulations is relatively small compared with, for instance, the number from large-scale structure simulation or from terrestrial plasma physics simulations. We anticipate that total data volumes will not exceed several terabytes (TBs) by 2020 and perhaps 10–50 TB by 2025. These data products include simulated spectra from reactive and radiation hydrodynamics simulations, as well as multi-messenger signals, such as neutrino signatures and gravitational wave templates. Web-based interfaces to access, visualize, and interact with stored data would be useful for science analyses and for disseminating the spectral output to the community. Given the limited number of full simulations that will be available, it would be helpful to field statistical tools (e.g., emulators based on Gaussian processes) that interpolate over the populated parameter space and find best fits to observational data. However, most of the investigators in the field perform related simulations on a multitude of platforms, both ASCR facilities and NSF machines (and, in some cases, on NASA resources as well). Because the size of individual checkpoint files can grow to be large (see below), being able to efficiently transfer these files over the wide-area network becomes a challenge.

Data storage capacity is a marked need. Sheer capacity is a necessary, but not sufficient, condition to enable science from nuclear astrophysics simulations. The spatial resolution requirements and multiphysics nature of nuclear astrophysics simulation lead directly to this need for considerable storage. For example, the largest, highest-fidelity CCSNe simulations today produce roughly 500 TB per second of physical evolution. Increasing local spatial resolutions (via AMR) by two to three orders of magnitude, coupled with an increase of 10–100× in the number of degrees of freedom at each grid point (e.g., from fully angle-dependent neutrino transport and/or from increasing the number of nuclear species evolved) will lead to data volumes of roughly

100 petabytes (PB) per simulation by 2025. In addition to the immediate requirement for capacity for long-running simulations, these data need to have relatively long persistence on resources that can enable interactive data analysis; for example, sufficient read bandwidth and low latencies for read. The timescale associated with the peer-reviewed journal publication cycle—including addressing the comments and questions of referees — is the minimum persistence time required. Although increased physical fidelity will require significant additional computational load, it should be noted that although data volumes are expected to increase significantly in the exascale environment, the rate at which data will be pushed to disk will not increase as quickly. Therefore, capacity requirements far outpace I/O bandwidth requirements for the multiphysics simulations that characterize nuclear astrophysics.

Nuclear astrophysics simulation codes already universally make extensive use of checkpoint/restart. Indeed, considerable code engineering has been performed in many cases to ensure that maximum bandwidth/minimum write times are achieved. Ideally, the addition of burst buffer technology in future platforms would present use methods that would allow this significant investment in code to persist. Although small additions or changes to application programming interfaces (APIs) would present few problems, these changes would hopefully be portable and are necessary to achieve good performance. An alternate burst buffer use case is using non-volatile random-access memory (NVRAM) as out-of-core storage for large, multi-dimensional data sets during runtime. Typical uses of such an out-of-core capability would be storage and read of tabular physics, such as neutrino and photon interaction cross-sections and thermonuclear reaction rates. These data sets would be write-once (at runtime), read-many in nature, and large enough to exceed single-node memories.

More bandwidth is also needed between the new, more powerful nodes that will become available. The combination of powerful nodes and high-fidelity multiphysics will lead to large potential communication payloads, perhaps as much an order of magnitude larger than the nearest-neighbor communications performed today. Latency is also an important consideration, but the requirements for global reductions are less stringent and generally do not grow with increased physical fidelity (e.g., global timestep determination scales with physical resolution, at worst).

Deep memory hierarchies will present significant challenges for multiphysics nuclear astrophysics simulation. Because large, multidimensional data structures are ubiquitous in astrophysics implementations, there are issues with both memory layout — how data structures are arranged in memory to ensure that memory can be accessed efficiently (e.g. for vectorization) — and memory placement. Memory placement issues include placing data structures in particular tiers of the memory hierarchy to allow the best locality (e.g., in GPU memory if the operations are to be performed by the GPU). Most importantly, the nuclear physics community needs a portable way to efficiently handle both memory layout and memory placement. Libraries are likely the best method because they do not require direct compiler support (helping with ubiquity) and can be “carried around” with individual codes, if need be. Completely ceding control of memory placement and layout to the runtime system is not considered viable or desirable because layout and placement needs are physically motivated and can be time dependent.

To most effectively use the various processing elements that will be available at the exascale and in the era leading up to it, a more-or-less portable programming model for node-level parallel processing is desirable. Several groups have already adopted a more descriptive approach to this problem (e.g., through the use of OpenACC), but more prescriptive approaches are also used. Whatever the nature of the exascale programming model, reasonable default assumptions about data layout, data placement, execution, and dependencies will be the distinguishing characteristics of a successful model. Wide availability of compilers and runtime systems that implement this model will be key to its ultimate adoption by the community. Significant amounts of development on smaller resources — from laptops to local clusters — will be required to produce efficient codes

for the exascale. Students and postdocs at universities will undertake much of this development, tying the need for wide availability directly to workforce development.

Most current nuclear astrophysics codes are written primarily in Fortran, and the continued support of this language is important for most groups. Ultimately, it is the robust support of arrays in Fortran that make it the preferred language for most of the codes in the field, regardless of the historical reasons for its use. Nevertheless, although C++ is becoming more and more popular as a “framework” language for nuclear astrophysics, there are decades of development encapsulated in hundreds of thousands of lines of Fortran in the field.

Figure 3-8 summarizes the computing resources required to accomplish the planned nuclear astrophysics program. In all computational resource dimensions, nuclear astrophysics has significant exascale computing requirements.

[NP 3.2] and [NP 3.2.1]

3.2 Experiment and Data

3.2.1 Scientific Challenges and Opportunities

The scientific challenges and opportunities associated with nuclear physics Experiment and Data span the full science program of DOE’s NP program. These challenges and opportunities were well articulated in 2015 Long Range Plan for Nuclear Science (DOE and NSF 2015) and include the following:

- Unfolding the quark and gluon structure of hadrons and nuclei;
- Realizing a predictive model of nuclei and their role in the cosmos;
- Testing the particle–antiparticle nature of neutrinos and other fundamental symmetries and neutrino research that open new doors to physics beyond the Standard Model (BSM); and
- Studying the properties and phases of quark and gluon matter in the high temperatures of the early universe, and exploring the spin structure of the proton.

Rather than pushing the energy frontier, nuclear physics experiments are driven by precision to access the multi-dimensional and multi-channel problem space, requiring beam intensity, polarization, and exquisite control of backgrounds and systematics. Over the next decade, this science will be addressed at the current and planned experimental facilities in the nuclear physics portfolio. NP supports three national user facilities:

1. The Argonne Tandem Linac Accelerator System (ATLAS) accelerator at Argonne National Laboratory (Argonne) enables experiments in low-energy nuclear structure and astrophysics.
2. Science at the Continuous Electron Beam Accelerator Facility (CEBAF) at JLab focuses on nucleon structure and QCD.
3. The RHIC at BNL is used to study the properties of the QGP and the spin structure on the nucleon.

The nuclear physics community also uses the Fundamental Neutron Physics Beam Line at Oak Ridge National Laboratory (ORNL) to perform precision studies of the neutron and its interactions. FRIB, under construction at MSU, will produce the short-lived isotopes needed to explore nuclear structure, nuclear astrophysics, and fundamental symmetries. NP is planning to deploy a ton-scale instrument to explore $0\nu\beta\beta$ decay, providing a sensitive search for physics BSM that leverages current efforts such as the Cryogenic Underground Observatory for Rare Events (CUORE) experiment depicted in Figure 3-9. Beyond the next decade, there has been an expressed need from the community for construction of a new powerful EIC to precisely elucidate the role of gluons in nucleons and nuclei.

The successful realization of science outcomes at these facilities critically depends on the optimization of experimental infrastructure. HPC plays an important role in the reliable and efficient execution of the experiment and in the rapid and well-grounded treatment of the resulting data.

The particle accelerators at the heart of the nuclear physics user facilities rely on computationally intensive simulations to optimize accelerator performance. These include simulations of the beam transport through the machine, including the self-field interaction within the beam itself and the beam–beam interaction from the opposite colliding beam. Design of the future EIC will offer significant computational challenges. There is a near-term need to assess competing accelerator designs for technical pitfalls and feasibility, as well as a longer term need for a full-fledged design study to optimize cost/performance. Both these needs can only be met with HPC, because the selffield effects and the electron cooling that need to be considered for the EIC design are based on self-consistent particle simulations that demand exascale computing resources.

Reliable and safe accelerator operations performance has elements that benefit from HPC. The overhead associated with the tuning of a complex particle accelerator can be significantly reduced, and the performance of the accelerator can be significantly improved, by employing machine learning and real-time multi-physics beam transport simulations, both of which are computationally intensive. The high-power particle beams can produce prompt radiation, and verification of the shielding for personnel protection makes use of radiation transport codes that demand significant computational resources to achieve results that have the necessary statistical significance. The damage caused by the interaction of high-power beams on targets and beam dumps can severely limit facility operations. Radiation transport and commercial engineering codes are both used to simulate material properties under such severe conditions, where again significant computational resources are needed for the requisite statistics.

Highly efficient detector systems are necessary to achieve the science goals outlined above, especially when the underlying experiments aim to detect rare events. Detector simulations are performed by using standard codes such as Geant, Fluka, and virtual Monte Carlo (VMC). As detector complexity has increased, the compute needs for the simulations have grown correspondingly. Of course, the detector simulation itself is just one part of the work flow for simulated and observed event reconstruction. Event generators follow different paradigms depending on the physics models involved. Their computational needs vary, but the demand for statistically significant results typically translates into computational requirements that comprise more than half of the total computational resources needed for experimental data analysis.

The approach to the acquisition of data from the highly efficient detector systems has not changed much in the last 20 years. Detector signals are digitized by electronics in front-end crates, and data readout is initiated by the trigger electronics. Data are transported to an event builder, and the subsequent built events are distributed for filtering, monitoring, display, etc. The event stream is also stored to disk for offline analysis. The trigger defines the data; therefore, experiment optimization depends critically on minimizing the time needed to validate data. Data throughput is determined by the physics event rate, detector response, and overhead associated with the data work-flow steps enumerated above. The latter depends on the performance of both the network and compute environment. A current trend is to move some functionality previously performed in software running on embedded processors into firmware on custom electronics. The ever-increasing expense and decreasing availability of beam time at user facilities drive such efforts. Loosening triggers to store as much data as possible for future data-mining, without putting the primary science goal of the experiment at risk, is becoming more commonplace.

In terms of data analysis and data-mining, there are a plethora of ways that data are treated by the different experiments. NP-supported experimentalists generally use common software building blocks (e.g., Geant, ROOT) for modeling and analysis of data. In general, HPC methods and advanced computing facilities are not being utilized to process the data. The compute infrastructure for data analysis and data storage/archiving usually resides within the host facility. The offline approach to data analysis is usually different from that performed online, which as stated earlier is “limited” by the need for rapid data validation. However, there is already motion in the field to employ higher performance online computing environments to narrow this gap between online and offline analyses. The nuclear physics experiment community should continue to explore other opportunities to leverage HPC resources to optimize data analysis, while still being able to employ established tool sets.

[NP 3.2.3.2, Item 3]

3.2.3.2 ASCR

Emerging Needs in Data Science

There are emerging scientific questions that will require advanced machine learning techniques, possibly over combined data sets from multiple experiments or observations. In NP, an example of this is 3-D nucleon tomography, an emerging research area that will begin to acquire significant data samples with the JLab 12-giga electron volt (GeV) program (see Figure 3-13) and other experiments coming online, as well as integrating older data. The core of this program will be global fits of worldwide data to theories with five or more parameters. Collecting and preparing the data for analysis and maintaining them for use over several decades will be required, as will provisioning for work flows that capture the parameters of the fits, the version of the theory codes, and any phenomenological preparation of the experimental data. Other disciplines will face similar challenges as theoretical and experimental advances open novel lines of inquiry. Appropriately capturing the data is one of the more pressing issues, because knowledge can be irretrievably lost. The NP community understands that neither ASCR nor NSF is currently positioned to provide large data repositories for the individual program offices. These program offices are likely to have to integrate and federate across their facilities’ storage and/or integrate with the community repositories to meet the long-term scientific goals. There is a need to establish best practices and reference architectures (similar to the ESNet DMZ architecture) as a way to guide the creation and maintenance of cost-effective and usable data repositories.

[NP 3.3]

3.3 Nuclear Structure and Reactions

3.3.1 Scientific Challenges and Opportunities

The field of nuclear structure and reactions (NS&R) research is now undergoing a renaissance, as a result of experimental developments and exciting theoretical progress in the low-energy nuclear many-body problem. The general challenge for this interdisciplinary field is to understand the principles of building up nuclear complexity out of fundamental degrees of freedom, which — when inspected at higher resolution — have a complicated structure of their own. The overarching questions that drive NS&R are as follows:

- Where do nuclei and elements originate?
- How are nuclei organized?
- How can nuclei be exploited to reveal the fundamental symmetries of nature?
- What are the practical and scientific uses of nuclei?

Complete answers to these questions require a much deeper understanding of atomic nuclei than is currently available. Both low-energy experimentation and nuclear theory address these questions in a synergic manner.

The planned NS&R research program at the exascale is perfectly aligned with the national nuclear physics program described in the 2015 Long Range Plan for Nuclear Science (DOE and NSF 2015) and the NRC decadal study (National Academies Press 2013) reports. A new generation of large-scale experiments and facilities using exotic nuclei to probe nuclear and particle physics is being constructed in the United States. FRIB is designed to study the structure of nuclei important for the processes that create all the elements in the universe. Nuclear theory at the exascale will calculate the properties and decays of atomic nuclei to understand the mechanism of nuclear binding and the formation of the elements and to extend these results to astrophysical environments that cannot be created in the laboratory. It will also address nuclear physics programs at JLab, which will delineate nuclear properties at short inter-nucleon distances and provide data on neutron distributions in nuclei.

$0\nu\beta\beta$ -decay experiments Majorana and nEXO, using ton-scale germanium and xenon detectors, are a high priority in the 2015 Long Range Plan for Nuclear Science (DOE and NSF 2015; National Science Advisory Committee [NSAC] 2014[a]) and are being designed to (1) determine whether neutrinos are their own anti-particle (Majorana neutrinos) and (2) probe the neutrino's absolute mass scale in the laboratory. The envisioned exascale NS&R program will calculate and quantify uncertainties in nuclear $0\nu\beta\beta$ -decay matrix elements to guide next-generation $0\nu\beta\beta$ -decay experiments.

The Deep Underground Neutrino Experiment (DUNE) is being designed to accurately measure neutrino oscillations and the neutrino mass hierarchy (Fermilab undated[a], [b]). Here, nuclear calculations will describe neutrino and anti-neutrino scattering from argon to provide a more accurate determination of neutrino oscillation parameters, including the parity and CP violating phase.

Maximum success for these several hundred million- to billion-dollar experimental investments requires substantial simultaneous investment in computations of the structure and reactions of nuclei. As multiple DOE and NRC reports recognize, exascale computations are required to fulfill the critical role of guiding and making predictions for new experiments, and important computational advances are required to solve challenges expected at the exascale.

Figure 3-15 shows the degrees of freedom of NS&R research in the context of the theoretical roadmap of the nuclear many-body problem. QCD — the underlying theory of strong interactions — governs the dynamics and properties of quarks and gluons that form baryons and mesons; hence, QCD is also responsible for the complex inter-nucleon forces that bind nuclei. In this area, significant progress is being made by computing properties of the lightest nuclei and nuclear interactions within lattice QCD. A powerful suite of ab initio approaches based on inter-nucleon interactions provides a quantitative description of light- and medium-mass nuclei and their reactions. For medium-mass systems, global configuration-interaction (CI) methods that employ microscopic effective interactions offer detailed descriptions of nuclear excitations and decays. For heavy complex nuclei, the tool of choice is the nuclear DFT and its extensions.

The NS&R exascale program will use different methods to perform validation and verification on nuclei in the regions of models' overlap to ensure accurate coverage of the entire nuclear landscape. The combination of these methods provides a coherent picture of the structure and reactions of all atomic nuclei over a wide range of energies and momenta. Each method has its major advantages: they can first be compared and validated in specific problems, and then each can be employed in the most appropriate regime. Although the methods described here are similar to those used for other

strongly correlated quantum many-body systems, they are unique in that nuclear codes must solve for interactions that strongly couple different spin (up and down) and isospin (neutron and proton) states. NS&R computations involve extreme-scale codes that face common issues in efficient and reliable task-based parallelism, hierarchical computation and associated memory usage, and important issues in optimization at the largest scales. The community has successfully met such challenges over the past several generations of large-scale computing facilities through the SciDAC (DOE 2015) and INCITE programs (DOE Leadership Computing 2016).

What can modern nuclear structure theory do to improve its efficiency, enabling it to respond nimbly to opportunities for scientific discovery? One mechanism is to team up in large, multi-institutional efforts involving strong coupling among nuclear physics theory, computer science, and applied mathematics. While experimentalists are no strangers to large collaborative efforts, nuclear theorists have eventually come to the realization that “the whole is greater than the sum of its parts.”

In this context, one representative example is the NUCLEI SciDAC collaboration (Ohio State University undated) involving nuclear structure theorists, computer scientists, and mathematicians from 16 institutions, including 10 universities and 6 national laboratories. The scope of nuclear science and math/computer science is quite wide-ranging. Over the course of NUCLEI and its Universal Nuclear Energy Density Functional (UNEDF) SciDAC predecessor, collaborations across domains have grown and now involve many direct connections. Figure 3-16 illustrates the present status of collaborations within NUCLEI. As is apparent from this network diagram, the math/ computer science participants are directly embedded in the various nuclear theory efforts. In each partnership, these participants collaborate with physicists to remove barriers to progress on the computational/algorithmic physics side. This partnership has proven to be very successful, resulting in many excellent outcomes, some initially unanticipated. The UNEDF/NUCLEI case emphasizes the importance of assembling agile theory teams working on important questions and programmatic deliverables that would be difficult or impossible to tackle by individual investigators or atomized small groups.

As in other areas of science, NS&R research uses a cycle of “observation-theory-prediction-experiment” to investigate phenomena, build knowledge, and define future investigations. Such an approach guides the relationship between theory and experiment: theory is modified or rejected based on new experimental data and the improved theory can be used to make predictions that guide future measurements. The positive feedback in the experiment-theory cycle, illustrated in a schematic way in Figure 3-17, can be enhanced if statistical methods and computational methodologies are applied to determine the uncertainties of model parameters and calculated observables. In partnership with applied mathematics and computer science, modern nuclear structure theory strives to estimate errors on predictions and assess extrapolations. This is essential for developing predictive capability, as theoretical models are often applied to entirely new nuclear systems and conditions that are not accessible to experiment. Statistical tools can be used to both improve and eliminate a model or better define the range of a model’s validity.

[NP 3.4]

3.4 Cold Quantum Chromodynamics

3.4.1 Scientific Challenges and Opportunities

A precise understanding of the proton, the neutron, and the forces between them is critical to many aspects of subatomic science — from the earliest moments of our universe and the fundamental conditions necessary for life, to national security, the production of carbon-free energy, and the transmutation of nuclear waste. In each of these areas, and beyond, there are clear needs to improve our present understanding of, and to enhance our predictive capabilities in, nuclear systems — from the very smallest nucleus and its reactions to nuclear matter found in the dense interiors of CCSNe.

In designing the next generation of HEP proton colliders, the particle production rates and their discovery potential depend sensitively on the structure of the proton. The present uncertainties in the structure of the proton translate into uncertainties in machine and detector design parameters and projected experimental costs.

In designing experiments to search for the violation of lepton number (L) — that may manifest itself, for instance, in the $0\nu\beta\beta$ decay of nuclei — uncertainties in the expected nuclear decay rates currently result in significant uncertainties in estimates of the amount of the active material required to accomplish the experimental objectives. These uncertainties have implications for precision atomic physics. The electronic structure of atoms is sensitive to the distribution of charge within the nucleus at its center, and recent experimental results have demonstrated that the radius of the proton is not as well determined as scientists had once believed. The approximately 7σ discrepancy between electronic and muonic measurements is under active investigation (see, for example, Figure 3-25), and the computational technology now exists to calculate the proton radius from first principles, complementing the experimental program. The present-day uncertainties in nuclear reactions and the behavior of matter under extreme conditions continue to be addressed by ever-increasingly precise determinations of the spectra of hadrons, their structure, and their interactions. Quantities that are the foci of investigations in cold QCD — the structure and reactions of mesons, nucleons, and light nuclei — have an impact on an array of systems, from uncertainties in fusion reactions that play an important role in the operation of the NIF (Lawrence Livermore National Laboratory undated); to uncertainties in isotope and neutrino production in nuclear fission reactors; to neutrino interaction rates in experiments, such as DUNE (Fermilab 2016), designed to determine the fundamental properties of neutrinos; to the behavior of matter in explosive astrophysical environments; to expected features of the gravitational waveforms from the inspiral of binary NSs that are expected to be seen using LIGO (Figure 3-26). Results in light nuclei, combined with efforts of many-body theorists, are expected to have a direct impact on the studies of many-nucleon systems, including those studied at the FRIB and through astrophysical observations and those used in accelerator neutrino and $0\nu\beta\beta$ -decay experiments. The anticipated advances in high-performance computing and data ecosystems (HPCDEs) that are expected in the exascale era are required to provide precise calculations of key quantities and systems necessary to accomplish the scientific objectives of the cold QCD program.

The drive toward a precise theoretical understanding of cold QCD systems has traditionally spurred the development of HPC. It was the desire to perform increasingly more realistic QCD calculations in HEP that led to the development and fabrication of the QCD on digital signal processors (QCDSP) then QCD on a chip (QCDOC) by the Lattice QCD Group at Columbia University, in collaboration with the University of Edinburgh and BNL, and ultimately to IBM's Blue Gene/L supercomputer. This trend continued with the deployment of large, GPU-based supercomputers, such as the one deployed by the USQCD consortium at JLab in 2010 and dedicated to lattice QCD calculations.

Protons, neutrons, nuclei, and all strongly interacting particles (hadrons) emerge from the dynamics and interactions of quarks and gluons, fundamental building blocks of our universe, dictated by QCD. Hadrons are composite systems comprising “valence” quarks, immersed in a “sea” of quarks, antiquarks, and gluons that are quantum mechanically fluctuating in and out of the vacuum. They are entangled quantum systems with an indefinite particle number, and the intrinsically nonlinear and quantum mechanical nature of QCD confines the quarks and gluons into hadrons so that they have never been observed in isolation. In analogy with electromagnetism, gluons carry the force between quarks in a way that is similar to photons carrying the force between electrically charged particles. The nonlinearity in QCD arises from the gluons carrying “color” charges and therefore interacting with themselves, with significant consequences. For instance, the strength of interaction

between quarks becomes stronger at larger distance, to the point where the vacuum will “spark” to produce quark–antiquark pairs when their separation becomes on the order of 1 fm.

Further, symmetries of the underlying laws responsible for the dynamics of the quarks and gluons are masked by the structure of the strong vacuum, and hadrons and nuclei are emergent phenomena resulting from the simple laws obeyed by the quarks and gluons. Reproducing and predicting the spectra, structure, and interactions of the hadrons, and the nuclei that are bound states of protons and neutrons, without the use of HPC has so far eluded scientists. The impact of the nonlinearities is greatly enhanced by quantum fluctuations, rendering analytic techniques of limited utility for quantitatively describing low-energy processes. Soon after the discovery of QCD as a candidate theory describing strong interaction processes, Wilson (1974) formulated lattice QCD as a numerical method to provide the solution to QCD in the low-energy regime.

The United States operates world-leading and complementary accelerator facilities for research into cold QCD. The CEBAF at JLab is the world’s most powerful facility of its kind for studying cold QCD systems, delivering multi-GeV electron beams onto nuclear targets for a precision probe of “cold” nuclear matter (JLab undated[a]). This facility recently underwent an upgrade, at a cost exceeding \$300 million, to increase its beam energy from 6 GeV to 12 GeV. The facility enables precision investigations of the valence structure of the nucleon and nuclei and of the spectra of mesons and baryons and promises to uncover and characterize new exotic states of matter. The \$730million FRIB is currently under construction at MSU and is scheduled to become operational in 2022 (MSU undated). The experimental program that is planned for FRIB will probe the limits of stability of neutron-rich nuclei and allow for a refinement of the nuclear forces that dictate the nature of high-density matter, which is present in the interior of NSs and which imprints itself onto the gravity-wave signals from the inspiraling NSNS systems that are expected to be observed using LIGO (California Institute of Technology undated). The NSAC (DOE SC undated) 2015 Long Range Plan for Nuclear Science (DOE and NSF 2015) recommends the evolution of the JLab and RHIC (BNL undated) programs into one centered around an EIC (Cartlidge 2015), with a focus on unraveling the gluonic and “sea” structure of nucleons and nuclei (Figure 3-27). This plan also recommended the construction of a ton-scale $0\nu\beta\beta$ -decay detector to search for lepton number violation manifesting itself through the $0\nu\beta\beta$ decay of certain nuclei.

Precision lattice QCD studies of the structure, spectra, and interactions of mesons, baryons, and nuclei, and of the fundamental symmetries of nature, are now crucial to the U.S. nuclear physics research program and also to components of the HEP research program. This field is entering an era in which lattice QCD calculations will be routinely performed at the physical quark masses and include electromagnetism. The program of computation is closely tied to the U.S. experimental program in nuclear physics and to the scientific milestones established by the NSAC and provides crucial support to other areas of nuclear and particle theory. The component of the program focused on the spectrum and structure of the mesons and baryons, including the study of exotic states, is crucial in optimizing the scientific productivity of the 12-GeV experimental program at JLab. It is evolving to include anticipated needs of a future EIC. The program on light nuclei, hypernuclei, and nuclear forces complements and supports the FRIB facility and provides input for studies, combined with nuclear many-body theory, of dense astrophysical environments. The main focus is to refine the chiral nuclear forces, including multi-neutron forces, defined through low-energy effective field theories (EFT), which are used as input into nuclear many-body calculations of nuclear structure and reactions. The fundamental symmetries component — which is intimately connected to the hadron structure and nuclear forces program — is essential to understanding the origin of time-reversal (T) violation once a neutron electric dipole moment (nEDM) is observed and to searches for nonstandard model physics at neutron facilities at ORNL (ORNL undated) and Los Alamos National Laboratory (LANL undated). It is also important to improving estimates of L-violating $0\nu\beta\beta$ -decay rates of nuclei in support of the planned ton-scale $0\nu\beta\beta$ -decay experiment.

[NP 3.4.4]**3.4.4 Computing Needs and Requirements**

Lattice QCD calculations in cold QCD for nuclear physics proceed through three distinct stages, as shown in Figure 3-39, each of which requires large-scale HPC resources, and a final data analysis phase, which is less computationally intensive but can utilize advanced data management technologies. The first phase of these calculations involves the generation of representative samples of the strong vacuum, known as gluon-field (or gauge-field) configurations. The second phase is the computation of the propagation of quarks in these fields to produce data objects known as “quark propagators.” The third phase, known as contractions, combines the quark propagators into correlation functions and physical observables. Each phase places a different emphasis on parallelism.

Gauge configuration generation is a strong-scaling problem admitting only data parallelism and is responsible for roughly 25% of the computational cost of current nuclear physics calculations. Because of its strong-scaling needs and linear nature, it is ideally suited to run on long campaigns on capability supercomputers. In contrast, quark-propagator calculations can admit a large additional degree of ensemble-level parallelism. Propagators can be computed independently on each gluon configuration, and many propagators can be computed in parallel on each configuration. Hence, the most effective way to produce quark propagators is in high-throughput ensemble calculations. The majority of this phase of computation is spent in sparse linear solvers, which are bound in terms of on-node performance by memory bandwidth in the throughput regime, and also by communications interconnect bandwidth when running on many nodes in parallel. In the extreme strong-scaling regime with small problem sizes on each node, the bottlenecks also include latencies of the cache and memory subsystems and that of the interconnect fabric. The contraction stage reduces the 4-D data of the quark propagators, where the dimensions are the three space and the time dimension, to 1-D correlation functions in which the remaining dimension is the time dimension of the original lattice. Additional trivial parallelism can be exploited in the reduction of the 3-D spatial data because it can be performed independently for each value of the time dimension.

Accomplishing the cold QCD science objectives in the exascale era will require substantial computational investments. These estimates are based on (1) the aggregated needs of both the gluon field generation program needed to supply the initial data and of analysis projects identified in Section 3.4.2 and (2) the cold QCD case studies. It is estimated that this program will need approximately 4.5 billion node hours in 2020, as measured in GPU node hours of OLCF’s current Titan system (OLCF undated) and 17.2 billion node hours in 2025. The breakdown of these requirements between capability and capacity resources is shown in Figure 3-40. A modest need for growth of the memory per node in future systems is anticipated. Overall memory needs per run will grow from 8 TB today to more than 400 TB over 10 years. One primary factor in this growth is the increase in the number of sites in the space-time volumes required for the finer discretization of space-time, which drives other factors such as an increase in the size of tensor objects to be combined in the final contraction phase.

Meeting the target of spending less than 20% of the computational time performing I/O will require achieving a maximum I/O rate of around 100 GB/sec. It is anticipated that scratch-file system size needs (hot data) will grow to more than 8 PB over 10 years. This growth is primarily driven by the increasing throughput of jobs moving from the systems. Finally the permanent long-term storage needs and archival needs (cold data) will grow to around 31 PB and 315 PB, respectively, by 2025. The breakdown of these requirements between “hot data” and “cold data” is shown in Figure 3-40.

With the increasing data needs, I/O and data management will become increasingly challenging. Gluon-configuration generation produces relatively small amounts of data, and typically only about

5% of the calculation is spent performing I/O, while propagator and contraction codes are expected to spend approximately 20% of the runtime performing I/O. Further, the nature of the data changes between the three phases of calculation. Gluon-configuration generation provides large ensembles of lattices that are typically long lived and are akin to raw experimental data. They can have “hot” lifetimes during phases of propagator and correlation function generation and can remain archived in between such periods. Propagator calculations produce data that are too large to keep long term and are typically discarded after the correlation functions have been produced in the contraction phase. Finally, the contraction phase produces data that are logically stored in databases for random access in the final stage data analysis. Given these factors, it is necessary to improve both that I/O and data management strategies, for example, by taking advantage of burst buffers. The final-stage data analysis needs focus on the selection of portions of correlation function data and fitting these to known models. This may be different from, and perhaps simpler than, some typical present-day data science scenarios of searching for patterns, or classifying data according to identified features. The expertise of ASCR data scientists with the algorithms and technologies in managing, organizing, and analyzing large-scale data can provide fresh insight into our methodology, and is a natural point for an exploratory collaboration to improve data management and analysis techniques currently used in cold QCD calculations.

Effective exploitation of architectural components in modern computer systems (caches, fast and slow memories, and interconnects) in a portable fashion across diverse architecture is challenging, and programs offered by ASCR facilities for application readiness such as NESAP, Theta ESP, and the Center for Accelerated Application Readiness (CAAR) are invaluable, especially because they each feature a performance portability component allowing multiple facilities to be targeted. Partnerships and relationships formed with facilities and their staff, such as consultants/catalysts in these programs, carry forward into production and into long-term algorithmic and performance development partnerships in a natural and organic way. Nuclear physics can also contribute to feeding back experiences to the facilities and other users through participation in workshops, code “hackathons” and training webinars.

Laboratory and university researchers play a crucial role in accomplishing the cold QCD science objectives. At present, those associated with the cold QCD effort are supported by institutional DOE and NSF grants and by SciDAC. Without these researchers, the scientific goals that have already been accomplished would not have been and the objectives in the exascale era will not be realized. We estimate that more than a doubling in the currently SciDAC-supported (or equivalent) researchers in the area of cold QCD is required to accomplish the objectives outlined in this and a previous (DOE SC 2016) report. It is anticipated that the recently supported ECP in lattice QCD will meet some of this need.

The capacity-computing hardware that is operated by USQCD through DOE support continues to be essential to the cold QCD program. In addition to providing half of the computing resources available to cold QCD in the United States, it has proven itself to be essential in “standing up” small projects and new ideas for algorithms and physics thrusts, and bringing them to the point at which they can be integrated into the large-scale physics program supported on leadership-class compute platforms. This is intimately related to the concept of the “speed of science,” in which having a component of the ecosystem dedicated to taking an idea from the moment of conception through to the first attempt at execution in a short time, is important for the project and in nurturing the curiosity and agility of junior scientists. The USQCD project has been at the forefront of supercomputer hardware design with the QCDSF, QCDOC that led to the Blue-Gene series of IBM machines, and it deployed the first large-scale GPU cluster (located at JLab). Needs have been identified for an expansion of the capacity-computing hardware available to the cold QCD program. The 2015 Long Range Plan for Nuclear Science (DOE and NSF 2015) contains an initiative to address this and other points.

The cold QCD program has proven itself to be an outstanding arena in the training of junior scientists. This is because of the diverse skill sets that junior scientists have access to and can choose from, ranging from the domain sciences, to HPC, to algorithm creation, to software production, to work-flow design and, of course, the actual running on the largest supercomputers in the world. Junior scientists who have been trained within USQCD are successfully in career paths in universities, national laboratories and private industry, for example, data-science and related companies and hardware producers, such as IBM and NVIDIA. It is critical to increase the resources available for the training of junior scientists in this multidisciplinary environment that arises naturally on the road to accomplishing the scientific objectives of cold QCD. The training provided in this area is well matched to significant workforce needs of the national laboratories.

With the volume of data expected to be produced and retained, there is a growing need for implementing community-wide data-curation and data-provenance practices. The present team-based structure of the community, each working toward a number of science objectives, while sharing some of the intermediate results, such as gluon configurations and quark propagators, has led to the community-wide practice of noncompete agreements between scientists as a way to protect their intellectual property, consistent with data management plans. The data and results, in many instances, reside in repositories accessible by the larger community. It is desirable to have a system and set of protocols that allow usage agreements and data provenance to be openly and readily available to the community. Ideally, a system would evolve toward a complete tracking of data from creation, through verification and storage, through analysis, to publication. An integrated data genealogy that is accessible to the wider community would be of benefit.

[NP 4]

4 PATH FORWARD

This review addressed the extensive, diverse, and evolving requirements of the exascale HPCDE required to meet the scientific challenges of the DOE and NSF nuclear physics research programs. The diversity of the nuclear physics HPCDE emerges from the need to solve QCD in the hadronic and low-energy regimes and at extremes of temperature and density, calculate the structure and reactions of nuclear many-body systems, perform large-scale nuclear astrophysics simulations, and acquire and analyze large data sets generated in the nation's laboratories using real-time dataprocessing capabilities. This HPCDE includes the following:

- Exascale capability computing resources comprising heterogeneous architectures with commensurate hot- and cold-data storage capabilities;
- Enhanced infrastructure for data-handling capabilities — including transfer, archiving, curation, and large-scale databases;
- Development of a highly skilled and sustainable workforce, including the creation of permanent positions at the interface between nuclear physics and applied mathematics and computer science; and
- Increased capacity computing resources at universities and nuclear physics laboratories.

Increased investment in each of these components of the ecosystem is required to allow nuclear physicists to optimally address the scientific challenges using exascale-computing resources. The integration of an exascale HPCDE into the nuclear physics research program, in coordination with advances in nuclear theory, experiment, applied mathematics, statistics, and computer science, will provide unprecedented predictive capabilities in subatomic physics and is assured to spark scientific discoveries that we cannot imagine today.

[NP 4.1]**4.1 Designing an Exascale Ecosystem for Nuclear Physics**

For researchers to successfully achieve the scientific goals described in this report, they need access to a rich computing and data ecosystem. This joint NP-ASCR review identified items of importance to the NP community. Scientists performing research of interest to NP need to (in no particular order):

- Solve computational problems of extreme complexity and magnitude across a very wide range of physical scales;
- Explore parameter spaces and quantify uncertainties using ensembles of calculations;
- Verify and validate algorithms, codes, and models;
- Develop and optimize new codes, algorithms, models, and workflows for exascale computing architectures;
- Read, write, manage, analyze, curate, and track data of complexity and scale never before encountered;
- Grow and sustain a workforce to carry NP computational science through the exascale era, including enhanced collaboration between NP, ASCR, and NSF, and creation of permanent positions at the interface between NP and ASCR; and
- Enhance local capacity hardware at laboratories and universities that is capable of efficiently integrating with ASCR infrastructure and facilities.

Each of these tasks is described in the paragraphs that follow.

Solve computational problems of extreme complexity and magnitude across a very wide range of physical scales

The computational resources required during the next decade by nuclear physicists have been established. Figure 41 shows the requests for compute resources during 2016, along with projected resource requests for 2020 and 2025. All areas except Experiment and Data have very large requirements for heterogeneous resources, either GPU or many-core. Experiment and Data and all areas have modest projected capacity computing needs for conventional core resources. By 2025, nuclear physics is projected to need exascale data storage for both cold- and hot-data at the levels shown in Figure 4-1, and the associated infrastructure to transfer these data for subsequent use and analysis. Nuclear physics research has a clear need for both exascale capability computing resources and exascale capacity computing resources and the associated global data-handling infrastructure. The nature of some of the calculations required to be performed in the ecosystem have large in-node memory needs, along with the need for fine control over memory hierarchies. Further, these nodes will be required to have high-speed I/O capabilities.

Computing at exascale requires access to exascale hardware, supporting software, and applications that can run efficiently in an exascale environment. ASCR facilities already host pre-exascale systems and, in conjunction with domain scientists (NESAP at NERSC, CARR at OLCF, and ESP at ALCF), they have developed successful and effective application readiness programs and will continue to do so. ASCR facilities will host exascale systems when they become available, with configurations, options, and engineering-developed capabilities informed by the needs and requirements laid out in the Exascale Requirements Review reports.

The facility application readiness programs are focused on achieving performance of specific algorithms on specific hardware architectures. But that is not enough to take full advantage of the opportunities afforded by exascale. New algorithms and mathematical formulations will be needed,

and new computing and runtime paradigms need to be explored. There is great need in all areas of nuclear physics for enhancing the collaboration between NP domain scientists and ASCR applied mathematicians, statisticians, computer scientists, and facility experts – along with HEP, BES, and other physicists. The goal of such collaborations, which will build on and enhance the very successful SciDAC collaborations, will be to develop codes and work flows to operate in an exascale environment.

Explore parameter spaces and quantify uncertainties using ensembles of calculations

Modern HPCDEs give scientists the ability to explore the appropriate parameter space in calculations and simulations. This capability is required in order to, for example, test theories and assumptions, find optimal solutions, propagate uncertainties, validate models, and quantify uncertainties. These endeavors require access to large numbers of core-hours distributed across a spectrum of job scales on HPCDEs. Allocations of adequate size — often mid-scale computing allocations — are needed to support what are often referred to as “ensembles” that are run to enable solution of exascale-size challenges with a complete quantification of associated uncertainties.

Ensemble calculations and simulations, and the experimental program, will require rapid-access and addressable databases. Such databases will also be required for data curation and longterm effective staging of hot and cold data for post-production analysis. Currently, database are integrated into some production codes; an enhanced effort, through collaboration with ASCR in developing databases suitable for conducting nuclear physics in the exascale era, is now required.

Verify and validate algorithms, codes, and models

On-demand access to hardware of the same or similar architecture as the exascale resources is vital for code development, testing, and scaling studies. Enhanced coordination between NP, NSF, and ASCR facilities in planning for the requirements of nuclear physics would be beneficial in assuring that such resources are built into procurement and NSF, ERCAP, INCITE, and ALCC allocation planning. Capacity computing resources at universities and national laboratories have traditionally provided the infrastructure on which such development and validation occur, primarily because of the “speed of science” and the ability to gain access without the typically long-delays associated with peer review and committee evaluation.

Develop and optimize new codes, algorithms, models, and work flows for exascale computing and architectures

Creating and implementing efficient, high-performance scientific applications, and the work flows used to execute them at exascale, are critical to accomplishing the science objectives of nuclear physics. Of increasing importance, as computer architectures evolve, is the development of abstractions and tools for performance portability. Developing close collaborations between nuclear physics researchers and ASCR researchers and staff at the facilities and building upon those that have been set in motion with the SciDAC projects, is essential to achieving this level of computational sophistication.

Read, write, manage, analyze, curate, and track data of complexity and scale never before encountered

Issues related to data, and data science, are now at the forefront of HPC, and this trend is expected to continue in the exascale era. These issues include (1) integration with existing (non-HPC) work flows and methods for data acquisition, transfer, and analysis in order to leverage the scale of resources available at HPC centers, and (2) working seamlessly across computational and data environments at all scales and locations relevant to nuclear physics.

As calculations and simulations reach exascale, the time needed to produce output data and checkpoint files with current technology and algorithms grows to unacceptable levels. A close collaboration between ASCR and the nuclear physics community can address this issue through development of algorithms and technologies like burst buffers and identification of system balances that accommodate the required I/O rates.

Experimental nuclear physics science is embracing HPC to solve its significant needs for additional computational resources and extreme data acquisition, storage, sharing, management, and analysis. The nuclear physics experimental community has a mature and widely accepted software, data management and movement, and analysis infrastructure that will have to interoperate with HPC facilities. Significant effort will be required to prepare nuclear physics experimental software and work flows and develop portable software stacks for HPC. Such efforts will also be required to accommodate existing and evolving nuclear physics infrastructure at the HPC facilities and the fading distinction between online and offline analysis. Moving forward, ASCR, NSF, and NP scientists working together can help ensure that state-of-the-art analysis techniques are employed and are able to run efficiently at ASCR and NSF HPC facilities.

Further, there are emerging needs for enhanced real-time data capabilities. Deploying appropriate computational hardware at the nuclear physics experimental facilities that have architectures similar to those of ASCR facilities may be of benefit in interfacing nuclear physics experimental data requirements with ASCR HPCDEs.

First-of-their-kind HPC and data systems are — by their very nature — unique, with special environments. The diversity across ASCR and NSF HPC facilities — each with their special environments, requirements, and practices — can make it difficult for researchers to operate across facilities. But by the facilities working with each other and the science community, the key obstacles can be addressed and minimized. Examples include implementing common identification and authorization mechanisms, defining and implementing a standard base set of HPC software and tools, standardizing procedures and namespaces for using data transfer nodes, and developing a common set of tools to collate data with their corresponding data management plans.

There is a need for enhanced collaboration between ASCR and NP in data management across all areas of nuclear physics, including database optimization, coordination with approved data management plans, access control, data tracking from creation to deletion, and data curation.

Grow and sustain a workforce to carry NP through the exascale era, including enhanced collaboration between NP, ASCR, and NSF and creating permanent positions at the interface between NP and ASCR

The SciDAC projects provided the organization within the nuclear physics community to optimize the science output from tera- and petascale compute resources. The funding from NP and ASCR was used to support junior scientists — typically postdoctoral fellows and graduate students and fractions of senior scientists — in collaborating to develop new algorithms and optimize codes on evolving capability and capacity computing hardware. Without the SciDAC programs, the nuclear physics community would not have codes running efficiently on, for instance, GPU and many-core architectures nor on the Blue-Gene series. It is critical that such support for scientists and such collaborations continue through the exascale era. As identified in *Reaching for the Horizon: The 2015 Long Range Plan for Nuclear Science* (DOE and NSF 2015), the *Nuclear Physics Workforce Report* (NSAC 2014[b]), and the National Academies nuclear physics decadal study report (National Academies Press 2013), a significantly larger HPCD-skilled workforce is required than exists today to meet the grand challenges facing nuclear physics.

Computational scientists in NP require significant support in developing and optimizing codes, algorithms, and work flows for next-generation systems. Building upon the developments and organization within the nuclear physics community that have emerged from the SciDAC projects, the NESAP program at NERSC, CARR at OLCF, ESP at ALCF, and ECP have initiated this process. Lessons learned and the best practices from these programs are being disseminated to DOE and NSF HPC users, in general, through the facilities training programs. All these efforts are valuable and are expected to continue. In addition, general training in HPC is being carried out, with the Argonne Training School on Extreme-Scale computing being the current premier example. The ASCR facilities will team with the ECP to continue and expand programs like this.

Establishing viable career paths to facilitate the transition of junior scientists to senior scientists who remain associated with nuclear physics research efforts — either at universities or at the national laboratories — is becoming urgent. The larger question of workforce development, achieved through programs for career development, recognition, acceptance, and placement, is recognized as a major concern for science in general.

The major areas of collaboration that have proven fruitful for nuclear physics have been with high-energy, condensed matter, and atomic physicists and applied mathematicians and computer scientists. These collaborations must continue and become significantly stronger through the exascale era. In addition, collaborations with statisticians are forming, and they are expected to become essential considering the need for uncertainty quantification in theoretical computations and the dramatic increases in data volume and complexity that are expected in the exascale era.

Enhance local capacity computing and associated infrastructure at laboratories and universities that is capable of efficiently integrating with ASCR infrastructure and facilities

Data acquisition, processing, and analysis, which are key elements of the nuclear physics experimental program, require a diverse exascale ecosystem. Significant elements of the ecosystem — those required to acquire data from the detector, apply acceptance criteria, rapidly move accepted events to disk, and format the data for distribution to scientists — is accomplished at the national laboratories on local capacity-computing hardware. In the exascale era, petascale capacity-computing resources will be required at the national laboratories to accomplish the tasks that are necessarily local to the experimental detectors. Comparable-scale resources will be required at universities to perform the subsequent analyses and code optimization; such resources are critical to the education of junior scientists who will form the future workforce.

Capacity-computing resources at universities and national laboratories have proven invaluable in developing the codes and algorithms used for large-scale simulations and calculations in nuclear physics. In addition, such resources permit exploration of new ideas at a modest scale without the need for writing a proposal to be peer reviewed and possibly supported — a process that typically takes months. The nuclear physics long-range plan (DOE and NSF 2015) has identified a need for enhanced mid-scale capacity-computing hardware at universities and national laboratories to meet these increasing needs.

[NP C.2.4]









C.2 White Paper Addressing Experiments and Data

Nuclear Physics White Paper on Experiment and Data

Amber Boehnlein (Thomas Jefferson National Accelerator Facility),
Jason Detwiler (University of Washington), Paul Mantica (Michigan State University),
and Jeff Porter (Lawrence Berkeley National Laboratory)

4. The most important needs for computing and data in 2020 and 2025 relative to today

The most significant computing need in nuclear physics experiment will be for effective real-time processing capability to manage the large data rates expected in the 2020–2025 time frame. Rates at over 1PB/day will become the norm, outpacing the long-term storage capacities that are predicted to exist for the community. Furthermore, at those interaction rates, conventional trigger systems used today to select the most interesting events for storing and later analysis will, in many cases, not be able to adequately disentangle those events from the rest of the data stream. The nuclear physics experiment community recognizes this challenge, and effort is being directed to develop such near-line capabilities for triggerless data streams. However, the community is simultaneously engaged in ongoing experiments that make it difficult to do the research and development work needed for such a large paradigm shift in online operations. The community would benefit greatly from the development of better algorithms (such as Machine Learning methods) and data-processing tools for lossless real-time data reduction near the beam line. As can be seen in Table 1, nuclear physics experiments do not require large amounts of overall computational capabilities to reach our scientific goals. At the same time, the data-processing capabilities required by the community are illustrated by the 10s to 100s PB of online storage capacity needed for managing access to the very large analysis-ready data sets. The community expects that the new methods and tools developed for processing online data streams can also be used to enhance the throughput in offline data analysis. Use of such methods will require an HPC ecosystem that supports network-friendly adaptive work flows needed in such large data-processing tasks.



APPENDIX L

Sections Cited in This Report
from the Exascale Requirements
Review for Advanced Scientific
Computing Research (ASCR)

This page is intentionally left blank.

SECTIONS CITED IN THIS REPORT FROM THE EXASCALE REQUIREMENTS REVIEW FOR ADVANCED SCIENTIFIC COMPUTING RESEARCH (ASCR)

The following sections from the Advanced Scientific Computing Research (ASCR) Exascale Requirements Review report are cited at least once in this Crosscut report.

[ASCR ES.3]

Appendix F contains the ASCR Executive Summary.

[ASCR 2.1]

2.1 ASCR Mission

The mission of the ASCR program is to advance applied mathematics and computer science; deliver the most advanced computational scientific applications in partnership with disciplinary science; advance computing and networking capabilities; and develop future generations of computing hardware and software tools for science, in partnership with the research community, including U.S. industry. The strategy to accomplish this has two thrusts: (1) developing and maintaining world-class computing and network facilities for science; and (2) advancing research in applied mathematics, computer science, and advanced networking.

The ASCR Research Division underpins DOE's world leadership in scientific computation by supporting research in applied mathematics, computer science, high-performance networks, and computational partnerships such as Scientific Discovery through Advanced Computing (SciDAC).

The *Applied Mathematics* program supports mathematical and computational research that facilitates the use of the latest HPC systems in advancing our understanding of science and technology. More specifically, this program develops mathematical descriptions, models, methods, and algorithms to describe and understand complex systems, often involving processes that span a wide range of time and/or length scales.

The *Computer Science* program allows innovative advancement in a broad range of topics, including programming systems, system software, architectures, performance and productivity tools, and many others. In particular, the program focuses on effective use of very large-scale computers and networks, many of which contain thousands of multi-core processors with complex interconnections and data movement.

The *Next Generation Network for Science* program in ASCR conducts research and development (R&D) activities to support distributed high-end science. It focuses on end-to-end operation of high-performance, high-capacity, and middleware network technologies needed to provide secure access to distributed science facilities, HPC resources, and large-scale scientific collaborations.

SciDAC brings together computational scientists, applied mathematicians, and computer scientists from across application domains and from universities and national laboratories to ensure that scientists are using state-of-the-art technologies to solve their increasingly complex computational and data science challenges.

[ASCR 3.2.1.1]**3.2.1.1 Needs Identified from Breakout Session**

This breakout session, coupled with the superfacility special session, focused on the role of data (analytics, movement, storage, and visualization) in emerging architectures. The main finding of these sessions can be summarized simply as the overwhelming desire to create modern, HPC-enabled data centers as we approach the era of exascale computing. Today's petascale simulations produce scores of terabyte (TB)- to petabyte (PB)-sized data sets. Today's experimental facilities, including light sources, genomics centers, and telescopes, produce equally large data sets. With the strong scientific desire to confront observational and experimental data sets with those produced by simulations, and because each of these will grow by more than an order of magnitude over the coming decade, it is readily apparent that investments in both new types of centers, as well as R&D in a variety of data-centric efforts, will be required to handle this deluge of data.

Modern, HPC-Enabled Data Repositories/Centers — Superfacility (1)*Description and Drivers*

A 'superfacility' is a network of connected computing and other scientific facilities (e.g., NERSC, Spallation Neutron Source [SNS]), software, and expertise that will enable new modes of discovery. Each of DOE's experimental and observational facilities is a unique national resource. Seamlessly connecting these resources to HPC facilities has already enabled new discoveries and will be transformative in the future, offering the potential to benefit many more fields in the coming decade. As we move to an era when exascale simulations are directly confronted by, and analyzed in conjunction with, multi-PB experimental and observational data sets, a superfacility will be needed to handle this analysis and thus is our highest priority.

The idea behind having a modern, HPC-enabled data repository or center has emerged already on a per-domain basis, but domain centers are based on (arguably) old technology and are not able to meet the expanding needs of the scientist in terms of capacity, speed, coupled analysis and other data-centric operations, fine-grained-access policies for the compute facilities and data, etc. Several experiments have created their own high-throughput computing (HTC) and HPC "pipelines" to meet their current needs; however, these pipelines are often constricting the potential science that can be done. Consider for instance the two very different science cases of a light-source experiment and a cosmological telescope survey to highlight this point. Both are constructed to answer a specific question, and a pipeline is developed to address this point. Confrontation of the data generated from these pipelines with results from a simulation is completed after the fact; new science questions can only be posed ex post facto, with no ability to steer the experiment while it is collecting data or enable real-time follow-up resources.

Thus, alleviating many of these constraints for the domain scientists could have a large impact on future science. And while here we provide examples from joint experimental and simulation-based analysis, the same holds true for the simulation-only analysis more typically encountered by ASCR scientists. Simulations are reaching the point, even now, where it is infeasible for a given HPC facility to store all the data generated. Thus, analysis needs to be performed in situ in order to stay ahead of the deluge of data, or the data must be moved from one facility to another for subsequent analysis. In addition, researchers are just now asking the question: what could be gained by "steering" a simulation while it is running?

What do users want in terms of a "superfacility"? They would like to bring together aspects of HPC, distributed computing, networking, and facilities personnel in a secure environment, where HPC and data are on equal footing. They would like to see multiple facilities (including compute and experimental) joined together and operated together by means of a high-performance network

where geography is not a constraint, with ubiquitous access and a common software stack at all locations. In this superfacility, the network is explicitly part of the whole, not “just there” or invisible, and federated identity is a requirement to access all resources.

Because certain simulations or experimental runs and subsequent analysis will often generate “bursty” (i.e., non-smooth) resource allocation requirements, the facility must be able to handle elastic compute and storage requirements. In addition, coordination and co-scheduling will often have to be carried out at multiple-component facilities (e.g., a large simulation is run at OLCF and/or ALCF with data streamed to NERSC, where a scientist can perform analysis using one of the real-time queues). In addition, what is “good” or “right” to run at a given location on one day may be different from the best scenario tomorrow. The superfacility would allow the flexibility to make those decisions in a way that is useful for the scientist.

Persistence is a key component of a superfacility. In order to support experiments, the compute, storage, and other resources necessary for its success must be available for the lifetime of the experiment (e.g., decade-long experiment run). This does not imply that an individual center needs to persist for the length of an experiment or be available 24/7. However, the superfacility itself needs to continue to run over this timeframe. Thus the resources that constitute a superfacility should, by necessity, have a long-term funding model.

Discussion

There are many interesting challenges (and accompanying R&D) that need further effort and discussion at the institutional, scientific community, and facility levels. Such challenges include the following.

- Providing high-performance data repositories for data products so that it is faster to access a dataset from a high-performance repository over the network than it is to access it from a local disk (DOE 2015).
- Data properties need to be expanded to include additional attributes, such as value, lifetime, etc., and a policy concerning data retention and “value” needs to be developed.

The time is ripe to create a call for white papers from scientists on how they would make use of a superfacility in the exascale era (or before). Subsequently, we could imagine partnering domain and ASCR scientists with NERSC, ALCF, OLCF, and ESnet to develop demonstrations of a superfacility and discover the bottlenecks. This would likely lead to potential R&D efforts in workflow management, edge services, math/analytics, HTC on HPC, containers, and compression, among other many others.

[ASCR 3.4.2]

3.4.2 Production Systems

Session participants included a broad range of ASCR-supported researchers who develop software throughout various levels of the stack (ranging from low-level system software through intermediate layers of numerical packages and higher-level scientific application codes), as well as tools for profiling and analysis. The group identified two overarching themes in the needs of ASCR facilities that could support software development on production systems:

1. A holistic approach to computing systems provided by DOE facilities, including access to a variety of systems at each facility (to support both small-scale testing and at-scale debugging and tuning) and coordination across facilities; and
2. A cohesive, sustainable cross-facility software ecosystem that enables developer productivity and promotes encapsulation of ASCR research for use by scientific application teams.

Table 3-10 lists the research needs identified for software development in production systems, the scientific drivers associated with those needs, and the potential for impact of each identified need (low, medium, high). The table and the sections that follow elaborate on these themes from the perspective of seven fundamental needs. Most of these needs also apply to early delivery systems. Likewise, many of the needs discussed in Section 3.4.1 (Software Development on Early Delivery Systems) also apply to production systems, although the relative priorities for various topics shift depending on machine and software lifecycles. Several of the needs related to software deployment are only briefly introduced here, because they are discussed in more detail in Section 3.6 (Software Deployment and Support).

[ASCR 3.4.2.1]

3.4.2.1 Needs Identified from Breakout Session

Access to a Variety of Systems (1)

Description and Drivers

In contrast to domain scientists, who need computing resources for large-scale production runs (i.e., access to many nodes for many hours), ASCR-funded researchers need access to computing systems in a variety of modes that explicitly support the development of software that encapsulates cutting-edge advances in applied mathematics and CS, which in turn supports domain science applications. Because DOE facility policies traditionally tend to favor large-scale production runs, ASCR researchers often encounter difficulties in obtaining consistent access to systems for software development. ASCR researchers need DOE computing facilities to balance allocation of computing resources to support both production runs and software development, including software installation, testing, debugging, performance analysis, and enhancement. Development nodes should be expected to be “more available” (short wait time) in comparison to production nodes.

Developers often need interactive access to hardware (as opposed to batch access for production runs). Programmers are more productive when they can debug and profile in real time rather than wait for machine availability. Batch access typically is not sufficiently fast for “cognitive turnaround” needed in many development activities, such as debugging.

Achieving robust, efficient, and scalable performance is more difficult than ever before due to the complexities of emerging architectures; ensuring good performance at scale requires resources for debugging and performance tuning at scale. Short-term access to many nodes (as opposed to long-term access for production runs) is a very different (but important) usage model. Drivers include debugging (may need many nodes for bugs that appear only at scale), profiling, performance optimization, and scalability testing.

In addition, developers may need special privileges on machines for testing different machine configurations, such as setting cache/memory mode, and other changes that require root access or kernel changes.

Discussion

These needs primarily span the later stages of software development (during performance analysis and optimization); earlier phases of development are typically done on laptops and small clusters before transitioning to DOE facilities. An important consideration is that ASCR researchers often need to debug software that has already been deployed (and is already supporting application users) at DOE facilities and elsewhere.

We advocate at least two approaches to address these needs: (1) provide developer access to small clusters that match the environment of production machines, and (2) provide developer access to the full machine for scalability testing. Because these needs for software development span all scientific research areas and apply to early delivery systems, we believe that addressing these issues will increase programmer and research productivity for both ASCR investigators and domain scientists.

Software Ecosystem (2)

Description and Drivers

ASCR research supports science programs across DOE and has the potential to provide a broad range of new math and CS functionality needed for next-generation science, as well as to improve the performance, programmability, reliability, and power efficiency of production applications. However, the impact of research done in ASCR is limited because ASCR's funding model and reward structure do not scale to the large user base of ASCR facilities. Many ASCR research products never see use in real applications because of lack of interest or awareness from application developers, lack of clearly defined guidelines for deploying and supporting software at LCFs, and lack of funding for researchers to make their tools usable in production applications. ASCR researchers need a clear path for transitioning research software to production usage, as well as clear models for maintaining and extending software products that are valuable to application codes. Without these resources, it will be impossible to build a sustainable exascale software ecosystem for use by the broader community.

Discussion

ASCR research software is currently fragmented — many research artifacts are produced by separate groups, but few are used widely in production at facilities. The basic software ecosystem on HPC systems today has remained unchanged for many years. For ASCR software to broadly impact DOE science discoveries, it must be used in production codes. However, application teams are wary of relying on research software. Many have had bad experiences in the past when the software on which they relied ceased to work; some teams opt to implement their own versions of complex techniques and algorithms to avoid relying on external groups. With exascale machines arriving soon, and with the increased architectural diversity they are expected to bring, it will be critical for application developers to leverage research in advanced algorithms, performance portability, tuning, resilience, and power efficiency, to make the best use of future machines. At exascale, application teams will not have the resources to do all algorithmic development and tuning themselves. The working group identified two areas where the facilities could help to build a sustainable software ecosystem:

1. Identify key research products for facility users and help to sustain and grow their development; and
2. Provide a clear set of requirements and processes for software deployment on production systems.

These areas, and their associated drivers, are discussed below.

Identify Key Research Products for Facility Users and Help to Sustain and Grow Their Development

Description and Drivers

In the startup world, companies budget for growth, and the success of a product is measured in terms of the number of users it attracts. Within ASCR, the incentive structure is heavily biased toward paper publications, and researchers have little incentive to develop tools for widespread use.

Indeed, maintenance is 80% of software development, but makes up only a tiny fraction of ASCR funding, and research project budgets are typically flat from year to year. For application developers to adopt software, they need to know that it is stable, reliable, and has long-term support.

Discussion

Software sustainability is an issue that extends beyond the facilities, and productizing research is a very resource-intensive process. The working group recommends that a larger discussion be initiated to establish a viable path for software as it moves from research to production. ASCR and the facilities can help by aggressively lowering barriers to good software practices and by helping research teams transition to production roles. Facilities should help researchers make connections with production application teams from the beginning of their projects. ASCR and the facilities should also help by identifying products with growing adoption and by allocating funding, staff, and hosted collaboration tools so that these projects can easily establish release processes, issue tracking, automated regression testing, version control, and other best practices. Funding and staff allocated to such support efforts should be commensurate with uptake by application developers, so that projects can grow with demand. That is, ASCR and the facilities should plan for success, growth, and long-term sustainability from the start. Facilities should view themselves as supporting a true developer ecosystem in addition to HPC users.

Provide a Clear Set of Requirements and Processes for Software Deployment on Production Systems

Description and Drivers

For research products to be taken up by application developers, the developers must be made aware of the research projects, and the research products must be made available for easy use. Currently, the majority of publicly installed software on facility systems is installed by the facilities, and the procedures through which users can deploy their software and advertise it to other facility users are inconsistent and poorly documented. Researchers have expressed frustration at not being able to find ways to reach potential users at these sites.

Discussion

To remedy this situation, facilities should clearly document requirements for deploying software for public use on HPC systems, and they should provide mechanisms that allow users to advertise their research tools. ASCR researchers should be encouraged to deploy their software widely from the start of application projects and to seek out application users with help from the facilities. Lowering the barriers for software deployment will allow researchers to more easily update and support software for facilities. These topics, as well as automation issues associated with deployment, are covered in more detail in Section 3.6.

Coordinate across Facilities (3)

Description and Drivers

Access to computing systems and installation of research software on these systems are currently handled independently (and typically differently) by each DOE computing facility. In addition, policies for access have traditionally been formulated primarily to serve the needs of science teams that are conducting simulation campaigns. For example, an approach frequently used by developers of numerical libraries is to apply for access to a specific facility for purposes of algorithmic research, testing, and performance tuning. However, access to systems at one DOE computing facility does not extend to other DOE computing facilities; in addition, substantial paperwork is required to re-apply for access during subsequent allocation cycles. Consequently, facility policies

and procedures do not fully address the needs of researchers who develop reusable software that is employed by applications teams on systems at multiple facilities. Unlike science teams, who may naturally focus on just one machine at just one facility at a given time, and who can justify machine usage by the results of simulation campaigns, developers of software packages and CS tools need sustained access to all relevant computing systems at all ASCR computing facilities in order to fully debug and test software so that it is robust, performs well, and is ready for use by applications teams wherever they work.

Additional challenges arise because of the different processes and infrastructure each facility uses. Coordination of access and policies across facilities, including a holistic perspective on machines as a coordinated set of systems, would help to simplify work for ASCR-funded researchers. Session participants expressed the desire for a single federated login name and one-time password token that are valid across all facilities. Also desired are compilers installed in a consistent manner across facilities, common environment variables across facilities (e.g., \$SCRATCH_HOME, \$ARCHIVE_HOME), common queue names, common commands to view allocations and disk quota usage, and common job launchers. Researchers also want tools to be installed and accessible using common environment variables and modules, including portable tools for performance analysis, testing, and debugging.

Discussion

Steps toward coordinated and sustained access to computing facilities, as well as common environments across systems, would lead to improved productivity and tool usage for both software developers and applications users. A possible metric to evaluate the impact of supporting such coordinated access is the time required to transition software between sites.

Automated Software Testing on HPC Facilities (4)

Description and Drivers

Regular and extensive testing helps to ensure the correctness of scientific software, which tends to change frequently in order to address research needs (e.g., refactoring for better performance on emerging architectures and the incorporation of new functionality as needed by new science frontiers) (Bartlett et al. 2016). It is imperative that code changes do not inadvertently introduce new errors or reintroduce old errors, particularly in ASCR research software that supports a broad range of applications teams (Bartlett et al. 2016).

While software testing is time-consuming and challenging at any scale and in any computing environment, testing is especially difficult across DOE computing facilities due to the need to maintain portability across a wide variety of (ever-changing) systems and compilers. Each software package has a broad test space to adequately cover its functionality, and this space is multiplied across different platforms and their corresponding software stacks. Difficulties are compounded because computing resources at facilities are limited, while software researchers need frequent testing access and a relatively quick turnaround time. Two particularly difficult tasks are testing system software and handling performance variability (for performance testing).

Although automated testing is well recognized as a best practice, DOE facility policies do not allow this. ASCR-funded researchers need support for automated software testing on all systems at all computing facilities, including unit, regression, verification, integration, and performance testing. As HPC software packages have become increasingly complex, with multiple developers, support is also needed for continuous integration-style testing (to identify errors as soon as possible after they are introduced) and result aggregation (to help simplify the examination and understanding of test results).

Discussion

DOE computing facilities could consider trying to consolidate testing infrastructure in order to provide better and more consistent testing support across facilities and systems. A metric to consider is the time developers of a software package require to establish automated testing capabilities on a new system.

Debugging Capabilities: Variety of Tools, Debugging at Scale (5)

Description and Drivers

Debugging — the process of isolating and correcting errors or abnormalities in code to allow proper program operation — is part of the software testing process and is an integral aspect of the entire software development lifecycle. A variety of debugging tools are available on conventional systems to help identify and fix code errors, but debugging in parallel is very difficult due to the complexities of parallel programming. Debugging at DOE computing facilities is even more difficult because of challenges with machine access, limited parallel debugger functionality at scale, and the growing complexity of both hardware and software. More types of parallelism, deeper memory hierarchies, and more complex data structures contribute to numerous types of possible errors, including race conditions, memory leaks and access errors, implementation errors, degraded performance, and system errors or faults. While as a community we should be able to do better than printing information about variables (e.g., using ‘print f’ statements), this low-tech approach is the most common way of tracking down errors at scale. ASCR researchers need access to a variety of tools that support debugging at scale, across all systems at all computing facilities, for both static and dynamic analysis of code execution, so that we can identify programming and system errors rapidly and ensure robustness and correctness of software. We also need to capture the system state during execution to accelerate debugging.

Discussion

The opportunity cost of time spent debugging for users of DOE computing facilities is huge. The more effective the tools we have, the more effective we can be (and research and developer staff time is very expensive).

Containers for Portability from Laptops to Computing Centers, Common Software Environment to Improve Workflow (6)

Description and Drivers

ASCR researchers need (1) improvement in portability across the many systems our software uses, (2) pre-built software environments that reduce compilation issues and provide rapid access to many software capabilities, and (3) workflow improvements that are compatible with the broader computing ecosystem. ASCR researchers face these problems in their role as infrastructure providers for scientific applications that often use a variety of third-party software.

Container technologies have emerged as a promising technology to address these needs. Containers allow an application, its dependencies, and their associated run environment to be packaged into a binary image. Investing in container technologies can ease integration, help with portability and regressions, and encourage increased adoption of component-based application development.

Discussion

Container technologies with proven performance and portability are being adopted rapidly in the broader software ecosystem. Containers promise to revolutionize workflows by simplifying the use

of many third-party software components, improving portability, and providing a common software environment for developers and users. Standardizing on containers would also prepare LCF efforts to blend more easily with those of data sciences communities, where these technologies are already being deployed.

Challenges with containers center on the unique needs for DOE computing facilities relative to the mainstream computing community. DOE facilities need to address security issues, and DOE facility users need to access specialized LCF hardware. Industry standards such as Docker do not support these capabilities. Efforts such as Shifter should continue in order to satisfy these additional HPC requirements.

Requirements for Relevant Operational Data from the Facilities (Facility Monitoring Data, Usage Statistics) to Support Software Tool Development (7)

Description and Drivers

In order to effectively focus efforts for software R&D and to facilitate cross-layer performance analysis at scale, ASCR researchers need access to system-wide monitoring and logging capabilities. For example, we need to understand shared resource activity (I/O, parallel file system, and network), power/energy usage, and failure rates for hardware and software. The relevance of various types of data will evolve over time (e.g., data movement, fault tolerance, and power management will become increasingly important); this information will help to ensure that research addresses the most urgent issues. ASCR researchers also need detailed information about tools, compilers, languages, and libraries used on all systems at all facilities. Section 3.7 (Operational Data and Policies) provides further discussion.

[ASCR 3.6.1]

3.6 Software Deployment and Support

3.6.1 Production Systems

Lowering barriers to software deployment at ASCR HPC facilities is critical for DOE's mission and the success of U.S. exascale computing. Software tools produced by ASCR CS and applied mathematics researchers enable physicists, biologists, and other domain scientists to exploit the full power of machines at ASCR facilities and thereby increase scientific output. Often, scientists cannot reap these benefits because research software is not deployed for general use at ASCR facilities. Facilities have limited resources dedicated to software deployment, typically supporting only a core software stack of compilers, programming models, tools, and math libraries. There is currently no well-defined path for researchers (or any facility users) to make their tools available to scientists. Facilities are the nexus of the HPC community. Barriers to software deployment must be eliminated so that facility users can scale beyond the limited resources of facility staff and reach the broader community with their software.

The exascale software ecosystem will comprise a wide array of tools, libraries, programming models, and performance portability frameworks, all of which are expected to be used by DOE application developers. If application developers are to build their simulations using this software, it must be portable and reliably deployed at ASCR facilities. Without better support for software deployment, application scientists will not benefit from the exascale software stack, and scientific productivity will be sacrificed as scientists struggle to deploy and scale their codes on new platforms.

Table 3-13 lists the research needs identified for software deployment and support for production systems, the scientific drivers associated with those needs, and the potential for impact of each identified need (low, medium, high).

[ASCR 4]**4 PATH FORWARD**

For researchers to move forward in addressing ASCR research challenges, an evolving computing ecosystem must support them. While Section 3 describes detailed requirements in specific areas for ASCR research, common themes emerged across the different areas of research; these themes represent potential high-impact paths forward for ASCR research.

4.1 Cross-Cutting Need Areas

Many of the recurring or highest-priority requirements identified during the ASCR research exascale review address communication and collaboration between ASCR research and the ASCR facilities — from knowledge about the availability of resources to policy and process information sharing. Additional high-level requirements include improved access to hardware at all development stages and an ecosystem that is amenable to the ASCR researchers' requirements for development and testing of software. This section identifies impactful first steps to evolve the ecosystem based on ASCR research needs.

Table 4-1 summarizes the primary cross-cutting requirements identified for ASCR research. Most of the Crosscut topic areas fall into two larger areas: (1) hardware capabilities and (2) ecosystem access, capabilities, and policies. In Sections 4.1.1 and 4.1.2, we provide some specific examples of these two high-level, Crosscut areas.

Table 4-1 displays cross-cutting need areas, defined as general technology areas that encompass multiple needs identified by breakout groups in their individual requirements discussions. Crosscutting need areas do not reflect a specific need common across different breakout sessions; rather, they represent a spectrum of needs falling under a common technology umbrella. Each breakout group targeted a specific aspect of ASCR research requirements analysis for HPC facilities within a specific phase of development/deployment (emerging, early delivery, production). Cross-cutting need areas represent an attempt to group together the logically related needs of the individual breakout groups to facilitate a more system-wide approach in responding to those needs.

Such a holistic approach to addressing identified needs could provide greater benefit to the HPC computing ecosystem as a whole than simply addressing point solutions for each specific need. The six cross-cutting need areas listed in Table 4-1 are not prioritized; however, the table does indicate by letter (H for high; M for medium) how an individual breakout group rated a need in its session that was subsequently included in one of the cross-cutting needs areas. Table 4-1 contains only a subset of the needs identified by individual breakout groups in Section 3. Specific needs identified by a breakout group that did not logically fall into one of the cross-cutting areas stand on their own merits, as detailed in Section 3.

[ASCR 4.1.2]**4.1.2 Ecosystem Access, Capabilities, and Policies**

ASCR researchers rely on access to emerging, pre-production, and production HPC resources to execute their mission. Hardware resources also require capabilities and policies that are welcoming to CS and applied mathematics research. The traditional use models of HPC resources are focused on domain science campaigns that might not need the level of system interaction that ASCR researchers need. As a result, different consideration is needed to address how the HPC facilities operate.

- ASCR researchers require scheduling, interaction, and allocations that are more in line with their research needs (Sections 3.1.1, 3.2.2, 3.3.1).
- Distributed computing support is needed at HPC facilities (Section 3.3).

- ASCR researchers must be able to interact, test, and interface with new software technologies on ASCR facility hardware. This need spans most research, including software development, data, and networking (Sections 3.3.1, 3.4, and 3.5).
- Growing data needs are driving the ASCR facilities to consider a different balance of capabilities for the future; the need for this new capabilities balance is echoed in this report (Section 3.2.1).
- Support — through policies and infrastructure for software development, data, and networking — and training are needed.
- Access to logs and operational data for research is required; also required are consistent cross-facility logging, monitoring, and sharing and privacy policies (Sections 3.1.3 and 3.7).
- Enhanced communication is needed:
 - Increased exposure to available facility resources and programs;
 - More regimented access to emerging systems, emulators, and early systems to facilitate research and reduce delays in software deployment (Sections 3.1 and 3.2.2); and
 - Communication among the facilities, vendors, and users of early delivery systems.

This page is intentionally left blank.

