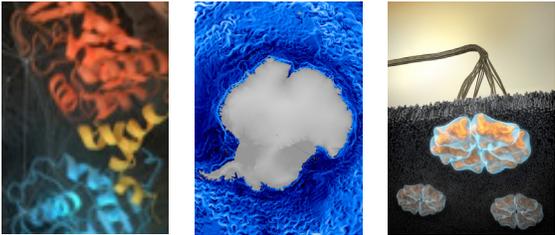
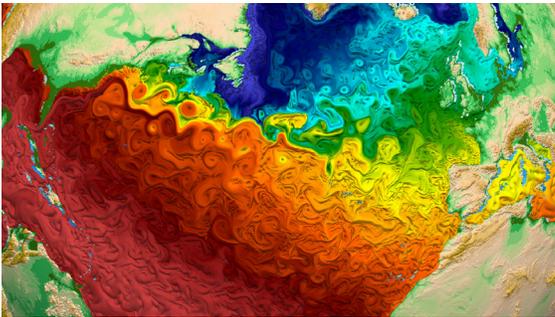


BER

BIOLOGICAL AND ENVIRONMENTAL RESEARCH



EXASCALE REQUIREMENTS REVIEW



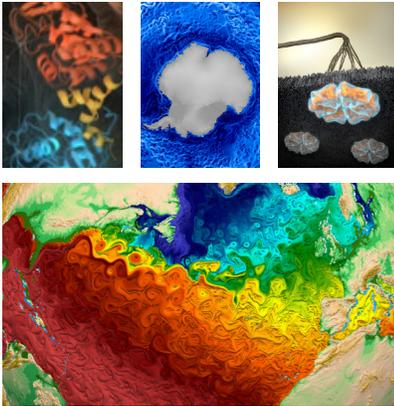
An Office of Science review sponsored jointly by
Advanced Scientific Computing Research and
Biological and Environmental Research

MARCH 28-31, 2016

ROCKVILLE, MARYLAND



U.S. DEPARTMENT OF
ENERGY



On the cover:

Top left: Multi-timescale simulation of a single protein. A phosphoglycerate kinase protein was subjected to MD simulations on various supercomputing architectures. The relative motions of the red and blue domains of the proteins are highly complex and can be described in terms of the motion of a configurational point on a rough energy landscape (illustrated). The transitions of the structure between energy minima on the landscape can be described in terms of a network (illustrated), which was found to be fractal (self-similar) over 13 decades of time (Image credit: Thomas Spletstoesser, scistyle.com).

Top middle: Ocean currents and eddies in a high-resolution global ocean simulation. Colors show speed, where white is faster. Detailed turbulent structures are visible throughout the Southern Ocean, where the Antarctic circumpolar current flows eastward around the globe. Large eddies are particularly visible in the Agulhas current at the southern tip of Africa. These ocean simulations are validated against satellite and shipboard observations. The domain includes 100 vertical layers and 1.5 million horizontal grid cells ranging from 10 to 30 km in diameter, and was run on 8000 processors. (Image credit: The Model for Prediction Across Scales-Ocean [MPAS-Ocean], a component of the U.S. Department of Energy's new Accelerated Climate Model for Energy [ACME] and developed at Los Alamos National Laboratory [LANL]. Image by Phillip Wolfram and Mark Petersen of the MPAS-Ocean team, which also includes Todd Ringler, Xylar Asay-Davis, Mathew Maltrud, Luke Van Roekel, Milena Veneziani, and Jon Wolfe, all of LANL.)

Top right: Model of the cellulase synthase enzyme, Cesa, derived by integrating neutron scattering and high-performance computing (Image credit: Thomas Spletstoesser, scistyle.com).

Bottom: The paint-like swirls of this visualization depict global water-surface temperatures, with the surface texture driven by vorticity. Cool temperatures are designated by blues and warmer temperatures by reds. Trapped regions of warmer water (red) adjacent to the Gulf Stream off the eastern coast of the United States indicate the model's ability to simulate eddy transport of heat within the ocean, a key component necessary to accurately simulating global climate variability. (Image credit: The Model for Prediction Across Scales-Ocean [MPAS-Ocean], a component of the U.S. Department of Energy's new Accelerated Climate Model for Energy [ACME] and developed at Los Alamos National Laboratory [LANL]. Image by Phillip Wolfram and Mark Petersen of the MPAS-Ocean team, which also includes Todd Ringler, Xylar Asay-Davis, Mathew Maltrud, Luke Van Roekel, Milena Veneziani, and Jon Wolfe, all of LANL.)

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.

BIOLOGICAL AND ENVIRONMENTAL RESEARCH

U.S. DEPARTMENT OF ENERGY

An Office of Science review sponsored jointly by
Advanced Scientific Computing Research and Biological and Environmental Research

Lead and Contributing Authors BER Organizing Committee

Adam Arkin, *Meeting Co-Chair*, Lawrence Berkeley National Laboratory
David C. Bader, *Meeting Co-Chair*, Lawrence Livermore National Laboratory

Tjerk P. Straatsma, Oak Ridge National Laboratory
Jack C. Wells, Oak Ridge National Laboratory
James J. Hack, Oak Ridge National Laboratory
Richard Coffey, Argonne National Laboratory
Eli Dart, ESnet

Report Section Leads: Biological Systems Science Division (BSSD)

Multiscale Biophysical Simulation from Molecules to Cells

Matt Jacobson, University of California-San Francisco
Jeremy Smith, Oak Ridge National Laboratory

Mapping Sequence to Models

Rich Bonneau, New York University
Dan Rokhsar, Lawrence Berkeley National Laboratory

Microbes to the Environment

Lee Ann McCue, Pacific Northwest National Laboratory
Tim Scheibe, Pacific Northwest National Laboratory

Biological Big Data Challenges

Eoin Brodie, Lawrence Berkeley National Laboratory
Kathy Yelick, Lawrence Berkeley National Laboratory

Report Section Leads: Climate and Environmental Sciences Division (CESD)

Atmospheric Simulation and Data Assimilation within the Earth System

Atmospheric Research/Simulation

Bill Gustafson, Pacific Northwest National Laboratory
Minghua Zhang, Stony Brook University

Data Assimilation, Model Initialization, and Reanalysis

Gil Compo, National Oceanic and Atmospheric Administration

Terrestrial and Subsurface Research

Dave Moulton, Los Alamos National Laboratory
Peter Thornton, Oak Ridge National Laboratory

Oceans and Cryospheric Research

Wieslaw Maslowski, Naval Postgraduate School
Todd Ringler, Los Alamos National Laboratory

Coupled System Integration – Earth System Models

Bill Collins, Lawrence Berkeley National Laboratory
Phil Jones, Los Alamos National Laboratory
Nathan Urban, Los Alamos National Laboratory

Integrated Assessment Modeling

Kate Calvin, Pacific Northwest National Laboratory
Andy Jones, Lawrence Berkeley National Laboratory

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

Large-Scale Heterogeneous Data Management

Dean Williams, Lawrence Livermore National Laboratory

Observational Data Processing: Retrieval Algorithms and Instrument Network Simulation

Pavlos Kollias, Stony Brook University

Uncertainty Quantification and Analytics of Observations and Model Results: Uncertainty Quantification, Statistics, Emulators, and Analytics of Model Observation Comparisons

Bert Debusschere, Sandia National Laboratories

Model Development Testbeds

Hsi-Yen Ma, Lawrence Livermore National Laboratory

Transforming Science through Exascale Capabilities: Algorithms and Computational Science

Algorithms

Katherine Evans, Oak Ridge National Laboratory

Programming Languages and Programming Models

Esmond Ng, Lawrence Berkeley National Laboratory

Software Engineering for Portability

Rob Jacob, Argonne National Laboratory
Pat Worley, Oak Ridge National Laboratory

Report Section Leads

Biological and Environmental Research: Vision and Grand Challenges

Ruby Leung, Pacific Northwest National Laboratory
Mark Taylor, Sandia National Laboratories

Path Forward (ASCR)

Katie Antypas, National Energy Research Scientific Computing Center
Richard Gerber, National Energy Research Scientific Computing Center
Katherine Riley, Argonne Leadership Computing Facility
Tjerk Straatsma, Oak Ridge Leadership Computing Facility

Sponsors and Representatives

Dorothy Koch, U.S. Department of Energy

Ramana Madupu, U.S. Department of Energy

Carolyn Lauzon, U.S. Department of Energy

Additional Contributors

Argonne National Laboratory Document Design and Development Team

Mary Fitzpatrick, Andrea Manning, Michele Nelson, and Vicki Skonicki, Communications, Education, and Public Affairs

Beth Cerny and Richard Coffey, Argonne Leadership Computing Facility

This page is intentionally left blank.

| | | |
|----------|--|-----------|
| 2 | Biological and Environmental Research: Vision and Grand Challenges | 15 |
| 2.1 | BER Vision | 15 |
| 2.1.1 | BSSD Vision | 16 |
| 2.1.2 | CESD Vision | 17 |
| 2.2 | BER Grand Challenges | 18 |
| 2.2.1 | BSSD Objectives and Grand Challenges | 18 |
| 2.2.2 | CESD Objectives and Grand Challenges | 20 |
| 2.3 | Mapping the BER Grand Challenges to the Computing Ecosystem and Exascale Mapping | 21 |
| 2.3.1 | Mapping BSSD Programmatic Objectives to Computing Ecosystems | 21 |
| 2.3.1.1 | OMICS | 21 |
| 2.3.1.2 | Computational and Data Challenges Associated with Integrative Approaches for Modeling Molecular and Cellular Systems | 21 |
| 2.3.2 | Mapping CESD Programmatic Objectives to Computing Ecosystems | 23 |
| 2.4 | Report Roadmap | 26 |
| 3 | BER Research Directions and Computing Needs/Requirements | 29 |
| 3.1 | Biological Systems Science Division (BSSD) | 29 |
| 3.1.1 | Multiscale Biophysical Simulation from Molecules to Cells | 31 |
| 3.1.1.1 | Scientific Challenges and Opportunities | 31 |
| 3.1.1.2 | Priority Research Directions | 32 |
| 3.1.1.3 | Cross-Cutting Research Directions | 36 |
| 3.1.1.4 | Computing Needs and Requirements | 37 |
| 3.1.2 | Mapping Sequence to Models | 39 |
| 3.1.2.1 | Scientific Challenges and Opportunities | 39 |
| 3.1.2.2 | Priority Research Directions | 39 |
| 3.1.2.3 | Cross-Cutting Research Directions | 43 |
| 3.1.2.4 | Computing Needs and Requirements | 43 |
| 3.1.3 | Microbes to the Environment | 44 |
| 3.1.3.1 | Scientific Challenges and Opportunities | 44 |
| 3.1.3.2 | Priority Research Directions | 45 |
| 3.1.3.3 | Cross-Cutting Research Directions | 47 |
| 3.1.3.4 | Computing Needs and Requirements | 48 |
| 3.1.4 | Biological Big Data Challenges | 50 |
| 3.1.4.1 | Scientific Challenges and Opportunities | 50 |
| 3.1.4.2 | Priority Research Directions | 50 |
| 3.1.4.3 | Cross-cutting Themes and Methods in Data Analytics | 56 |
| 3.1.4.4 | Computing Needs and Requirements | 60 |

| | | |
|---------|---|-----|
| 3.2 | Climate and Environmental Sciences Division (CESD)..... | 63 |
| 3.2.1 | Atmospheric Simulation and Data Assimilation within the Earth System..... | 63 |
| 3.2.1.1 | Atmospheric Research/Simulation | 63 |
| 3.2.1.2 | Data Assimilation, Model Initialization, and Reanalysis ... | 68 |
| 3.2.2 | Terrestrial and Subsurface Research..... | 72 |
| 3.2.2.1 | Scientific Challenges and Opportunities | 72 |
| 3.2.2.2 | Priority Research Directions | 73 |
| 3.2.2.3 | Cross-Cutting Research Directions | 73 |
| 3.2.2.4 | Computing Needs and Requirements | 74 |
| 3.2.3 | Oceans and Cryospheric Research | 75 |
| 3.2.3.1 | Scientific Challenges and Opportunities | 75 |
| 3.2.3.2 | Priority Research Directions | 77 |
| 3.2.3.3 | Cross-Cutting Research Directions | 79 |
| 3.2.3.4 | Computing Needs and Requirements | 80 |
| 3.2.4 | Coupled System Integration – Earth System Models..... | 81 |
| 3.2.4.1 | Scientific Challenges and Opportunities | 81 |
| 3.2.4.2 | Priority Research Directions | 81 |
| 3.2.4.3 | Cross-Cutting Research Directions | 82 |
| 3.2.4.4 | Computing Needs and Requirements | 83 |
| 3.2.5 | Integrated Assessment Modeling..... | 85 |
| 3.2.5.1 | Scientific Challenges and Opportunities | 85 |
| 3.2.5.2 | Priority Research Directions | 85 |
| 3.2.5.3 | Cross-Cutting Research Directions | 86 |
| 3.2.5.4 | Computing Needs and Requirements | 87 |
| 3.2.6 | Transforming Science through Exascale Capabilities: Model-Data Fusion and Testbeds | 89 |
| 3.2.6.1 | Large-Scale Heterogeneous Data Management..... | 89 |
| 3.2.6.2 | Observational Data Processing: Retrieval Algorithms, Instrument Network Simulation..... | 92 |
| 3.2.6.3 | Uncertainty Quantification and Analytics of Observations and Model Results: Uncertainty Quantification, Statistics, Emulators, and Analytics of Model-Observation Comparisons..... | 96 |
| 3.2.6.4 | Model Development and Testbeds | 99 |
| 3.2.7 | Transforming Science through Exascale Capabilities: Algorithms and Computational Science | 102 |
| 3.2.7.1 | Algorithms..... | 102 |
| 3.2.7.2 | Programming Languages and Programming Models ... | 104 |
| 3.2.7.3 | Software Engineering for Portability..... | 107 |

| | |
|---|------------|
| 4 Path Forward | 111 |
| 4.1 Methods Development | 111 |
| 4.2 Computational Environment and Resources | 112 |
| 4.3 Data | 113 |
| 4.4 Communication and Community Involvement | 114 |
| 4.5 Conclusion | 115 |
| 5 References | 117 |
| 6 Acronyms and Abbreviations | 123 |
| Appendix A: Biological and Environmental Research Organizing Committee and Review Participants | A-1 |
| Appendix B: Biological and Environmental Research Meeting Agenda | B-1 |
| Appendix C: Biological and Environmental Research (BER) White Papers | C-1 |
| Appendix D: Biological and Environmental Research (BER) Case Studies | D-1 |

FIGURES

| | | |
|-------------|---|-----|
| Figure 3-1. | 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. | 33 |
| Figure 3-2. | Multi-timescale simulation of a single protein. A phosphoglycerate kinase protein was subjected to MD simulations on various supercomputing architectures. The relative motions of the red and blue domains of the proteins are highly complex and can be described in terms of the motion of a configurational point on a rough energy landscape (illustrated). The transitions of the structure between energy minima on the landscape can be described in terms of a network (illustrated), which was found to be fractal (self-similar) over 13 decades of time. | 34 |
| Figure 3-3. | Model of the cellulase synthase enzyme, CesA, derived by integrating neutron scattering and HPC..... | 37 |
| Figure 3-4. | Projected estimate for file size of an MSI data set of a single cross-section through a human brain using various MSI techniques. As advances in spatial and mass resolution are combined with new dimensions, including ion mobility separation and MS2, raw data sizes will increase dramatically..... | 55 |
| Figure 3-5. | Schematic showing the interaction of algorithms among models in which they are used and the facilities (“hardware” above), languages, and libraries that define how they are designed and optimized..... | 104 |

TABLES

| | | |
|------------|---|-----|
| Table 3-1. | Projected HPC Needs for 2025 | 49 |
| Table 3-2. | Advantages and Disadvantages of the Various Programming Models..... | 106 |

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 Accelerated Climate Modeling for Energy (ACME) model, 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.

1 INTRODUCTION

1.1 The DOE Exascale Requirements Reviews Initiative

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 and applied mathematics. Meetings were held for each of DOE's six Office of Science (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 Science (BER) review was held on March 28–31, 2016.
- The Nuclear Physics (NP) review was held in June 2016.
- The Advanced Scientific Computing Research (ASCR) review was held in September 2016.

The overarching goal was to determine the requirements for an exascale ecosystem that includes computation, data analysis, software, workflows, high-performance computing (HPC) services, and other programmatic or technological elements that may be needed to support forefront scientific research.

Each Exascale Requirements Review has resulted in a report prepared by DOE for wide distribution to subject matter experts and stakeholders at DOE's ASCR facilities, including the Argonne and Oak Ridge Leadership Computing Facility centers (ALCF and OLCF, respectively) and the National Energy Research Scientific Computing Center (NERSC); and the Energy Sciences Network (ESnet).

1.1.1 Previous DOE Requirements-Gathering Efforts: "Lead with the Science"

DOE has experienced definite value in implementing its previous requirements-gathering efforts. As noted by Helland (2016), such review meetings have served to:

- Establish requirements, capabilities, and services.
- Enable scientists, programs offices, and the facilities to have the same conversation.
- Provide a solid, fact-based foundation for service and capability investments.
- Address DOE mission goals by ensuring that DOE science is supported effectively.

1.1.2 National Strategic Computing Initiative (NSCI)

The National Strategic Computing Initiative (NSCI) was established by Executive Order on July 30, 2015. Helland (2016) identified the NSCI's following four guiding principles:

1. The United States must deploy and apply new HPC technologies broadly for economic competitiveness and scientific discovery.
2. The United States must foster public-private collaboration, relying on the respective strengths of government, industry, and academia to maximize the benefits of HPC.

3. The United States must adopt a whole-of-government approach that draws upon the strengths of and seeks cooperation among all executive departments and agencies with significant expertise or equities in HPC while also collaborating with industry and academia.
4. The United States must develop a comprehensive technical and scientific approach to transition HPC research on hardware, system software, development tools, and applications efficiently into development and, ultimately, operations.

NSCI's objectives echo plans already under way in DOE's current exascale computing initiatives. In fact, DOE is among the NSCI's three lead agencies (along with the U.S. Department of Defense and the National Science Foundation), which recognizes these agencies' historical roles in pushing the frontiers of HPC and in helping to keep the United States at the forefront of this strategically important field (Helland 2016).

1.2 BER Workshop, Subsequent Report Preparation, and Purposes

DOE SC convened an Exascale Requirements Review for the Biological and Environmental Research (BER) division, which took place on March 28–31, 2016, in Rockville, Maryland, and brought together leading BER researchers and program managers, scientific and HPC experts from the ASCR facilities and scientific computing research areas, and DOE's BER and ASCR staff and ESnet personnel. (See Appendix A for the list of participants.) During the review, participants:

- Identified forefront scientific challenges and opportunities in biological and environmental research that could benefit from exascale computing over the next decade.
- Established the specifics of how and why new HPC capabilities will address issues at various BER frontiers.
- Promoted the exchange of ideas among application scientists, computer scientists, and applied mathematicians to maximize the potential for use of exascale computing to advance discovery in biological and environmental research. (See Appendix B for the meeting agenda.)

Outlines and input from white papers and case studies (Appendices C and D, respectively) authored by the participants and submitted to the BER Organizing Committee chairs in advance of the meeting guided the discussions in general sessions and topical breakouts. Committee members and review participants collaborated at the meeting to identify the grand challenges, priority research directions, and computing requirements for their fields of research — communicating these requirements to the DOE SC offices and ASCR facilities. This report therefore reflects extensive and varied forms of input from many voices in the BER community regarding HPC requirements for BER's world-class initiatives.

The review afforded a rare opportunity for nearly 100 participants to interact and learn about each other's areas of expertise, challenges faced, and the exciting opportunities to be made possible by the exascale computing environment.

1.2.1 Post-Review Involvement of the BER Organizing Committee

Since the March 2016 review, members of the BER Organizing Committee have met regularly via conference call and in-person meetings to continue shepherding the BER Exascale Requirements Review report to completion. This effort — led by committee members Adam Arkin (Lawrence Berkeley National Laboratory), Dave Bader (Lawrence Livermore National Laboratory), and Tjerk Straatsma (Oak Ridge National Laboratory) — has involved collaborating with lead authors on drafts of various sections, circling back to review participants to solicit further input and clarification, and elaborating upon the material, as necessary.

1.2.2 Exascale Requirements Reports Will Meet Multiple Needs

DOE managers will use the Exascale Requirements Review reports to guide investments and budgeting, complete their strategic planning, and respond to inquiries, including specifically in their efforts to:

- Articulate the case for future upgrades to DOE and SC management, the Office of Management and Budget, and Congress.
- Identify emerging hardware and software needs for SC, including for research.
- Develop a strategic roadmap for the facilities based on scientific needs.

BER program managers may also use the reports to inform their work. Although balancing such varied end uses can present challenges, the reports are intended as an information tool that can be used by many stakeholders.

1.3 Report Organization

In the balance of this Exascale Requirements Review, Section 2 provides an overview of the BER vision and grand challenges facing the fields of biological and environmental research. Section 3 addresses key scientific challenges and opportunities, along with the priority and cross-cutting research directions and computing needs and requirements associated with each. Section 4 outlines a path forward for successful collaboration to occur among the DOE's ASCR facilities (i.e., the LCFs, NERSC, and ESnet). References and the acronyms/abbreviations used in the report are listed in Sections 5 and 6, respectively, followed by the appendices.

This page is intentionally left blank.

2 BIOLOGICAL AND ENVIRONMENTAL RESEARCH: VISION AND GRAND CHALLENGES

The Office of Biological and Environmental Research supports fundamental research and scientific user facilities to achieve a predictive understanding of complex biological, climatic, and environmental systems for a secure and sustainable energy future. The program seeks to understand the biological, biogeochemical, and physical principles needed to predict a continuum: from molecular- and genomic-scale processes where reactions occur at rates faster than 10^{-10} seconds, to environmental and Earth system processes at regional and global scales, where changes occur over decades. Outlined in this document are some of the current scientific challenges within BER research programs where exascale computing could have a major impact on BER science.

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 Accelerated Climate Model for Energy (ACME), 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.”

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.

2.1.2 CESD Vision

CESD supports fundamental science and research capabilities that enable major scientific developments in the coupled Earth system, with particular focus on atmospheric, land surface, and subsurface ecosystem process and modeling research in support of DOE's mission goals for basic science, energy, and national security. This support includes research on clouds, aerosols, and the terrestrial carbon cycle; large-scale climate and Earth system modeling; the interdependence of climate effects and ecosystems; and integrated analysis of climate impacts on energy and related infrastructures, with a view toward informing U.S. energy needs. It also supports subsurface biogeochemical research that advances fundamental understanding of coupled physical, chemical, and biological processes controlling both the terrestrial component of the carbon cycle and the generation of methane and carbon dioxide, as well as hydrologic processes that significantly affect the environmental fate and transport of a variety of inorganic and organic contaminants and other energy by-products. This integrated portfolio of research from the molecular level to field scales emphasizes the coupling of multidisciplinary experimentation and advanced computer models and is aimed at developing a predictive, systems-level understanding of the fundamental science associated with climate variability and other energy-related environmental challenges.

CESD continues to advance the science necessary to further develop predictive climate and Earth system models targeting resolution at the regional spatial scale and at intraseasonal to centennial timescales, including focus on areas of critical uncertainty, while also modeling at scales appropriate for various processes and considering how best to couple or transfer knowledge gained and modeled at one scale to another scale or process. Connected with this multiscale challenge is the need to trace predictive uncertainties to process uncertainties and to determine the observations needed to improve process and larger-scale predictive models.

CESD research advances understanding of how the Earth's dynamic, physical, and biogeochemical systems (e.g., the atmosphere, land, oceans, ice, and subsurface) interact, and the research also projects how these will influence future climate effects and environmental change to inform plans for future energy and resource needs. BER's Earth system modeling trends toward ultra-high resolution and variable, regionally refined resolution (ocean, atmosphere, cryosphere, land, and energy/societal elements) coupled climate simulation, in which Earth system processes are increasingly resolved in high-resolution regions. Cloud-resolving, land-process-resolving, and energy-relevant human-systems models inform and may be coupled or even embedded within high-resolution regions of the global model.

To achieve these objectives, BER (in collaboration with ASCR) will be addressing challenges for computational performance, accurate and appropriate component coupling, and climate predictability. This effort requires that advances be made in applied mathematics, computer science, and software to improve coupling methods, time-stepping, and load-balancing for multiphysics and multiscale systems; methods to accelerate system initialization of Earth systems; algorithms for physics and biogeochemistry for high or variable resolution; mathematical methods to determine system predictability; methods to derive statistics and diagnostics during simulation; algorithm designs that map to new architectures and memory structures; and coupled model performance

and portability. To effectively develop next-generation Earth system modeling for new computer architectures, multidisciplinary teams of committed mathematicians and computational and climate scientists will need to work together closely to design appropriate algorithms and code structures.

2.2 BER Grand Challenges

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.

2.2.2 CESD Objectives and Grand Challenges

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. The Division has particular focuses on regions and systems where uncertainties are largest, such as polar clouds, deep convective systems, mesoscale dynamics, cloud feedbacks, aerosol-cloud indirect effects, effects of organic aerosols, permafrost, boreal and tropical ecosystem changes, interactions of local hydrology with climate, and implications for water availability. These challenges are important for reducing the uncertainty in climate sensitivity and variability, as well as assessing the impacts of climate on local regions.

CESD's program activities support these goals, with emphasis on developing and improving the predictive capability of Earth system models; on advancing targeted studies of atmospheric, terrestrial, and subsurface biogeochemical system processes; on using CESD's facilities to achieve unprecedented understanding of Earth's dynamic processes through targeted experimental studies; and on strengthening engagement with energy and environmental stakeholder communities.

In coming years, the CESD goals and capabilities will be oriented toward particular grand challenge research problems that are critical to energy on the one hand, and that require coordination of the multidisciplinary activities in CESD on the other.

One challenge is to advance the understanding of natural and anthropogenic drivers and impacts of changes in the coupled Earth-energy-human system, including improved discernment of regional information for drivers and impacts, and the interplay with energy systems.

A second challenge is to better understand high-latitude processes, feedbacks, and interactions with lower latitudes, with a goal of reducing uncertainty concerning the factors governing the high-latitude changes. High-latitude changes are expected to have important effects on local energy-resource potentials, as well as impacts on lower latitude climate and coastal regions.

A third challenge involves improving understanding of the integrated water cycle by studying relevant processes in atmospheric, terrestrial, oceanic, and human system components and their interactions and feedbacks across local, regional, and global scales, thereby improving hydrological predictability.

A fourth challenge is to advance a predictive understanding of coupled biogeochemical processes and cycles across spatial and temporal scales by investigating natural and anthropogenic interactions and feedbacks, and associated uncertainties, within Earth's climate and environmental systems.

A cross-cutting grand challenge involves the development of conceptual and computational frameworks for data management, analysis, and visualization; model-data integration; and hierarchical model interconnectedness to support integration of models and data across scales and complexity to address CESD grand challenge science.

2.3 Mapping the BER Grand Challenges to the Computing Ecosystem and Exascale Mapping

2.3.1 Mapping BSSD Programmatic Objectives to Computing Ecosystems

BER biological research seeks to fully understand microbial and plant systems relevant to DOE missions and their interdependencies in order to develop efficient systems for producing biofuels and bioproducts and to address biology-based environmental questions. To achieve mission goals, BER researchers must be able to develop predictive, quantitative system biology models. Such models can lead to new insights, be useful in generating and testing hypotheses, guide experiments and aid interpretation, and inspire new theories.

Many different experimental techniques are used in BER biology research, each with different data structures and procedures for obtaining useful information out of the data. These techniques generate large amounts of data, often gigabytes' to terabytes' worth per day: in addition to the genome sequencing and -omics techniques (noted in previous sections), large amounts of data are generated at high speeds for structural biology techniques such as serial crystallography and cryo-electron microscopy (cryo-EM), and volumetric imaging such as tomography or scanning electron microscopy.

2.3.1.1 OMICS

Meta-omic analytic techniques for studying microbial community dynamics, function, and metabolic potential involving sequencing, proteomics, and metabolomics produce terabytes' worth of data per experiment. While these techniques can provide a compositional snapshot of the species, genes, metabolites, and activities of microbial communities, such high-throughput data sets are complex; and thus, making inferences across data sets is challenging because combining biological data sets across modalities and time/space scales for comparative analysis is largely application dependent. In order to answer questions posed by the scientists/users for analyzing these complex data, new tools will be required. In addition to the enormous data analysis challenges, the biggest risk of running biological applications on exascale systems is that the analysis tools we rely on will not port easily to the emerging architectures. All of these problems require scalable computational solutions and machine learning approaches. Thus, the increasing amount of data generated raises important issues for:

- Data storage and interaction with increasing volumes of data.
- New visualization tools that allow scientists to explore multimodal data sets and ask new questions.
- Biological applications that run efficiently on the energy-efficient architectures needed for exascale systems.
- Seamless data sharing modalities for the scientific user community.
- The enabling of the scientific community to run computations that combine DOE data with user-generated data sets on DOE compute resources.

2.3.1.2 Computational and Data Challenges Associated with Integrative Approaches for Modeling Molecular and Cellular Systems

The data for 3D structural models generated by microscopy, crystallography, and imaging techniques yield temporal information. A comprehensive, biologically realistic characterization of a system relevant to bioenergy and the environment integrates meta-omic data as well as 3D and 4D molecular structure data. Bioimaging technologies utilize electron, ion, neutron, optical, and/or X-ray sources to image and probe biological samples across multiple spatial, temporal, and chemical scales.

As the suite of tools used for multimodal or correlative bioimaging expands, so does the complexity of information that must be indexed, registered, and overlaid in a meaningful manner so scientists can holistically interpret the resulting data. Development of user-friendly and broadly accessible exascale bioimaging data exploration platforms is very much needed. KBase has demonstrated the utility of a computational platform for hosting multimodal -omics data in a coherent fashion. However, to identify emergent phenomena in biosystems from bioimaging data sets, new exascale platforms and methods for multimodal and multiscale data hosting, archiving, indexing and registration, visualization, and exploration are needed. These platforms are not intended to simply act as graphical interfaces but should also directly link to or incorporate automated processing schemes and virtual cell/biosystem models that are constantly updated as users add new images and data. The ability to computationally survey all static images, spectra, movies, and models generated for the same or similar samples across scales and with ease of use will permit DOE researchers to access currently hidden or overlooked data and create a unified understanding of molecule-to-mesoscale organization and dynamics that can be exploited to enable rational design of biosystems for superior bioenergy and environmental applications.

Large data frames of data generated at a super-fast rate from the Linear Coherent Light Source (LCLS) or from cryo-EM, as examples, can currently be measured and stored for later processing and analysis. However, faster algorithms and compute power are needed to enable scientists to process that data “on the fly,” with the ultimate goal of visualizing the structure as it emerges from the data that is being processed and analyzed as it is collected. Methods developed for these newer instrument platforms are informing new ways of conducting experiments at more mature facilities such as synchrotron experimental stations, so the high-speed processing, analysis, and visualization algorithms will be needed there as well. Meanwhile, new approaches to interpreting and visualizing large data representations such as volumes from scanning electron microscopy or serial blockface imaging are also needed—specifically, automated algorithmic segmentation of those image volumes. This advance may involve neural network or other machine learning approaches.

As novel technologies are developed, the relevant computational needs increase substantially. New advancements in high-throughput data acquisition, automated processing, efficient image registration/indexing, visualization/exploration, modeling, data hosting/storage/archiving, and compression are required to address challenges facing both functional and emerging bioimaging instrumentation.

Advanced computing systems are needed to collect, process, analyze, and combine these different types of experimental data to achieve meaningful results in BER research. An example of where the utility of an integrated model for bioenergy interests will be especially useful is biomass recalcitrance and developing the pretreatments necessary to generate cellulosic biofuel from lignocellulosic plant matter. Models and simulations of lignocellulose structure have informed the development of pretreatments that can disassemble this composite into its valuable constituent monomers without loss or damage to the sugar monomers released.

A “complete” molecular dynamics (MD) model of cellulose, lignin, and hemicellulose, as well as their interactions with enzymes and chemical and thermal reactions, has yet to be realized. As modeling and simulation methods continue to improve and more studies are carried out, our understanding of the molecular basis for recalcitrance will lead to improved technologies for producing biofuels and bioproducts (Davison et al. 2013).

Detailed molecular models built from integrating many kinds of experimental and theoretical information from other sources can potentially generate mesoscale-level (10- to 100-nm) models of subcellular organelles and systems, viruses, or whole cells. Calculations and computational approaches differ markedly across this expanse of scales, from molecule to cell, represented in both

the spatial and temporal domains. The different levels of approximation, and therefore algorithms, used at different scales will need to be harmonized (see Im et al. 2016, a review that originated from a discussion at the 2014 meeting on Modeling of Protein Interactions).

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.

2.4 Report Roadmap

The preceding discussions represent the proverbial “view from 30,000 feet” of the state of BER science and the grand challenges we are tackling as a community. From such a vantage point, it is exciting to contemplate the discovery science that tackling these grand challenges may lead to as we seek to understand BSSD and CESD complexities across a range of scales.

The next stage is to bring these “bird’s eye” views to ground level and concretize them through greater context and detail. Thus, in Section 3, we expand on four BSSD topic areas and seven CESD areas, highlighting particular challenges that BER scientists are trying to resolve; the computational roadblocks that are inhibiting further breakthroughs; and specific resource, theoretical/intellectual, and people needs that can help us bypass these roadblocks.

Section 4 concentrates specifically on requirements, recommendations, and requests that appear in Section 3. These are grouped into four categories of opportunities for collaboration among domain scientists, academics, computer scientists, and the ASCR facilities and encompass:

- Methods development
- Computational environment and resources
- Data
- Communication and community involvement

These items constitute the requests or “ask” that BER is making at this time so DOE can closely and actively support a path forward for BER science initiatives and thereby promote an evolving computing ecosystem leading to exascale.

This page is intentionally left blank.

3 BER RESEARCH DIRECTIONS AND COMPUTING NEEDS/REQUIREMENTS

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.

3.1.1 Multiscale Biophysical Simulation from Molecules to Cells

3.1.1.1 Scientific Challenges and Opportunities

Obtaining a predictive understanding of complex plant and microbial systems of primary importance in bioenergy and environmental sciences requires the translation of diverse experimental data into time-dependent, *three-dimensional* physical descriptions of dynamic, interacting elements. Success in this endeavor will furnish a rational approach to understanding genotype/phenotype relations and their response to environmental change. The natural means of obtaining physical descriptions is biophysical computer simulation, a broad field that originated in the 1970s with picosecond-timescale, atomic-detail molecular dynamics (MD) simulations of small single proteins (McCammon and Karplus 1980). This field is now expanding to development of models of large ensembles of biological macromolecular systems and even of cell compartments, whole cells, and communities of cells.

The advent of exascale computing would, in principle, allow the MD simulation of a whole cell at atomic detail, that is, simulations of the trajectories through space of every atom in a living cell. Although acquiring this capability appears seductively within the realm of the possible, the likelihood that it becomes feasible in a meaningful way in time for exascale hardware is very low because of two problems. First, our experimental description of a living cell is nowhere near being atomistic in detail, thus precluding the generation of an atomic-detail “starting structure” that is required for MD; and second, the timescales over which MD can be performed are severely limited, extending to only the micro- to milliseconds; and even with exascale computing, the timescales for large, direct, classical atomistic simulations will remain constrained to fractions of a second. Furthermore, the absence of explicit chemical reactions in classical MD precludes simulation of chemical reactions, such as, for example, phosphorylation by kinases.

To overcome the above challenges and derive a physical description of the cell that will be of formidable predictive power will require expansion of the field to incorporate *multiscale* approaches. This expansion is an area of great promise but also presents considerable challenges.

Biophysics was an early pioneering field of multiscale simulation, with the quantum mechanical/molecular mechanical (QM/MM) approach, which led to the 2013 Nobel Prize in Chemistry (Smith and Roux 2013). In QM/MM, areas of an enzyme that undergo chemical reaction are modeled quantum chemically, with the environment treated in a more approximate (molecular mechanics) form. This approach has allowed mechanisms and energetics of enzyme reactions to be determined. However, systems biology goes beyond the consideration of single macromolecules to obtain holistic information about interacting biological systems. Unfortunately, current bioinformatics and biophysical tools examine only narrow windows of spatial, temporal, or chemical information, leading to fragmented information. The challenge, then, is to integrate these tools in a multiscale framework capable of visualizing, conceptualizing, and testing genomic and imaging information with physical models on relevant timescales. Realizing this integration will require the development of the key missing capability — the integration of -omics information and “big” biological data into physically realistic, 3D, time-dependent pictures of the systems involved. This 3D view will require the melding of disparate and voluminous biological data with knowledge of macromolecular structure and dynamics; the positions of subcellular structures and the macromolecules with respect to those expressed structures; and the distribution of smaller species such as solvents, metabolites, and ions. The ability to calculate the functional dynamics of these molecular systems at successively linked system scales will be required. Biophysical multiscaling will expand to include coarse-grain particle simulations at the level of the cell while incorporating data from such disciplines as proteomics, transcriptomics, metabolomics, dynamic flux balance models, imaging, and thermodynamic speciation. The challenge will be to use physics, chemistry, and extreme-scale computation to convert these data into models that can provide predictive understanding

of microbial and plant systems, subsuming the present need for “brute force,” trial-and-error experimental approaches.

Success in systems-level biophysical modeling will have considerable ramifications in energy-related biological and environmental sciences (see Figure 3-1). Biophysical modeling at both the molecular scale and above will be essential in synthetic biology efforts, from structure-based protein engineering to establishing thermodynamic and kinetic descriptions of metabolite, energy, and electron flow. Finally, integration of the methods described here with environmental modeling as described in Section 3.2 of this report will enable predictive understanding of the cycling of carbon, nitrogen, and environmental contaminants in the global Earth system.

3.1.1.2 Priority Research Directions

The “workhorse” codes of computational biophysics are MD and quantum chemistry, and it will be critical that the major MD codes are all successfully ported to exascale architectures to ensure the continued leadership and productivity of DOE BER efforts, as well as those of the broader scientific community (these are also two of the major workhorse codes of materials science, for example). This goal is achievable but will require DOE to “own” this challenge, with BER, ASCR, and DOE’s Office of Basic Energy Sciences (BES) taking cooperative leadership roles, providing needed resources, and collaborating in an unprecedented manner. While MD and QM form the backbone of biophysical simulation, their integration with systems biology will require careful planning of a multiscale toolset that scales up to coarse-graining, compartments, cells, and communities in an integrated and self-consistent fashion and in an exascale-savvy way. The interoperability of 3D cell-scale cellular codes with molecular and mesoscale modeling — a key requirement for multiscale modeling — will need to be addressed.

Methods for simulating different scales at once will be at the frontier of future computational biophysics. Coupling both temporal scales and length scales will be required. Length-scale coupling can be accomplished by: (1) “serial coupling,” in which coarse-grained parameters are derived from finer-level simulations; and (2) “parallel-coupling” multiphysics codes. The development of both of these approaches will be required.

3.1.1.2.1 Molecular Dynamics

The mainstay of molecular-scale biophysics is MD, which now can be routinely applied to 500K-atom systems for timescales over 1 μ s. MD has been a major user of supercomputers for 50 years and is a natural application for extreme-scale supercomputing (Figure 3-2). The advent of exaflop capability will permit lengthening accessible MD time and length scales, thus enabling the discovery of new phenomena involving dynamic, non-equilibrium, and hierarchical assembly with rapid time-to-solution. Weak scaling of MD, in which the system size is scaled, is likely to be practicable at the exascale, leading to system sizes of up to 10^{10} atoms. This scaling will permit the MD simulation of systems of considerable interest to the energy biosciences, such as plant cell walls, microbial membranes, and multiple interacting protein machines. In contrast, strong scaling, in which longer timescales are accessible to systems of a specific size, is highly nontrivial owing to the complexity and heterogeneity of post-Moore’s Law extreme-scale architectures. Instead of simple porting or relatively uncomplicated vectorization, extensive effort is now needed to adapt MD codes to heterogeneous architectures. Furthermore, this effort will require not only engaging simply in “software development” or “coding” but rather in the co-design of hardware and software driven by the conceptualization of application-specific simulation methodologies that seek to find the optimal machine-dependent, time-to-solution of a given physical problem. Optimizing codes for maximum single-node performance while striving for good, strong scaling presents one of the most significant software and algorithm problems facing the MD community.

Four MD codes — AMBER, CHARMM, GROMACS, and NAMD — are commonly used in the community, and each of these codes has thousands of users. To ensure the greatest impact of exascale computing on MD-enabled science, *all four* of these codes must be implemented and optimized on each of the new architectures. This requirement will not represent a duplication of

Example system: Lignocellulosic biomass. The “grand challenge” class of problem of understanding the capture of energy and carbon dioxide (CO₂) and its cycling through the environment, including through biofuel/bioproduct pathways, will require obtaining a detailed understanding of the plant cell wall and all aspects of its generation and degradation. This understanding includes the deconstruction of natural cell walls, the ability to co-design less recalcitrant plants, and the methods and organisms capable of most efficiently converting them to fuels and renewable products. There is considerable interest in genetic modification and natural variation that render plant cell walls more amenable to deconstruction. The challenge here is to understand how genotypic changes in plants are translated into changes in the physical properties of the cell wall. This is essentially a physicochemical problem, driven by the complexity of the plant cell-wall hierarchical structure (Figure 3-1). To understand both the organization and deconstruction of this complex composite material, the associations and morphology of the component polymers and the enzymes acting on them require characterization across multiple length scales ranging from nanometers to micrometers, up to and including cell communities and the plant:microbe interface.



Figure 3-1. 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: Thomas Spletstoesser, scistyle.com).

effort, as each code has different strengths and weaknesses. The realization of this resource need (which has parallels in other science fields) is critical to the success of the overall exascale science endeavor. The MD codes need to be optimized for the pre-exascale machines Summit, Aurora, and Cori, with a sustainable path toward exascale computing. This work will make exascale supercomputing accessible to a large computational science community.

However, simply porting existing codes will not allow MD simulations to take full advantage of the new exascale architectures; rather, methodological development is also required. Methods for overcoming the treatment of long-range electrostatics, a major hurdle in extreme-scale MD scaling, should be explored and implemented, resulting in an extension of present MD capabilities to the point of reaching simultaneously the μm length scale and μs timescale. The standard Particle-Mesh-Ewald (PME) method relies on Fast Fourier Transforms (FFTs), which require expensive global communications; alternatives should be implemented. Short-range electrostatics will also need to be improved, including the use of polarizable force fields optimized for the expected hardware.

Of particular importance to biology will be the development, implementation, and testing of *ensemble methods*, in which extreme-scale supercomputers are efficiently partitioned to increase accessible sampling times for smaller MD systems well beyond the μs . Ensemble methods lessen the requirement for fast communications given that the communication between ensembles is infrequent. For example, metastable states, rather than atoms, can be placed on nodes, which can then be connected with Markov-like networks. Architecture-specific replica exchange, possibly using multiple dimensions such as temperature and pH, and path sampling will also be required.

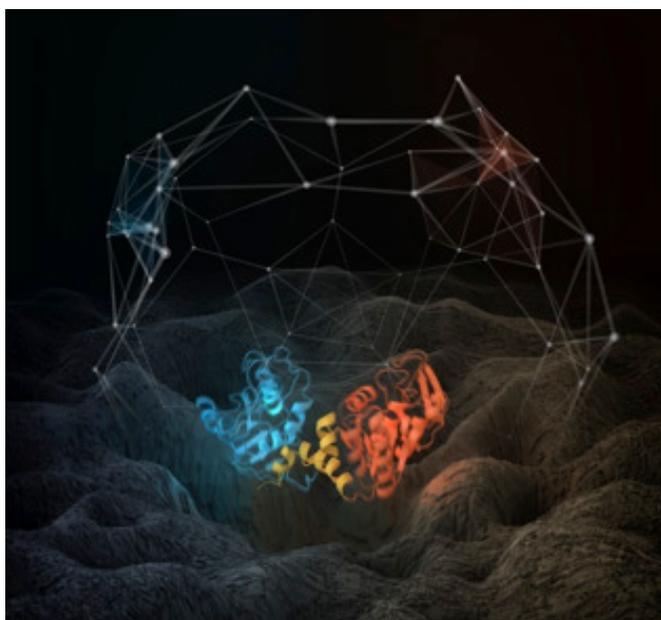


Figure 3-2. Multi-timescale simulation of a single protein. A phosphoglycerate kinase protein was subjected to MD simulations on various supercomputing architectures. The relative motions of the red and blue domains of the proteins are highly complex and can be described in terms of the motion of a configurational point on a rough energy landscape (illustrated). The transitions of the structure between energy minima on the landscape can be described in terms of a network (illustrated), which was found to be fractal (self-similar) over 13 decades of time (Image credit: Thomas Splettstoesser, scistyle.com).

3.1.1.2.2 Quantum Chemistry

QM methods will remain critical tools for BER research for two primary reasons: (1) QM is required to provide parameters for other simulation methods like molecular dynamics; and (2) chemical reactions, such as occur in enzymes, and electron transfer processes can be treated only by using quantum methods. Combined QM/MM methods are particularly attractive for the latter purpose. Techniques that combine quantum and molecular mechanics to compute reaction mechanisms and associated energetics need to be improved and their development accelerated to provide missing information in metabolic pathway maps and to guide rational engineering of enzymes.

Successfully porting QM codes to exascale computing systems would enable entirely new applications, such as modeling data from time-dependent crystallography and predicting enzyme functions, and would support BER missions by examining critical biochemistry, including lignin chemistry, transferase mechanisms, contaminant metalloenzyme action, and engineered enzymes.

Despite this potential, supercomputing has rarely been used to apply quantum chemical methods to biological systems owing to challenges of scaling codes to HPC. Thus far, larger systems (e.g., 1,000 atoms) and higher levels of theory (e.g., coupled cluster) scale better than smaller systems (e.g., 100 atoms) and lower levels of theory (e.g., density functional theory [DFT] or Hartree-Fock). At the petascale, this challenge needs to be addressed through several strategies, including the ability to use both graphics processing units (GPUs) and many integrated cores (MICs). Priorities include further improving the efficiency of existing QM algorithms and optimizing the total time-to-solution. Code improvements will exploit the inherent sparseness of Hamiltonian representations using local formulations. Multiscaling will also be needed *within* the QM codes, that is, “QM/QM” methods will use a hierarchy of increasingly complex QM methods that will be talking to each other.

3.1.1.2.3 Coarse-Grained Simulations

Cellular events on the millisecond timescale or longer and system sizes beyond 100M atoms call for simulation methods more simplified than atomistic MD, averaging out the unimportant degrees of freedom to preserve long time- and length-scale properties. This “coarse-graining” (CG) can lead to a multiphysics description of biological phenomena. The challenge is to filter phenomena on short time and length scales that have mesoscopic consequences so that important data are preserved in coarse graining, all while maintaining self-consistency. Coarse-grained, particle-based methods exist and will be of increasing value in cell modeling, permitting cell-scale simulations on timescales of up to one second.

Brownian Dynamics (BD) CG simulations are becoming increasingly important for modeling “mesoscale” structures. BD simulations model molecular diffusion using the theory of Brownian motion and have been traditionally used to model intermolecular binding. New applications will include complex cellular environments. As with atomic-scale MD, a primary challenge for BD parallelization is the treatment of long-range interactions, which include electrostatics and hydrodynamic interactions; thus, significant effort will also be required to adapt BD codes to massively parallel and heterogeneous architectures so as to fully unleash the power of BD toward larger biological systems such as signaling, membrane complexes, organelles, and whole cells.

With the cell-level time and length scales accessible to CG simulation, tracing the diffusion of macromolecules and metabolites across the cell is feasible, including in the crowded cellular environment, providing information on system-dependent diffusion constants and associations between multiple molecules in the cytoplasm and at membranes. Moreover, CG methods scale efficiently on a variety of supercomputers. Integration of CG with systems biology data and scaling on supercomputers are relatively unexplored pursuits yet hold tremendous promise for whole-cell modeling.

3.1.1.2.4 Cell-Level Simulations

An emerging challenge for computational biophysics is modeling in 3D the structures and dynamics of large macromolecular complexes and machines, cell membranes/walls, organelles, entire cells, and communities of cells. Such efforts demand the use of data-driven, multiscale approaches. A viewpoint on this development appears in Mapping Sequence to Models in Section 3.1.2.

To simulate systems of biochemical reactions inside a cell, handling spatial heterogeneity and efficiently simulating timescales on the order of the cell cycle (minutes to hours) require the

integration of CG methods (above) with the development of stochastic modeling techniques, such as the reaction-diffusion approach. These techniques will allow characterization of complex events such as signaling cascades, transcription, translation and degradation, biofilm formation, and cell division.

One of the most pressing needs is to develop new methods to model reactions in cells that include both reaction dynamics and energetics. For example, in synthetic biology it is critical to determine the thermodynamic feasibility of engineered pathways and molecular products before embarking on an expensive experimental campaign. In addition, the mathematics underlying the simulation approach should ideally be equally useful for data analysis — integrating experimental data in a physically principled manner, such as direct integration of metabolite concentrations from metabolomics assays into thermodynamic or kinetic propensities (likelihoods) that can be used for simulations. Moreover, because of feedback and other issues, optimization and simulation solvers should be capable of handling nonlinear problems — as nonlinearity is the rule and not the exception in biological systems.

Cell-level simulation must be closely linked with experimental efforts such as the imaging of synthetic biology. Once a molecular-level 3D model of the cell is established, the evolution of this system will need to be followed over time, merging computation with experimental multiresolution imaging capabilities that can both identify individual macromolecular and small-molecule species in the cell and follow the distribution of these components with a high level of temporal resolution.

3.1.1.3 Cross-Cutting Research Directions

3.1.1.3.1 *Biophysics and Genomics*

Biology originates from the molecular scale, and rational engineering of biological systems in the energy and environmental biosciences will need to drill down to this level. Likewise, knowledge of thermodynamics and energetics at the single-protein level is critical to understanding cellular function. Determining the functions of all gene products is a major challenge of the post-genome era. Current estimates are that functions cannot be confidently assigned to 30-40% of putative proteins encoded by a typical genome. This challenge can only be addressed by deploying multiple, complementary approaches that include those described in Section 3.1.2, Sequence to Models, and Section 3.1.4, Big Data. Historically, the sequence-to-structure problem (i.e., protein folding) has been a central concern of computational biophysics. Although much progress has been made in recent years, methods to predict protein structures reliably (in the absence of a close structural homolog) need to be developed, as well as protein:protein and protein:ligand associations. Success in this endeavor will lead to much-improved gene annotation accuracy. Computational methods are also needed that enable reliable understanding of the mechanisms of enzyme functions and complex, multimolecular machines. The ability to compute binding free energies and reaction barriers using methods that include molecular dynamics and quantum mechanics (see Sections 3.1.1.2.1 and 3.1.1.2.2) will also play a central role in protein structure-function inferences. Thus, the computational power available in the future will need to be fully leveraged to understand folding, allostery, binding, and reaction, which, in turn, will permit deeper insights into and derivation of principles of biological functions, from the molecular to whole-organismal level. Exascale computing can in principle enable application of the above strategies on a genome scale.

3.1.1.3.2 *Supercomputing and Large Experimental Facilities*

DOE is ideally equipped to support molecular-scale research at its large-scale facilities for next-generation synchrotron radiation and neutron scattering. Using these facilities optimally for molecular-scale research will be increasingly valuable and will require integration with HPC.

For complex systems, interpretation of both the X-ray and neutron experimental results is greatly enhanced by detailed data from simulations (Figure 3-3). Combining the experimental data with high-level quantum mechanics and molecular mechanics computations promises to provide unprecedented detail on enzyme function and structures of large molecules and complexes in solution. To enable this advance, a software interface between neutron and X-ray biological science and extreme-scale MD supercomputing is needed that involves optimally calculating the scattering functions directly from the MD trajectories, thereby enabling the deployment of neutron/X-ray experimental and exascale simulation as a single, fully integrated research tool.

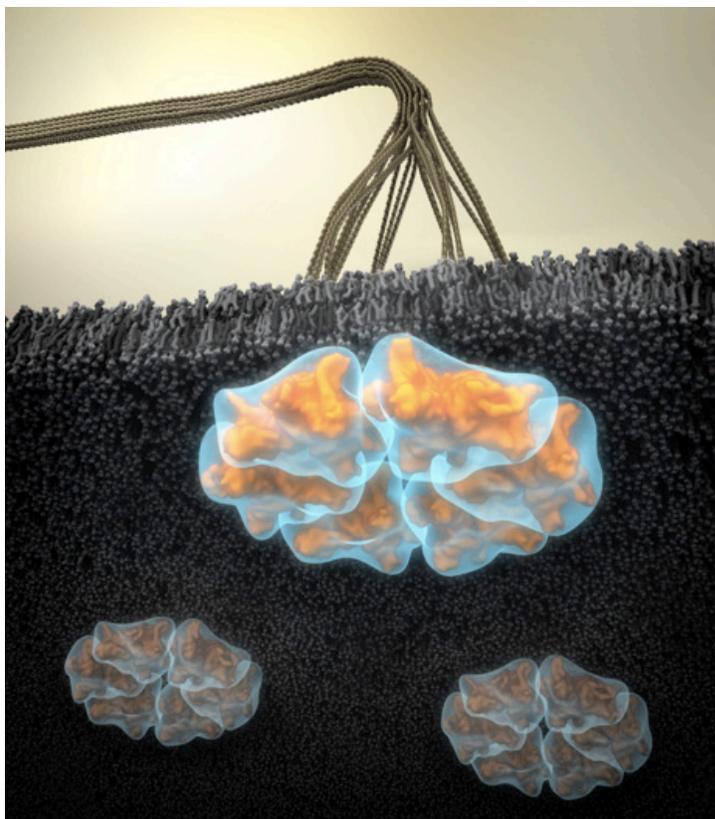


Figure 3-3. Model of the cellulase synthase enzyme, Cesa, derived by integrating neutron scattering and HPC (Image credit: Thomas Splettstoesser, scistyle.com).

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.

3.1.1.4.3 Cellular-Scale Simulations

A wide variety of cell-scale models can be envisaged and thus a wide variety of software and hardware requirements. Information on details of the subcellular structure will need to be held in memory as well as the state of the cell, consisting of all the metabolite, protein, and possible nucleic acid concentrations modeled by the simulation technology and broken out by each subcellular compartment. Communication requirements will depend on the level at which reactions are described and modeled, whether diffusion is deterministic vs. stochastic, and how the simulation domains are broken out by substructures. If there is a large difference in the timescales between reaction events and diffusion/advection in the cytoplasm, then there may be the opportunity to hide some of the computational latency. It is likely that the number of floating point or integer operations per second will not be the limiting factor for serial multiscale reaction-diffusion types of simulations. However, for more detailed whole-cell dynamical simulations, FLOPS may become of critical importance.

3.1.2 Mapping Sequence to Models

3.1.2.1 Scientific Challenges and Opportunities

The Sequence to Models breakout session encompassed a range of computational biology activities that collectively all aim to build dynamic and predictive whole-cell models. Identified challenges, although diverse, span a scale/information axis that consists of three essential and semi-insulated phases of research, including: (1) understanding the catalog of genes, proteins, and metabolites in a given organism, population, or community and how they evolve; (2) annotating and characterizing these parts with respect to function in the cell-system; and (3) building models of whole cells, including efforts to model signaling and regulatory networks with biophysically detailed network models.

Identifying the content of genomes is complicated by horizontal gene transfer, for microbes, and by exchange and variation in genes in eukaryotes resulting from sexual reproduction and mutation, respectively. Thus, assembly and annotation of gene content remains a highly relevant and active area of research. In several cases, participants identified areas where the increasing pace of genomic and metagenomic data generation will push these problems into the exascale regime. New work to quantify and model the content and evolution of populations was also discussed, involving the need for new algorithms, new biological thinking, and exascale computing. With outbred populations or microbial communities, even the “parts list” can be a moving target that is dependent on assumptions, data density, and biological models.

Multiple participants identified the annotation of gene function as a key problem in large-scale biology. Examples that support the need for improved protein function prediction systems include results from recent efforts to make synthetic organisms with reduced genomes. Even in these reduced genomes, we find that approximately one-third of genes have unknown functions. The same is true when we examine bacterial metagenomes or complete genome collections from outbred populations (humans, Poplar, switchgrass, etc.). Two overall types of effort aimed at obtaining better protein/gene characterization include (1) data-intensive integrative methods for protein function, and (2) compute-intensive methods aimed at characterizing protein structure and developing detailed models of biomolecules (also see Section 3.1.1, Multiscale Biophysical Simulation).

Last, the group discussed motivations and opportunities for building comprehensive models of gene regulation, cell function, and even integrated models of complete microbial cells. Current models are limited to abstracted representations of conscribed subsystems of eukaryotic cells and smaller microbial cells, and thus, scale is subexascale; however, participants discussed scaling both network and whole cell models to less-abstracted representations of multicellular organisms and microbial communities — efforts that will clearly push these systems’ biology efforts into the exascale regime.

3.1.2.2 Priority Research Directions

Three research threads underpin initiatives to advance biological understanding from the molecular to cellular levels:

- Genomics, including genome and metagenome assembly, and structural annotation of genes and operons. This thread includes homology searches and evolutionary analyses.
- Protein structure and function prediction. As the completeness of our structural annotation increases, these tasks dovetail into protein design tasks and the modeling of larger complexes.
- Network reconstruction and modeling, cell-level simulation, modeling, and prediction. As this field advances, these efforts will increasingly connect with synthetic biology and the interpretation of genetic variation in wild populations.

These activities integrate data from genomics, transcriptomics, proteomics, and other high-throughput data produced by researchers from DOE BER and elsewhere. Researchers interface with each other through shared long-term goals and by transferring basic information from one level to the next, and the activities can be carried out independently as part of modular research efforts. The ultimate goal of these activities is to produce whole-cell models that can be used to support modeling efforts and synthetic biology design efforts.

3.1.2.2.1 Genome and Metagenome Assembly and Analysis

DNA sequence data sets are growing exponentially owing to rapid improvements in sequencing technologies, a development driven in part by biomedical applications but equally applicable to problems of energy and environment. Mature sequencing technologies are limited to reading short DNA fragments (<250 base pairs), far shorter than the typical microbial genome (~5 megabases) or plant chromosome (100–1,000 megabases). Counterbalancing the limited length of sequence reads is massive parallelization that allows hundreds of millions of reads to be sequenced at once. The genome assembly problem bridges the length scales from sequence “reads” to chromosomes and genome by recognizing and utilizing overlaps among short sequence fragments to reconstruct longer contiguous sequences. Particularly challenging are large plant genomes, which feature heterozygosity, repetitiveness, and polyploidy, and metagenomes, which include heterogeneous mixtures of hundreds or thousands of different microbial species and their interstrain variation.

The DOE Joint Genome Institute, for example, currently generates >100 terabases of raw genome sequence per year. We project that by 2020, the complete genome sequences will be generated for hundreds of thousands of microbes, tens of thousands of accessions of key plant biofuel feedstocks, and hundreds of thousands of environmental metagenomes. Efforts to organize and access this data while preserving provenance and metadata present significant challenges, especially if the data are produced by multiple providers. While genome and metagenome sequencing is currently carried out as discrete projects, we envision that in the near future, continuous or near-continuous streams of DNA sequence data from environmental and field studies will be brought online, requiring new ways of thinking about assembly and DNA data processing.

In parallel, new sequencing technologies are maturing that read longer DNA sequences (10–100 kilobase pairs) from single molecules. The reliance on single molecules puts a fundamental limit on the accuracy of the resulting long sequences, and it is unlikely that error rates per read will be reduced below ~10%. Long, single-molecule data are currently more expensive per base and have far lower throughput relative to the mature short-read sequencing. However, we anticipate that over the next 5–10 years, throughput and cost will improve. It remains to be seen whether the assembly problem will be solved by incorporating long error-prone reads with short reads, or will be performed entirely with long reads, but it will certainly evolve to leverage this new data type as part of mixed-sequencing platform data sets.

Metagenome Assembly

Additional computational challenges are presented when assembly of multiple genomes from environmental samples is considered (a key DOE BER activity). In this situation, pipelines need to take account of the larger complexity and corresponding variable sequence coverage, depending on the environmental abundance of different taxa and their strains. In spite of these challenges, assembly is aided by new experimental types of data that can provide scaffolding constraints, such as long-range, mate-pair-like data (Putnam et al. 2016). Additional information can be brought to bear to group the resulting assemblies based on intrinsic sequence properties (e.g., base composition, word frequencies), with similarity to an ever-growing database of isolate genomes, and even comparisons with other microbiome samples (e.g., different time points or spatial samples from a given microbial community). Thus, we can expect large-scale genome (significantly larger than mammalian and plant genomes) assembly tasks from microbiology and environmental

metagenomics efforts. While genome and metagenome assembly have traditionally been performed with large shared memory nodes (e.g., 2 TB RAM), recent developments (Georganas 2016) have shown excellent scaling with distributed memory algorithms.

Protein Structure and Function for Proteome Annotation

As we build toward having genomes for nearly every biological organism of interest, we increase the need for automated pipelines for gene annotation. Herein we focus on proteome annotation.

Protein Structure Prediction

Knowing a protein's structure can give us insight into the function of that protein, as often protein structural patterns are conserved across much greater evolutionary distances than sequence patterns. Computational approaches to structure-function relationships are key to making further progress, as single-protein structures often cost upwards of \$250,000–500,000. It is also important to realize that the technologies that enable increases in sequencing throughput are outpacing the technological developments we have seen in the field of structural biology (although that may be changing as cryo-EM technologies advance rapidly). Several participants described advances in both protein structure prediction and its use in predicting protein function. These methods are typically very CPU intensive and very light on storage and I/O requirements. Thus, many structure prediction tasks can be run with very low priority on HPC facilities, along with more I/O-intensive applications.

Functional Annotation

State-of-the-art methods for protein function prediction integrate many data types in attempts to predict function for proteins and (for newer methods) protein regions. Much of the work in this field has been driven by many simultaneous advancements: (1) blind tests of methods like critical assessment of protein function annotation (CAFA) algorithms and mouseFunc, (2) the increase in the availability of genomics data for key species, and (3) the massive increase in sequences of related genomes and proteins for nearly any species of interest. Top-performing methods typically integrate upwards of 20 distinct data types using machine learning methods to predict functions for all proteins in a given species. The aim of these methods is self-consistent transfer from reliable functional annotations; this objective is complicated by the diversity of function types (large label space) and data types (as well as their quality and context dependence). Thus, functional annotation requires careful attention to data provenance, quality control, and a “virtuous cycle” of data cleaning and mapping. Truly integrative methods are typically run on one or a few species at a time to restrict the scale of the learning task, but newer methods (that include advances from the machine learning and transfer learning fields) could benefit from carrying out function prediction of a much larger, perhaps all-organism, scale. This effort would require much larger I/O, active memory, and storage than are typically available, as well as HPC and parallelization expertise to carry out all-organism function prediction calculation. Functional annotation of both protein- and ribonucleic acid (RNA)-coding genes is a true multiple-genomewide computational task and will continue to require the large computing facilities. Current algorithms typically analyze each protein family independently and therefore are easily parallelized. We foresee that new, more powerful methods will depend on covariance between protein sequences, requiring more integrative computation whose parallelization will be more challenging.

3.1.2.2 Network Analysis and Integrative Inference

Network reconstruction and downstream modeling, cell-level simulation, and prediction of cellular dynamics lie at the apex of BER-relevant activities, as these models integrate and rely on all of the methods outlined above, including primary annotation of genomes and basic bioinformatics associated with genomics and proteome annotation (for identifying functional consequences of regulation and completing our set of regulatory factors prior to network modeling or reconstruction). There are several computational activities associated with learning and modeling

networks that span a wide range of biological activities. Specifically, network inference is a key part of several computational pipelines associated with large-scale genomics efforts aimed at many species. Thus, defining the objective and the associated level of biophysical detail needed is a first step in stratifying network inference and modeling techniques. All participants agreed that increasing the overall level of biophysical detail used in network models was a fieldwide goal. The desired increase in the level of biophysical detail includes (but is not limited to) efforts to model cell dynamics, connect metabolic and regulatory networks, and model the regulatory combinatorics that result from protein-protein interactions.

Learning Networks

Several examples of large-scale network modeling and its uses were presented that spanned bacterial, animal, and plant species. As genomic technologies progress, new computational methods are keeping pace, and more accurate and comprehensive models of chromatin structure regulation, transcriptional regulation, and even the regulation of latent regulatory factor activity across cell conditions are within reach. Network inference algorithms, broadly, must carry out multiple subtasks, including: (1) estimating the latent activity of regulatory factors (as proteins and RNA modulate activity as a function of a large number of environmental and cell inputs); (2) deriving and using different data types to develop informative priors on network structure; (3) selecting the correct model via the use of global constraints on network sparsity and structure; and (4) connecting the resulting network models to downstream tasks such as interpreting genetic variation, visualization, synthetic biology, etc. These network inference tasks have not typically presented as difficult HPC tasks, as they are often broken down into easily parallelizable subtasks; however, newer methods will require greater HPC support and larger, integrated learning. Computing does currently place limits on this work; for example, there are many metrics by which we judge network models that we cannot use during learning because of computational expense (such as the stability of the network in detailed dynamic simulations of cell state). As the biophysical detail of network models increases (supported by both better methods and better data), we will see the computational demands of network inference tasks cross into the exascale regime.

Multilevel Network Modeling and Whole-Cell Modeling

Several of the participants in this breakout described whole-organism and whole-cell modeling efforts. Ongoing efforts were discussed where modeling was integrated with efforts to complete the annotation of the genome or the completeness of the network of known interactions (learning plus modeling). Other efforts were discussed that focused on the integration of network models across multiple layers (e.g., transcriptomic, proteomic, metabolomics). The most impressive results from these whole-cell modeling efforts were primarily limited to model bacterial species where the system scale is small and the quality of the “known” networks (or regulatory and metabolomics levels) is quite high. Although the prior successes of whole-cell modeling have not typically required large-scale computing (beyond a typical academic cluster), several research directions are likely going to require significantly more computational resources for whole-cell modeling. Scaling to larger systems, including host-pathogen, multicellular organisms, and bacterial communities, will require significantly more compute power. Further, the current models are generally at a fairly high level of abstraction — largely homogeneous and deterministic differential equations. As it becomes more important to model the discrete, stochastic and spatial nature of cellular chemistry (and, for that matter, microbial populations), the scale of cost for maintaining and running such models (which may require advanced Monte Carlo methods) and parameterizing and testing them effectively with data will become far more costly. See Section 3.1.1, Multiscale Biophysical Simulation from Molecules to Cells, for a more detailed description of three-dimensional cell modeling on HPC.

3.1.2.3 Cross-Cutting Research Directions

Please see our response to Section 3.1.2.2.

3.1.2.4 Computing Needs and Requirements

Computing needs and requirements have been identified in Sections 3.1.2.1 and 3.1.2.2.

3.1.3 Microbes to the Environment

3.1.3.1 Scientific Challenges and Opportunities

This section seeks to identify the computational requirements that bridge the programmatic interests and strategic directions of the two divisions within BER. The Biological Systems Science Division seeks to obtain a predictive systems-level understanding of complex biological systems to advance DOE missions in energy and the environment. By integrating genome science with advanced computational and experimental approaches, BSSD seeks to gain a predictive understanding of living systems, from microbes and microbial communities to plants and other whole organisms. The Climate and Environmental Sciences Division focuses on advancing a robust predictive understanding of Earth’s climate and environmental systems to inform the development of sustainable solutions to the Nation’s energy and environmental challenges. This pursuit includes the specific goal of synthesizing new process knowledge and innovative computational methods that advance next-generation, integrated models of the human-Earth system.

While these stated objectives are highly complementary, they require that we address an overarching grand challenge: how to integrate information across the physical and temporal scales that span microbes, microbial and plant communities (BSSD), and ecosystem functions (CESD) into quantitative predictive simulations of integrated complex systems (i.e., “microbes to the environment”).

Two critical scientific challenges frame the need for a multiscale predictive simulation capability that integrates BSSD and CESD science:

1. *Feedbacks between the Earth’s land and atmosphere.* There are many important but currently poorly understood feedbacks between long-term global atmospheric processes and localized, shorter-term land surface/subsurface processes that are mediated or altered by microbes, plants, and human activity. Climate forcings, including water availability (e.g., impacts of drought), temperature, and atmospheric gas concentrations, are currently resolved by Earth system models at scales of tens to hundreds of kilometers. These forcings, mediated by local fluid flow and material transport (advection and diffusion) processes operating over a wide range of scales, interact in complex ways with plants, microbes, and other organisms (including humans), which themselves are influenced by heterogeneous distributions of mineral phases and organic matter. These larger-scale dynamical processes and interactions create highly heterogeneous and dynamic conditions at the cellular level (on the order of tens of micrometers on a mineral surface or within a single pore of a porous medium). These smaller-scale conditions drive the temporal regulation of microbial metabolic processes as overprinted by genomic controls, giving rise to the dynamic expression of diverse functions of individual microbes and communities. The growth and physiological behaviors of microorganisms, in turn, alter their local environments and affect interactions with other microbes and nearby plants, as well as organic and inorganic compounds. These interactions influence local flow and transport processes, thereby driving fluxes of critical constituents (e.g., water, carbon dioxide, methane, volatile organics, and aerosols), which then stimulate larger-scale atmospheric processes at local and regional scales and eventually at global scales. These feedbacks are overwhelmingly complex and critical to understanding the impacts of human activity on future climate conditions, as well as the impacts of future climate conditions on the integrated human-Earth system. Transformative developments in predictive capability are needed not only to understand impacts at both ends of the scale spectrum (global climate and microbial ecosystems) but to predict impacts on human and natural systems at intermediate scales (i.e., water, energy, and the environment) that are central to DOE missions.

- Fate and transformation of natural organic matter in the terrestrial environment.* As a focused subset of the broad feedbacks described above, the current and future directions of BER's research programs indicate that it is particularly essential to improve understanding and prediction of the dynamics of organic carbon and inorganic nitrogen and phosphorus fluxes between natural, managed, and disturbed land surfaces and the atmosphere, as mediated by microorganisms and plants, and their interactions with inorganic mineral surfaces in a dynamic hydrological environment. The net flux of greenhouse gases from the land surface and to the Earth's atmosphere is estimated to be on the same order of magnitude as anthropogenic fluxes but represents the difference between two much larger biologically mediated fluxes: (1) carbon dioxide uptake by plants and subsequent sequestration in soils, and (2) microbial respiration of organic matter in soils and release to the atmosphere primarily as carbon dioxide and methane. Because of uncertainties in the prediction of both major fluxes at large/global scale, the net effect on future climate conditions is a major source of uncertainty in global Earth system models (e.g., Friedlingstein et al. 2014). Reduction of this uncertainty requires transition from the empirically based descriptive models currently used at system scale to mechanistically based predictive models that integrate complex elements of systems. Because carbon fluxes are intimately tied to biological processes embodied in plants and microbes and their interactions (rhizosphere), the development of an advanced simulation capability integrating models of plants, microbes, and soil is essential to achieving this objective. This scientific challenge is explored in detail in the recent DOE workshop report "Building Virtual Ecosystems" (DOE 2015a).

3.1.3.2 Priority Research Directions

Four priority research directions (PRDs) were identified by the breakout session participants:

- Integration of molecular-scale information into models of microbial metabolism at cellular and community scales (see Cannon et al. white paper).* Simulating the response of individual microorganisms and communities of microorganisms in the context of heterogeneous and dynamic environments presents a critical opportunity and challenge for extreme-scale computing. Predictive simulation capability requires representation of complex metabolic pathways and regulation mechanisms to properly account for microbial responses to environmental changes, as well as to accurately represent the fluid flow and material transport mechanisms that control environmental conditions at the scale experienced by microorganisms. Acquiring this capability will require making significant advances in pore-scale modeling capabilities integrated with microbial growth and metabolism models. Metabolic simulations from first principles may be achievable with exascale computing but are both computation- and data-intensive. Integration with two particular types of molecular-scale information is critical: (1) multi-omics data, and (2) experimental and simulation-derived dynamic and thermodynamics information. Metabolic models should be flexible to scale from a few lumped reactions that summarize cell functionality to a full metabolic model with a complex reaction network, with integration of -omics data for model parameterization and validation. Improvements are needed to the quantitative quality of multi-omics data (or improved algorithms to support inference of quantitative model parameters from qualitative observations; see Wieder et al. white paper) to enable fuller utilization. Thermodynamics control reaction pathways and rates and can potentially be computed using computational chemistry methods. The complexity of microbial communities of interest (e.g., soil microorganisms) is very great and requires advanced computational techniques, as well as theoretical developments for significant advancement. The proposed model systems would represent organisms and their environment at relevant scales (submicrometer to pore scale), and thus the models would best be applied for development of fundamental understanding and to identify emergent behaviors that could be incorporated into larger-scale simulations.

- 2. *Molecular-scale models of natural organic matter and interactions with mineral surface linked to molecular characterization (see Andersen et al. white paper).*** Paralleling the first PRD that focuses on biological simulation is the simulation of submicrometer to pore-scale abiotic reaction processes that control the fate of natural organic matter (NOM) and its exposure to potential degradation by microorganisms. Primary among these are interactions of natural organic matter and biological molecules (e.g., enzymes) with mineral surfaces and solvents that can lead to either stabilization and/or decomposition of NOM. Simulation and modeling protocols — from quantum mechanics to classical atomistic and coarse-grained simulations describing the electronic-to-nanoscale regime of complex Earth systems — are proposed to provide parameterizations pertinent to reactive transport in pore- to global-scale modeling. The supramolecular structure of NOM is complex and largely unknown, with intra- and intermolecular interactions in self assemblies, including inorganic ionic species (e.g., metal cations $\text{Fe}^{2+/3+}$, $\text{Mn}^{2+/3+/4+}$, Cu^{+2+} , Ca^{2+} , Al^{3+} , and anions such as phosphate) and interactions with mineral surfaces (e.g., clays, metal oxide/hydroxides). The aggregation of NOM in supramolecular self-assembled structures and on mineral surfaces is hypothesized to protect NOM, including biogenic organic matter (e.g., microbial exo-enzymes), from decomposition by biotic and abiotic processes. However, transition metal-containing mineral surfaces may facilitate abiotic decomposition of NOM through catalysis. Structural characterization of NOM through computational spectroscopy (Mössbauer, nuclear magnetic resonance [NMR], electron paramagnetic resonance [EPR], and X-ray absorption spectroscopy [XAS]) is also an important part of developing models for further computation of reactive transport parameters.
- 3. *Microbially mediated biogeochemical cycling from pore to watershed scales (see Steefel et al. white paper).*** Watershed processes represent a critical nexus between human activities (land use, water use, energy production) and natural ecosystems. As geochemical and hydrologic conditions change in a watershed, dramatic shifts in microbial community composition are often observed as system perturbations overlay effects on complex interactions with plants and the atmosphere to drive selection from a diverse microbial “seed bank.” These perturbations and associated ecological shifts occur at a range of physical scales from that of an individual pore in a soil (submillimeter) to that of a watershed (many kilometers). It is well recognized that microorganisms within a community embody a range of functional traits that, when expressed, work together to complete important biogeochemical cycles (e.g., nitrate reduction). Gaining the ability to understand and predict the mediation of these complex biogeochemical reaction networks by dynamically changing microbial communities over this broad range of scales is a major scientific and computational challenge. The overwhelming microbial taxonomic diversity can be addressed through definition of functional guilds, each of which has a unique combination of traits, defined based on metagenomic data, that govern fitness under dynamic environmental conditions. The computational model simulates the thermodynamics of coupled electron donor and acceptor reactions to predict the energy available for cellular maintenance, respiration, biomass development, and enzyme production. In addition, the model allows for a faithful representation of the functional diversity of microbial populations, how microbial physiological traits affect fitness, how biogeochemical processes are affected by emerging microbial composition, and how biogeochemistry feeds back to alter microbial fitness and community assembly. Incorporating biologically relevant microbial reaction networks into fluid flow and reactive transport modeling frameworks applied to heterogeneous environments from pore to watershed scales represents a major exascale computational challenge. It requires distributing a dense set of ordinary differential equations (ODEs, the reaction network) across a heterogeneous domain and linking to a set of partial differential equations (PDEs) for flow and transport at multiple scales, resulting in systems of perhaps tens of billions of degrees of freedom. Incorporation of complex data (e.g., large genomic databases, high-resolution images of soil pore structure, remotely sensed images of watershed topography and distributions of

plant functional types and density) is also challenging. Therefore, this PRD encompasses both computationally intensive and big data analysis issues.

4. *Mechanistic simulation of integrated plant-soil-microbe-atmosphere systems.* Plants play a central role in large-scale hydrology, energy and food security and sustainability, generation of atmospheric aerosols, carbon-cycling feedbacks between the land surface (soil) and the atmosphere, and other processes relevant to DOE research initiatives. Process-based (mechanistic) simulation of plant form and function, in the context of the overall ecosystem in which it exists (including microbial, soil, and atmospheric elements), is an area that is ripe for high-performance computation but requires a major effort. The issues surrounding this PRD are well described in a recent DOE workshop report (DOE 2015a) and will not be repeated in detail here. In summary, key issues are (1) integration of mechanistic models of individual plant processes into an overall plant model; (2) incorporation of genomic information into predictive simulations of plant form and function; (3) coupling of plant simulations with process models describing soil hydrology, microbiology, and atmospheric processes; and (4) development of a community software framework based on advanced software architecture design methods.

3.1.3.3 Cross-Cutting Research Directions

The scientific challenges outlined above can be addressed from two complementary perspectives: (1) bottom-up, which uses new mechanistic understanding gained from molecular- and cell-scale experiments and incorporation of mechanisms, pathways, kinetics, and other understanding into parameterizations for models that could be used to inform larger-scale predictive simulations; and (2) top-down, which integrates observational system-scale data with simulations across physical and temporal scales to refine predictive understanding. These two approaches define the two major cross-cutting research directions:

1. *Multiscale model coupling methods and framework.* Currently, scientific simulations are mostly conducted at a single well-defined physical and temporal scale. The equations, numerical formulations, parameterizations, and supporting data are all designed specifically to operate at that scale. Relationships to models and/or data at other scales are, if considered at all, usually constructed in terms of parameterizations of subscale effects at the scale of interest, either through formal upscaling or ad-hoc parameter tuning. In the context of our defined scientific challenges, this approach is problematic for at least two reasons. First, while the tuning of model parameters may enable models to accurately represent observed phenomena (descriptive modeling), the ability of such models to predict behavior under modified future conditions is questionable at best owing to the combination of parameter non-uniqueness and hidden model structural inadequacies. Second, because of mutual feedbacks between large-scale climate and cell-scale biological processes, the information flow of model information is not purely from small scale to large scale. In fact, in many cases, it will be necessary to dynamically connect models employing different mathematical process representations across scales both up and down during the course of a single simulation. The range of length and time scales spanned by processes considered in this session presents a daunting challenge but also an opportunity to leverage activities across several science application areas given that all areas face this challenge. Consequently, the development of generalized model coupling approaches and tools, built within an extensible and sustainable computational framework, is a critical cross-cutting need. The Scheibe and Wieder et al. white papers provide additional information.
2. *Use of observational data to inform and evaluate fine-scale models.* As observational technologies advance, both the amount and quality of data describing states and processes in natural systems are proliferating. These data are extremely valuable to the simulation process in many ways that involve significant computational needs. Model parameterization through formal data assimilation processes typically involves performing large ensembles of forward

simulations. When these simulations couple multiple process models,¹ the computational complexity significantly increases. Although such an approach is not often used today, inclusion of multiple scales of simulation models (as described in cross-cutting topic 1) would further multiply computational demands. Integration of observational data is also essential to model evaluation and uncertainty quantification (UQ), again often requiring the execution of large ensembles of alternative simulations. The utilization of large amounts of observational data also introduces needs in the area of scientific workflow management. This need is highlighted in the Maxwell white paper, which notes the need to compare billions of model outputs to millions of observational data in the context of continental-scale integrated hydrologic modeling. The process of preparing observational data (formatting, geographic registration, units conversion, data quality assurance) for comparison to model outputs can be extremely time consuming in the context of extreme-scale simulation and thus requires new scientific workflow environments. However, while these workflows share some characteristics, often the workflow requires significant customization for each specific simulation problem and data set of interest, hindering the development of generalized frameworks for model evaluation.

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.

¹ See, for example, <http://www.emsl.pnnl.gov/emslweb/coupled-hydro-geophysical-inversion-river-water-intrusion-and-biogeochemical-transport-modeling>, accessed August 2, 2016.

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.

| 2025 Needs | Omics Data Processing | Metabolic Modeling Simulations | Microbial Community Simulations |
|----------------------------------|-----------------------|--------------------------------|---------------------------------|
| Computational Core Hours | | 1,000,000 | 1,000,000,000 |
| Parallel | Yes | Yes | Yes |
| Adapted to GPU or Accelerator | No | No | Yes |
| Memory per Node | 1 TB | 20 GB | 1,024 GB |
| Aggregate Memory | 1 TB | 20 GB | 1,024 TB |
| Data Read and Written per Run | 10 TB (read) | | 320 TB |
| Maximum I/O Bandwidth Needed | | | 50 GB/sec |
| Percent of Runtime for I/O | | | Less than 1% |
| Scratch File System Space Needed | 1 TB | | 320 TB |
| Permanent Online Data Storage | | | 0 TB |
| Archival Data Storage Needed | | | 0 TB |

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.

3.1.4.1 Scientific Challenges and Opportunities

One of the largest challenges in biology is the increasing diversity, volume, and velocity of production of biological data ranging from spectral properties through molecular structure, chemical concentration dynamics, sequence evolution, and structured cells and communities. These data are possibly temporal and highly condition/context dependent and are often noisy enough to require sophisticated experimental designs with complex replicate structures to control for the variation in measurements and biological response. The computational frameworks for analysis of the raw data feeds that turn them into assertions of identity, structure, abundances, and dynamics require sophisticated, nuanced, and ontologically deep information representation; scalable numerical algorithms; and sophisticated, often interactive, visualization.

Biological data have all of the characteristics of “big data,” often summarized as the four Vs: volume, velocity, veracity, and variability. Technology improvements in biological instruments from sequencers to imaging devices are continuing to improve at exponential rates, with data volumes in petabytes today and expected to grow to exabytes in the future. The data streams from some instruments flow at hundreds of gigabytes per second, challenging the networking and storage infrastructures and often requiring some on-demand processing to select or compress data prior to storing it. The data volume and velocity are expected to increase exponentially, as detectors continue to improve in speed and resolution at a rate that exceeds that of Moore’s Law for computing. In addition, the use of robotics at experimental facilities will increase experiment throughput and thus data rates. The veracity of the data is inherently low — errors arise from the environment in which the data are measured and from the instruments and measurement techniques. The data are also highly variable because of biological diversity and the need to combine data from a diverse set of sources. Finally, biological–omics data lack the kind of geometric structure that is known *a priori* in physical science measurements, which has implications for the underlying algorithms and system requirements for analysis.

3.1.4.2 Priority Research Directions

Here we outline the challenges of the key measurement technologies driving big data challenges in biology.

3.1.4.2.1 Genome Sequencing

High-Throughput Short Reads via Illumina Dye Sequencing

Illumina’s sequencing-by-synthesis approach revolutionized sequencing by providing very high-throughput output at low cost. While Illumina is still the dominant technology for sequencing projects, it presents challenges for data analysis, especially *de novo* assembly. Most important, the

read lengths are currently limited to 2×300 bp mate pairs at the high-end platforms, making it hard to resolve repetitive regions of the genomes. While there are software solutions to address resolving the short tandem repeats (STRs), repeats that are longer than the read length can be impossible to assemble. Illumina library preparation methods that are polymerase chain reaction (PCR)-free improve evenness in sequence coverage, allowing sequence depth to be more reliably estimated. Illumina data have very low (<1%) error rates that are mostly the result of substitutions. However, Illumina errors have specific bias patterns and occur in specific sequence contexts, so repeated reads of the same data and statistical approaches may not eliminate the errors. Genome assembly and analysis are challenges on short-read data, and most of the assemblers used in production (especially those used on metagenomes) run only on single-node, shared memory systems because of the need for random access to a large memory space. However, distributed memory assembly codes such as Ray and HipMer are providing options that use the aggregate memory in an HPC system and its high-speed network to perform the assembly, thereby allowing for direct *de novo* assembly of larger data sets and large single genomes or metagenomes. The data volumes for sequence data will continue to grow exponentially, with roughly an annual doubling expected at JGI this year.

PacBio: Long, Noisy Reads via DNA Polymerase Monitoring

The Pacific Biosciences sequencing platform uses zero-mode waveguides to monitor nucleotide inclusion events by a fixed DNA polymerase at the bottom of a nanoscale well. The N50 for individual reads in a run is often 10–15 kb, with the longest reads >50 kb. Base-calling with this technology is somewhere between 80–90%, and by observing the integration times of individual nucleotides, epigenetic modifications can be imputed without the need for secondary assays. Although the error rate is high compared to short-read technologies, the errors are very close to uniform random noise, meaning that multiple passes over the same molecule (made possible by a circulating library preparation strategy) or the comparison of multiple reads covering the same underlying sequence results in base-calling accuracy at least as good as the best short-read technology. The randomness of errors is highly attractive, as the Illumina platform admits systematic biases that lead to patterned errors in certain genomic contexts — including simple lack of coverage. The DOE's JGI uses PacBio sequencing for the assembly of bacterial genomes, and the technology has been used to assemble complex eukaryotic genomes with results superior to the best manual (bacterial artificial chromosomes [BAC]-based) assemblies. The technology is also being used for transcriptomics in organisms with alternative splicing, as whole transcript isoforms are acquired, and hence covariation between distal RNA processing events can be quantified. A major limitation of the technology is the requirement for large DNA inputs (5–100 ug for a single run). Throughput (~20K reads per ~\$700 run) is not yet suitable for high-volume metagenomic sequencing, and this lower throughput also reduces computing and networking requirements relative to Illumina. Although the technology has consistently scaled read lengths and throughput, the quality or error rate of the reads is still high compared to other technologies. The next generation of the technology will likely enhance throughput by 6–7 times. The machines are also very large and hence not portable, so they cannot be used for *in situ* environmental analysis. However, the large facility-based sequencers may simplify the provisioning of network and computational resources, given that the computing can also be centralized in one or a small number of computing and data facilities and a limited number of high-speed networks used to connect them. Assembly algorithms for PacBio data differ from those used on Illumina data owing to the different error characteristics and read lengths, but both are ongoing areas of research, especially for metagenomes. The rapidly evolving landscape of sequencing technology will drive the need for new computational methods and tools for assembly and analysis; and while current algorithms are likely to be adapted to these workloads, the new technology may admit very different algorithmic approaches and thus different computing system characteristics.

Nanopore Sequencing: Direct Detection of DNA and (Soon) of RNA and Peptides

Oxford nanopore and other emerging technologies enable the direct detection of DNA by feeding single-stranded molecules through an electrified nanopore and monitoring the change in voltage. In principle, the technology is directly applicable to RNA and proteins, although technical challenges remain in these applications. Sequencing accuracy in v9 chemistries is above 90%. Because sensing is direct, as with PacBio, covalent modifications to nucleotides can be imputed. Moreover, the devices are small (longest dimension 10.5 cm) and lightweight (85 g) and require almost no electricity (1 watt). Environmental monitoring strategies that deploy nanopore sequencers via drones have been widely proposed, and small-scale tests are under way. This capability may generate demand for hardware-accelerated, *in situ* data processing. The massive parallelizability of this technology will likely result in substantial increases in data velocity in the next five years, with commensurate increases in computation and storage required. Raw data outputs for a single PromethION sequencer are around 1.4 GB/sec, although this output is immediately reduced by *in situ* hardware-accelerated statistical compression to ~40 MB/sec. The current generation of statistical compression is “lossy” (e.g., epigenetic information is lost), and improvements to the *in situ* analysis will result in improved scientific insight at the cost of less significant compression. By 2025, it is likely that individual investigators will be generating data at more than 100 GB/sec. This rate will constitute a sea-change in the velocity of biological sequence data. Unlike the large Illumina and PacBio systems, the Nanopore technology will raise a set of issues from streaming, distributed data collection. Individual data streams are small enough that they will not perturb the network or place significant demands on computing; however, in aggregate, a system of environmental sensors, for example, will require additional provisioning. The ability to transfer and analyze large data streams may improve the type and quality of information that can be preserved, because less *in situ* compression will be required.

Microfluidics and Single-Cell Sequencing

DNA sequencing technology and single-cell sample preparation assays are advancing rapidly in cost reduction and scale. The result of these advancements is increasingly larger numbers of cells that can be profiled using a growing array of high-throughput assays (profiling DNA, RNA, and epigenetic features). At the heart of these advancements is the development of microdevices — including microfluidics, reverse-emulsion droplets, and microwells. These microscale platforms can be tailored to enable easy capture and processing of cells, reducing labor and costs, while improving efficiency and consistency over conventional approaches. For example, reverse-emulsion droplets make ideal microreactors as they are precisely sized to contain one individual cell and can be manipulated individually: they can be filled, steered, split, combined, detected, and sorted, all at rates of thousands per second. This emerging technology already enabled the profiling of thousands to tens of thousands of cells (e.g., mitochondrial RNA [mRNA] and histone marks) in recent studies — a total that is expected to increase even more.

As opposed to the traditional bulk-sequencing methods, single-cell genomics allows researchers to observe the full spectrum of cellular phenotypes, which, in turn, enables them to address new fundamental questions. First and foremost, it provides an unbiased way to discern distinct and hitherto unknown subtypes of cells. For example, single-cell RNA-Seq was recently used to characterize the heterogeneity of cells in the immune system, which led to the discovery of novel regulators of immune functions. In addition, single-cell profiling provides a powerful tool for studying the mechanism of gene regulation, taking advantage of co-occurrence of signals (e.g., transcript abundance of different genes or the transcript abundance of a given gene and the accessibility of the nearby chromatin). Single-cell technology also has the potential to alleviate the need for large amounts of microbial DNA in metagenomics studies. Finally, genetic variation between cells provides yet another appealing application of single-cell genomics. Here, the DNA sequencing of single cells can be used to decipher the pathological evolution in a tumor, as well as somatic mosaicism in healthy tissue and in noncancerous disease. In addition, concomitant

measurements of RNA and DNA from the same cell are expected to contribute greatly to studies of expression quantitative trait loci (eQTL). In this approach, every cell is now an individual, and dependencies between single-cell genome variants and transcriptome variation provide a causal basis from the former to the latter. This treatment can boost the power of eQTL detection and analysis of factors that affect expression variance and temporal dynamics. It can also provide a unique lens in the case of tumors, where in principle, the analysis of multiple single cells from one tumor sample can be used to identify the genetic basis of transcriptional state variation between malignant cells.

Single-cell sequencing has the potential to transform the study of microbial communities, because it can capture the diversity of strains that bulk sequencing will not. The computational steps in analysis are similar but use different algorithms to address the higher error rates that result from the selection and amplification processes and to distinguish those errors from true biological diversity. The data storage implications are not clear — the single-cell data sets are likely smaller than a whole metagenome data set, but the variations across cells may increase the amount of data to be stored and shared with the community.

3.1.4.2.2 Mass Spectrometry

Mass spectrometry (MS) is a technique for characterizing the chemical makeup of a sample by ionizing the sample and sorting the ions based on their mass charge ratios. A central challenge across MS applications is the identification of molecular entities and structural features from spectra (i.e., which molecule(s) constitute particular peaks). Because performing accurate molecular identification from complex spectra is challenging and time consuming, in practice, only a small fraction of peaks are typically identified via automated means.

Tandem Mass Spec

To enable the identification of unknown compounds, advanced computational methods are needed to infer structural features of compounds on the basis of tandem, liquid chromatography (LC), gas chromatography (GC), and ion-mobility mass spectra. Tandem mass spectrometry involves fragmentation of analyte-derived ions via a sudden infusion of energy into the molecule, and the technique may yield the first (or only) data obtained on unknown compounds in complex samples. Current computational methods have focused on tandem MS for calculation of inferred fragmentation spectra based on quantum chemistry and spectral tree methods for enumerating potential fragmentation paths via *in silico* fragmentation. Application of existing computational methods in practice is currently being limited by (1) the high compute cost for calculating *in silico* spectra and trees combined with a lack of resources for the efficient application of the methods, and (2) high error rates as a result of inaccuracies owing to simplifications and errors in the modeling process. In addition to computational power, novel methods and algorithms are needed to enable high-performance, fast, and accurate computational identification of chemical identities from MS-based experiments.

Mass Spectrometry Imaging (MSI)

Mass spectrometry imaging (MSI) is widely applied to image complex samples for applications spanning health, microbial ecology, and high-throughput screening of high-density arrays. In the most common embodiment, a mechanism for ionization (e.g., a laser beam) is raster scanned across a sample, thereby acquiring mass spectra at each image location. Each of the acquired spectra has thousands of sharp peaks from the ionization of proteins and metabolites. MSI has emerged as a technique suited to resolving metabolism within complex cellular systems, where understanding the spatial variation of metabolism is vital for making a transformative impact on science. Unfortunately, the scale of the MSI data and the complexity of analysis required together present a significant barrier to scientists. Already today, common MSI data sets are on the order of five to tens of GB and can range up to hundreds of GB to TBs in size for very large images; performing

comparisons of multiple images is computationally intractable with the algorithms and computing resources available to an individual scientist. To increase the sensitivity and specificity toward the detection of molecules of interest, MSI is combined with ion mobility separation and tandem mass spectrometry, further increasing data sizes and complexity by several orders of magnitude through the generation of hundreds to thousands of spectra at every image location. In addition, MSI is commonly combined with other imaging technologies (e.g., light microscopy). Although these data can be collected routinely using current instruments, the broad application of MSI is limited because of a lack of resources and methods to store, manage, and analyze such large-scale, hyperspectral data.

New types of software frameworks will be required at the computing stage to lower the barrier to conducting integrated data analysis and perform extraordinary computational tasks through interactive use of high-performance systems. For example, the OpenMSI project provides a web-based gateway for management and storage of MSI data, along with easy access to HPC resources for statistical analysis and visualization of the hyperdimensional contents of the data. However, many research challenges remain regarding the management, provenance, and analysis of large-scale, multimodal, and hyperspectral MSI data using advanced computing (Figure 3-4; also see Section 3.1.4.3).

LC and GC Proteomics

To increase the sensitivity and specificity toward the detection of molecules of interest, chromatographic separation is commonly used prior to mass spectrometry detection. The physicochemical properties of molecules can lead toward a greater or lower affinity for materials. By flowing the molecules over these materials, researchers can induce a separation of molecules with different physicochemical properties. In liquid and gas chromatography mass spectrometry (LCMS and GCMS, respectively), samples are injected, and the separation causes the molecules to arrive at the mass spectrometer earlier or later in a run. This “retention-time” is an additional distinguishing feature guiding the identification of molecules.

It is not surprising that the solvent and stationary material used for chromatography has a very large impact on the retention time. Each laboratory often uses its own subtly different methods, and therefore comparing measurements between labs presents a major challenge. What is needed are controlled vocabularies that describe methods; advanced tools for biological analysis; and computational tools for management of petabytes of scientific observations, real-time integration of new results with historical findings, identification of best practices for sample preparation/analysis, and dereplication of results to identify new insights from repetition.

Direct Infusion, Matrix-Assisted Laser Desorption/Ionization (MALDI), and Other Techniques

Although highly desirable for compound identification, chromatography increases the duration of sample analysis from milliseconds to tens of minutes. Consequently, many mass spectrometry experiments are performed without chromatographic separation. These methods acquire a relatively small number of spectra for each sample and can process thousands of samples per day. A recent experiment with NIMzyme screening of cellulose-degrading enzymes at DOE’s Joint BioEnergy Institute (JBEI) simultaneously measured and analyzed 10,000 samples in one day. The routine application of such high-throughput assays will enable new large-scale and data-intensive experiments and will be critical to cost-effective, biomass-to-biofuels processes. As often occurs in these measurements, the compound of interest is known, and the mass spectra are used as a screen for the compound of interest. Efficient data management, provenance of experiments and analyses, management of analysis workflows, and methods to quickly analyze and compare large collections of measurements are critical for these applications.

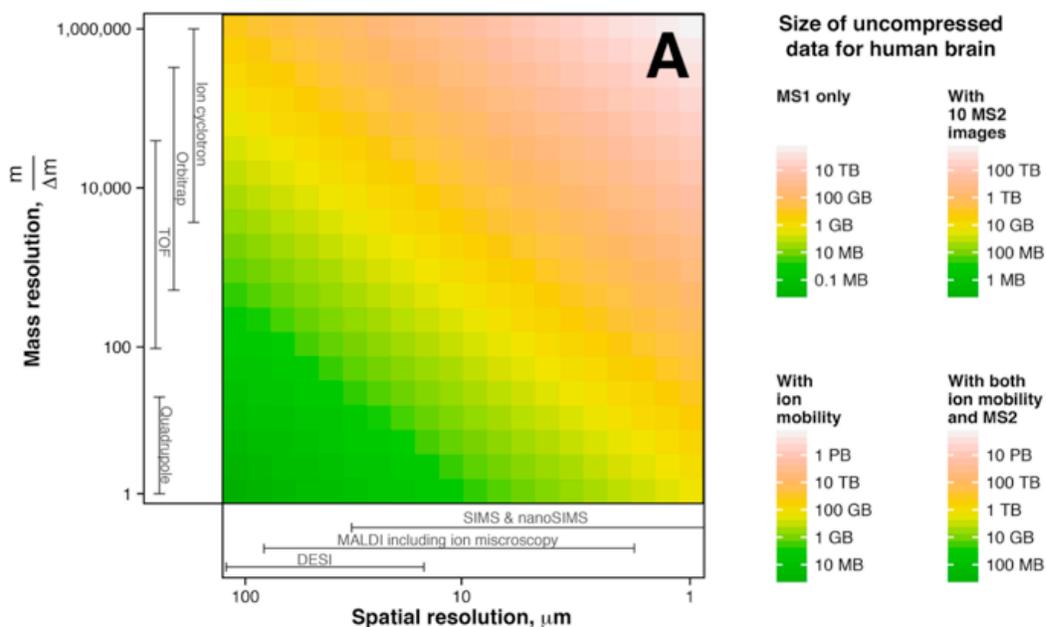


Figure 3-4. Projected estimate for file size of an MSI data set of a single cross-section through a human brain using various MSI techniques. As advances in spatial and mass resolution are combined with new dimensions, including ion mobility separation and MS2, raw data sizes will increase dramatically (image courtesy of Curt R. Fischer, Oliver Ruebel, and Benjamin P. Bowen, Lawrence Berkeley National Laboratory).

Deep Metabolome Annotation

Bioremediation, biomanufacturing, precision agriculture, biomolecular nanotechnology, bioenergy, and other concepts have emerged as industrial goals with horizons on a 20-year timescale. To apply robust engineering models toward realization of these goals will require greater understanding of the complex biological processes involved. To achieve these goals, MS-based metabolite assessment will likely be a necessary capability. Currently, samples can be routinely analyzed for metabolite characterization, and researchers are only scratching the surface in terms of the comprehensive measurement that is feasible and needed. Although identified as needs, areas that are missing include standardizations in sample preparation, data acquisition, quality control, data analysis, data management, and knowledge storage.

It has been estimated that the eukaryotic cells host around 200K distinct metabolites. Complex microbial communities (e.g., in soil) may be far more diverse. It is generally possible to identify a few hundred (up to ~1,000) metabolites using high mass-accuracy or chromatography-aided mass spectrometry. Hence, we know around only 0.5% of even the most complete plant metabolomes. At present, there are a few individual investigators and several consortium-scale proposals to remedy this situation and generate a deep, comprehensive map of the metabolome for several model organisms. These efforts will be as transformative for the interpretation of metabolomics data as gene annotation has been for genomics data. Success will also result in profound (>1,000-fold) compression for metabolomics data, currently a driver of biological data volumes.

The most reliable approach to matching unknown spectra to compounds relies on spectral databases. In this approach, the tandem mass spectra of a large number of pure authentic standards are acquired and compared to unknown spectra. This approach is useful only when pure standards are available. Unfortunately, these standards are typically not available commercially. Many tools have been developed to infer structural features of the unknown compounds on the basis of tandem mass spectra. These include (1) calculation based on chemical principles such as quantum chemistry, (2) spectral tree methods, and (3) databasing of spectra of pure compounds.

Although each of these three tools is amenable to being applied efficiently in an integrated manner, computational tools are needed to achieve this outcome efficiently.

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.

3.1.4.2.4 General Trends

The application of microelectronics to imaging (initially with visible light) began half a century ago and has followed the same explosive growth in data rates as other microelectronic-enabled technologies. State-of-the-art visible light, X-ray, and electron imaging detectors currently have data rates of ~10 GB/sec. Next-generation detectors will double or quadruple these rates in the near term, and rates of 100 GB/sec will be routine in the next decade.

High frame rates allow wide dynamic range (by addition of images) but, more important, they allow detection of single probe particles — which can then be used to improve image quality (e.g., by centroiding the position of the probe particle and determining its energy, its arrival time, etc.). This approach results in large numbers of sparse, low signal-to-noise images, which would ideally be preprocessed to vastly reduce data volume.

3.1.4.3 Cross-cutting Themes and Methods in Data Analytics

The complexity and size of the data generated by biological instruments, including multimodal and high-throughput scenarios, make advanced visualization and analysis methods indispensable. Advanced machine learning, dimension reduction, and statistical data analytics are commonly applied to identify and derive important features from the resulting data sets, and the growing scale of data necessitates the development of advanced, scalable methods that can efficiently process the data in a timely fashion. In addition, novel visualization and analysis methods are needed to enable the integrated analysis and comparison of data from multiple modalities and across experimental conditions and environments. Furthermore, methods for management of the complex analysis workflows will be critical.

3.1.4.3.1 Data Management and Provenance

Efficient data management in combination with data provenance are critical elements to enabling large-scale, high-throughput experiments. This capability includes the management of complex relationships between data from multiple modalities to enable integrated data analysis and fusion for investigations across spatial scales and fields of investigation. The size and complexity of the

data increasingly make high-performance computing indispensable, requiring the close coupling of data and compute capabilities. Methods to enable users to share and collaborate through data, analysis, and computing are critically needed to fulfil the full potential of these techniques to advance and transform metabolic and bio-energy research. The development and application of cutting-edge analytical methods is a core driver in genomics, MS, and other technologies, whether applied to medical diagnostics or energy applications. To ensure high-quality, reproducible, and nonredundant biological scientific discovery on massive amounts of data, the community requires the means to share, apply, and reproduce analyses. To address this central challenge, novel methods for shareable and reproducible data analysis that support standardized data and analysis storage and interfaces, provenance, and workflow management are critically needed. Ultimately, these methods must be easily accessible to application scientists with varying degrees of expertise in computing, which makes novel and advanced programming, web, and graphical interfaces indispensable.

In situ Analysis in Workflows

In situ analysis is ubiquitous in biological workflows. Virtually all sequencing platforms conduct *in situ* data transformation of one form or another to facilitate rapid reduction of large images, videos, genome sequences, or electrical signals into forms suitable for downstream analysis. How much processing should be performed directly in the device, on site with associated computing, or in a centralized high-performance computing center needs to be evaluated by considering overall costs. This evaluation includes the cost and speed of networking; the relative cost of centralized vs. onsite computing; and the possibility of specialized computing platforms based on field-programmable gate arrays (FPGAs), custom chip designs (ASICs), or nonstandard computing devices (e.g., neuromorphic processors). As distributed sensing is deployed via drones and other modalities, *in situ* processing and hardware-accelerated data reduction will need to generalize and scale to complex networks of lower-power sensors and computing devices.

3.1.4.3.2 Statistical Data Analytics, Machine Learning, and Inference

Statistical data analytics, machine learning, and inference are central to virtually all large-scale data analytics in the biological sciences. Supervised, semisupervised, and unsupervised methods are all common in parametric, semiparametric, and nonparametric formulations, with the latter beginning to dominate diverse application areas. Increasingly, machine learning frameworks are used to assess confidence or stability for imputations. These methods often have weak or nonexistent theoretical guarantees; and model misspecification, resulting in erroneous confidence regions, is a systemic problem. Best practices and standards are needed, and the BER-ASCR community is well positioned to take a leadership role in establishing these.

It is important that for most if not all methods, there is a strong trade-off between meaningful interpretability and high-quality prediction. By interpretability, we mean that one can interpret the output of the method in terms of processes generating the data: this ability typically requires identification of a small number of elements of the actual data; and by high-quality prediction, we mean the ability to optimize the performance of some measure such as precision, recall, etc. This problem is particularly acute for scientific applications, where the number of samples (n) is typically lower than the number of parameters (P). The reason is that major drivers for much of the recent interest in machine learning techniques has been Internet and social media, primarily focused on prediction accuracy. In scientific applications, however, researchers are typically interested in using methods to obtain insight into the world, and thus, it is desirable to interpret the output of algorithms; at the same time, high-quality prediction is important for diagnosis and forecasting. Identifying a small number of highly predictive, interpretable features is critical both to gain insight and to determine interventions and next-step causal experiments. Thus, there is a pressing need to develop statistical data analysis methods that are designed to be interpretable from the domain from which the data are drawn and that also enable high-quality prediction and are generally applicable to a wide variety of natural scientific data sets.

Deep Learning

The most rapidly advancing area of machine learning is undoubtedly the construction of deep architectures of learners, which have now shattered nearly every performance record in the community. These are currently the dominant methods for analyzing images, videos, and speech with clear applications to bioimaging data of various kinds. The most widely pursued architectures are deep neural networks (DNNs). At least as important as their predictive power, these techniques yield hierarchical representations of input data, which have had a substantial impact on particle physics and, likely soon, on biology. Convolutional neural networks (CNNs) are highly effective in video and image identification, segmentation, and registration. Recurrent neural networks (RNNs), particularly Long Short Term Memory (LSTM) machines, have been enormously successful for processing time-varying signals (e.g., speech and videos). Because their base calculations are linear (affine transformations), these methods are accelerated in GPU environments. Optimal design strategies for the architecture of deep networks are unknown and are likely application dependent. Industry leaders (IBM, Intel, Nervana, etc.) are putting considerable effort into hardware acceleration for these algorithms, both for training and deployment. However, the design space is constrained by trade-offs in, for example, power, speed, accuracy, etc., and so a single solution is unlikely to be optimal for all problems. Furthermore, in general, extracting interpretable features from data with DNNs is an area of ongoing research, and they require more training samples than are typically available from experimental data. In addition, rigorous theory for DNNs is lacking; however, recent advances have revealed, for example, the importance of second-order methods in evaluating the objective function: in high dimensions, the difference between local and global optima vanishes, whereas the prevalence of saddle points increases. Much more applied mathematics research is required, and deep connections to non-equilibrium statistical mechanics present opportunities for cross-cutting impact.

Ensemble Methods

Any time weakly dependent predictors (i.e., with different biases) can be fitted, an ensemble is assured to exceed the prediction performance of any individual base learner. Indeed, the state-of-the-art for many machine learning competitions (which emphasize prediction accuracy) is constructed from ad-hoc combinations of, for example, Deep Networks + Support Vector Machines + other methods. It is extremely difficult to interpret “what” is actually being extracted from the data in such cases. Random forests (RF), an ensemble of locally greedy, globally random decision trees, is a general, out-of-the-box-successful machine learning algorithm. Recently, iterative strategies have been formulated that enable the (nonparametric) discovery of high-order interactions at the same computational cost as interactions of order two (iterative Random Forests, or iRF), although currently, the discoverable order is limited by the log of the sample size — the maximum order is $\sim \log_2(N)$. Alternatively, the recently developed Union of Intersections (UoI) is a flexible, modular, and scalable framework to enhance both the identification of a small number of highly predictive features and the estimation of their contributions, and thus results in improved data prediction in linear and nonlinear regression/classification, while being algorithmically scalable and stable in high-dimensional and extreme-scale data sets. The principles behind both iRF and UoI are broadly applicable to ensembles of learners, and hence propagation to other base learners, particularly deep neural networks, may be fruitful. Hyperparameter optimization procedures have enabled the construction of weakly dependent deep architectures. Like all ensembles, these will be highly parallelizable; however, millions of CPU/GPU hours will be required to fit each base learner. The major compute costs associated with all machine learning methods is training. Hence, once fit, some will run on a laptop or a drone, enabling feature discovery and prediction, such as quantitative time-dependent measures of impact from drought, floods, or plant disease across large hard-to-reach regions. However, the computational costs of running many CNNs in parallel would be prohibitive for general mobile computing platforms, suggesting that data will be transferred to a centralized computing or custom ASICs required for field-deployable systems that perform online learning. For example, an imaging processing pipeline may be trained to recognize environmental features, such as the impacts of drought or infestations, using data that are constantly updating the model.

Inference and Uncertainty Propagation for Extreme-Scale, High-Dimensional Heterogeneous Systems

It is clear that modeling strategies for biological data are complex and hierarchical. Uncertainty enters into inferences at many levels, from biological variation, to technical noise, to sequential optimization procedures that often admit stochastic or otherwise nonconvex steps. Data inputs are high dimensional and heterogeneous, often with ambient dimension much larger than the number of observations ($N \ll p$). In many cases, particularly in the analysis of imaging data, even the ambient parameter space is initially unknown and requires imputation, for example, via deep machine learning architectures. Dependence structure is generally unknown, and nonlinear interactions of order $>\log(N)$ are common in some domains; their imputation constitutes a central challenge. Response structure is heterogeneous: interactions between parameters may differ across samples, particularly in spatiotemporally resolved settings in the environmental sciences, along a genome, or throughout a transcriptome. Integrative procedures with sufficient sophistication to apply different sets of predictive rules or functions in different contexts (e.g., cell types, temporal windows, conditions) are required. In short, we need to be able to handle data heterogeneity seamlessly, even when we have no hint as to the presence or structure of homogeneous processes *a priori*. The ecosystem of analytical procedures for the integrative analysis of biological data is expanding rapidly in scale, scope, and sophistication, and concomitant advances in uncertainty propagation are needed; these will require new foundational statistical theory. Emerging theoretical approaches focus on stability assessment through coupled data and model perturbation and require dependent sequences of simulations. It is becoming apparent that accurate confidence hypervolumes that appropriately propagate both biological and technical uncertainty in the intrinsic dimension will amplify the computational demands of integrative analysis procedures at least 100-fold and likely 1000-fold, and the necessity for sequential simulations will limit the parallelization of these procedures. Weakly separable models with long-range dependencies, including genetic and ecological networks, will require low-latency architectures. Pushing accurate statistical inference and uncertainty assessment into high-dimensional, heterogeneous settings will yield transformative advancements in the stability, reproducibility, and reliability of environmental science.

Dimensionality Reduction for Visualization and Understanding

As with the supervised methods described above, for unsupervised methods (e.g., dimensionality reduction [DR]), there is generally a trade-off between the accuracy of preserved high-dimensional structure in low-dimensional embeddings and the interpretability of those dimensions. Principal components analysis (PCA) is commonly derived from a small number of components under linear-Gaussian data assumptions, whereas t-Distributed Stochastic Neighbor Embedding (t-SNE) is a nonlinear DR method that accurately preserves high-dimensional distances in low-dimensional embeddings. However, individual components are usually complex combinations of physically meaningful features, which can aid visualization/clustering but severely hinder interpretability. Independent components analysis (ICA) attempts to “demix” signals, while CUR finds low-rank approximations of the data matrix that are explicitly expressed in terms of a small number of actual columns and/or actual rows of the data matrix. Both ICA and CUR result in more interpretable dimensions. Even the best of these algorithms are extremely scalable but are underutilized. For image/movie data, dictionary methods such as convolutional sparse coding (CSC) or convolutional sparse nonnegative matrix factorizations (csNMF) are methods that balance the preservation of local and global structure. These convolutional methods are typically nonconvex optimization problems, so scalable algorithms are needed and GPU acceleration would likely be fruitful.

Combinatorial Analysis

To associate genome structure and variance with physical and molecular phenotypes, comparative analysis across genomes is used. The goal is to determine the pleiotropic and epistatic relationships underlying cellular functions that will deliver further insights into the molecular basis of complex, multigenic interactions responsible for biological complexity and the emergent properties of

complex systems. Rather than full all-to-all comparisons, more sophisticated combinatorial approaches are being used. These computations involve massive databases of genomic data and are currently using petascale systems, including the GPU-accelerated Titan system.

Bayesian Analysis

Bayesian inference is a data inference method determining model parameters based on statistical analysis of a given data set and can incrementally update the parameters based on additional data. For example, it is a popular method for reconstructing phylogenetic trees and for constructing 3D structures from 2D projections in cryo-EM.

MrBayes is an example of software used for this problem, which has a parallel version based on MPI.

3.1.4.4 Computing Needs and Requirements

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.

3.1.4.4.2 Computational Requirements

Computational challenges in big data are somewhat different from those encountered in modeling and simulation and therefore place unique requirements on the computing facilities. At least two biological problems, cancer genomics (deep learning and other analysis, in collaboration with the National Institutes of Health [NIH]) and microbial analysis (assembly, annotation, and cross-metagenome comparisons) have been identified as exascale-level problems by the community. However, there is little experience with the new data sets, and the computational methods and implementation techniques are still evolving in significant ways, which makes more precise quantitative predictions of computing needs impossible. Qualitatively, the analysis methods have some characteristics that may drive system requirements in different directions:

- Some biological data analysis problems are currently performed using many independent compute nodes and thus do not take advantage of the high-speed interconnect and tight integration of an HPC system. For example, aligning a set of genome samples against a reference (especially common in human data given that there is an existing assembled reference genome that is small enough to replicate across nodes) can be carried out with this type of independent parallelism. Because there may be a large number of such problems, some type of job management layer that assigns work to nodes is very helpful, especially if the time taken by each varies. HPC features such as large NVRAM memories can still be useful, and the network can help with the distribution and management of jobs.
- At the other extreme, some data analysis problems are conducted using single nodes with large amounts of shared memory, and there are biological data sets that will take advantage of the largest possible shared memory. Genome assembly falls into this category, although both are being moved to distributed memory platforms. The assembly algorithms require random access to large shared data structures, and most implementations use some type of one-sided communication, such as UPC (unified parallel C), UPC++, or MPI-3.0's one-sided features. Machines with low overhead (minimal software overhead or hardware acceleration), low latency, and high injection bandwidth are important, because the algorithms tend to send small messages in irregular patterns and at unpredictable times. The lightweight one-sided communication, including remote atomic operations for synchronization, allow the aggregate memory of an HPC system to be used like a shared memory space.
- Genomic analysis typically does not use floating point arithmetic but is dominated by string manipulation and memory operations. Hashing, sorting, and histogramming are all common kernels.
- Machine learning algorithms (both classical and deep learning) currently use floating point arithmetic. Although these algorithms are implemented with floating point arithmetic, there are some cases (e.g., deep learning) when low precision (16-bit floating point) or some type of fixed-point arithmetic based on integers may be sufficient.

- Machine learning algorithms are based on linear algebra that is familiar to modeling and simulation experts, and both iterative and direct methods are used. The primary difference from a workload perspective will be in the sparse matrix structure: while physical modeling problems tend to have matrices with patterns of small dense blocks, bands near the diagonal, or other features that will lead to some locality of access, data analysis problems may have very large, sparse, and unstructured matrices, which will lower both spatial and temporal locality and make memory and network latency and bandwidth more important.
- Deep learning algorithms also represent an extreme point in the workload, dominated by dense matrix-multiplication and convolutions for the training phase. These are highly regular and very computationally intensive, at least at the node level, while all-to-all communication is used between nodes. GPUs are the current architecture of choice for these problems, because they offer highly efficient floating point operations with data parallelism. Specialized features or special-purpose architectures are also appearing to support this workload in the commercial market.
- As noted throughout the early science examples, having access to high I/O bandwidth to stream data into an HPC system from some external measurement device or the local storage system will also be essential, and large NVRAM memories may prove important for holding the large data sets.

3.1.4.4.3 Networking Requirements

Data rates from sequencers, imaging devices, MS, light sources, and other experimental platforms will drive the need for network bandwidth increases to connect these facilities with the computer centers for analysis and storage and to serve to the broader community. One example noted earlier is the 400 GB/sec network between NCEM and NERSC to send the imaging data directly into the computing nodes. Other experiments perform some amount of computing onsite, although as computing demands grow, that onsite support may also be expensive and difficult to manage because of both power/cooling infrastructure and personnel. Embedded networks of environmental sensors will require different types of network support with broader reach beyond the facility networks and less bursty data flows.

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.

3.2.1.1.2 Priority Research Directions

Balancing research priorities within climate prediction involves the appropriate use of the increased computing power, enabling us to balance increased resolution, domain size, and the use of ensembles. Each aspect adds value but not necessarily to the same degree for all applications. As reproducibility becomes more difficult on exascale computers, ensembles will become critical, both to identify outlier simulations that contain suspect calculations, as well as to gain a more statistical representation of climate. Whereas Earth system modeling has traditionally been deterministic, we will be forced to view climate simulations more like natural, real-world experiments where reproducibility is not expected, and sampling must be performed to determine the mean and variability of the measurements.

Convection Parameterization Development

Even with global cloud-resolving models (GCRMs) that use grid spacing near a few kilometers for horizontal resolutions and tens of meters for the vertical resolution, Earth system models will still rely on parameterizations to calculate the subgrid-scale processes of cloud microphysics, atmospheric turbulence, shallow convection, and impact of subgrid surface heterogeneity. In the next five to ten years, Earth system models will still primarily rely on parameterization to calculate deep convection. In the last several years, one promising approach of parameterizing deep convection has emerged in which three-dimensional, cloud-resolving models are embedded

within each grid box of a coarse-resolution global model. Earth system modelers should continue to pursue more parameterization approaches that place less demand on computational resources. These approaches should incorporate resolution awareness in the parameterizations so that they are valid in variable resolution models. They should also consider different parameterization components as systems so that interactions among the components are treated appropriately.

Development of improved cloud parameterizations involves two related but separate issues. The first is the ability to parameterize clouds in the “gray zone” where model resolution is sufficient to partially resolve significant portions of cloud systems yet remains unable to resolve their features enough to capture the important cloud processes that affect weather and climate. The second is the need for resolution-aware parameterizations that yield correct results across a range of grid spacings, and preferably also provide improved results at finer grid spacing. Both of these needs are of primary concern for development of the ACME model and the foreseen use of regionally refined grids.

Large-Domain LES for Cloud Systems

Improvement of cloud parameterizations requires a range of tools. Highlighted within the breakout session was the need for models capable of sufficiently resolving cloud processes with domains large enough to contain entire mesoscale convective systems throughout their lifecycle. This improvement would require large-eddy simulation models with grid spacings on the order of 100 m and domain extents on the order of 2000 km. Domains such as this are roughly 100 times larger than state-of-the-art domains used today and are comparable to one-off “hero” simulations currently published in the literature (Khairoutdinov et al. 2009; Schalkwijk et al. 2015). In addition, expanding from current, common LES domains to these larger domains represents a weak scaling problem in terms of computational need. Thus, with appropriate refactoring of LES models to take advantage of exascale-style computers, these large LES domains are achievable within the next 5–10 years, provided that I/O is able to keep up with the model needs and that support is provided for developing efficient models. Important for enabling true advances in the understanding of organized convection is the ability to use these large LES domains on a routine basis instead of as one-off simulations. Researchers need the ability to quickly and easily produce these simulations to test scenarios, compare different convective events, and build a library of simulations to interrogate for parameterization development purposes. The DOE ARM program is currently developing the LES ARM Symbiotic Simulation and Observation (LASSO) workflow (Gustafson and Vogelmann 2015; <http://www.arm.gov/science/themes/lasso>) to produce a library of LES simulations over its megasite observation locations, such as the Southern Great Plains site in Oklahoma. ARM’s current efforts focus on shallow clouds, which require only a small domain. Routine access to petascale and exascale computers could increase the ability of the ARM program to expand its modeling to generate a library of deep convective LES simulations to complement their observations. This modeling would greatly help in understanding how convection initiates and upscale growth sustains the convection over long time periods.

A critical aspect of achieving success with large LES domains for deep convection will be developing improved microphysics for handling the cloud formation. Solely increasing model resolution with today’s models will provide some insights. However, the fidelity of clouds will be hampered by the inaccuracies in the handling of ice for cirrus and mixed-phase clouds. In addition, efforts will be needed to improve spectral bin microphysics such that it will be cheap enough to run with the LES domains. The bin microphysics has been shown to greatly improve simulated convection at cloud-resolving scales (Fan et al. 2015), and this finding will become even more important for LES grid spacings. Current implementations of spectral bin microphysics roughly increase computation cost over bulk microphysics by an order of magnitude. Efforts to better handle transport of the many additional tracers required for bin microphysics, combined with yet-to-be identified algorithmic improvements, will be necessary and could benefit from ASCR’s research into algorithm development.

Improvements Expected for Global Modeling

GCRMs will have horizontal resolutions on the order of 1 km and vertical resolution of less than tens of meters in the atmospheric boundary layer and about 100 meters in the free troposphere and stratosphere (Miyamoto et al. 2013). These models will be able to encompass large convective systems and mesoscale weather systems that play key roles in the vertical transport of atmospheric moisture, energy, and momentum and in producing extreme precipitation and weather events. Relative to the current “high resolution” Earth system model resolution of 50 km in the horizontal and 50 layers in the vertical, GCRM will increase the computational amount by about 2,000 times. While the anticipated increase of computational resources by 100 times will not be sufficient to enable GCRMs to be used operationally for long-term climate variability simulations, GCRMs should be developed and used as a research tool for studies of extreme weather events, interactions of clouds and atmospheric circulations across scales, roles of atmospheric deep convection, and how the clouds should be parameterized in coarse-resolution models. This development will require improved dynamical cores that are nonhydrostatic and capable of scaling to the largest of DOE’s computers.

Because GCRMs will be too expensive to run for long-term Earth system simulation, model development is also needed to improve the multiscale modeling framework (MMF) method. One option for this is the Q3D-MMF (Jung and Arakawa 2014). This method has shown promise for small-scale pilot models, and it is at an appropriate stage in its development to be expanded to global scale. This method will improve upon current MMF techniques with 2D embedded LES models within each GCM column, and essentially will act as a sophisticated resolution-aware convective parameterization that converges to a GCRM as the resolution is increased. Q3D MMF will be expensive to use but will scale well in an exascale environment and should be capable of running climate-length simulations.

Data Assimilation

Data assimilation serves two purposes for Earth system 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 predictive simulations can be conducted. Advances in atmospheric data assimilation in the last several years have shown promising approaches to build data assimilation systems for Earth system models with minimal new algorithm and software engineering investments. An example is the DOE CAPT project in which reanalysis products from operational weather centers are used as a shortcut to directly assimilating observations (Phillips et al. 2004). Another example is Ensemble Kalman Filtering in which ensemble simulations are used with observations to enable statistical estimates of the climate states. Future research needs to develop data assimilation techniques at the appropriate spatial scales with the appropriate targeted observations that 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. Data assimilation will also be an important tool for the high-resolution models to help ensure that they are as realistic as possible. Applying data assimilation at LES scales is an open research area that needs further work to ensure that observations are not overly smoothed over large spatiotemporal scales that are inappropriate for the highly refined model grid spacing. (See Section 3.2.1.2 for expanded discussion of challenges and opportunities, priority research topics, cross-cutting directions, and computing needs and requirements related to this topic.)

3.2.1.3 Cross-Cutting Research Directions

Many aspects of atmospheric research share research needs with other aspects of BER computing. A recurring theme throughout this breakout and other breakouts was the use of ensembles. Current approaches to ensemble modeling involve outputting results from every ensemble member and developing the ensemble statistics after the simulations are complete. These require outputting

everything from all members and also require reading all of the information back into memory again to calculate the statistics. Current approaches to ensemble modeling involve outputting results from every ensemble member and developing the ensemble statistics after the simulations are complete. Today, this effort requires outputting everything from all members during the compute phase and reading all of the information back into memory for the post-computational data analytics phase. Ways to minimize the extra data movement — such as by executing all ensemble members simultaneously and allowing the ensemble statistics to be computed as the ensemble simulation proceeds — would be useful areas of exploration. Although there will always be circumstances where the size of an ensemble is too large to allow concurrent simulation, explorations of new hierarchical and nonvolatile, high-performance memory systems may open the door to more efficient end-to-end computational approaches that minimize the inefficient data movement solutions seen on today's architectures. In addition to improving the I/O efficiency for the entire workflow, another advantage of being able to run the entire ensemble at once would be the ability to query the ensemble during the simulation, allowing researchers to introduce new capabilities in the development process. For example, information about the ensemble spread could be used to inform aspects of the simulations, such as being able to test hypotheses and end a simulation once a given statistical confidence threshold is reached. In addition, errors arising from machine faults or programming flaws could also be detected more readily by intercomparison of the ensemble members and flagging those whose behavior deviates noticeably from the others.

Large LES domains will also be useful for understanding issues surrounding land-atmosphere interactions, such as those that require simulation of entire hydrological basins. Complementing the large LES domains will be an increased use of CFD modeling for urban environments. Simulating flow within cities requires grid spacings on the order of meters and thus cannot be performed with traditional atmospheric LES models. Of benefit will be the development of better methods to integrate CFD and LES for nested applications to understand city impacts over larger spatial scales.

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.

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.

3.2.1.2.2 Priority Research Directions

To develop a predictive capability of the Earth system, to quantitatively describe the past and monitor the current state of the physical Earth system, and to rapidly respond to human and natural events, data assimilation with Earth system models is needed. Accomplishing this assimilation will require considerable research in characterization of uncertainty and error in the observations, component and coupled model, and assimilation algorithms.

One goal of this effort is to determine the best and longest possible estimate of the probability distribution of the Earth system’s variability. To achieve this objective, research is needed on the methods and models to use and on maximizing their capabilities on HPC. Computational research areas will include load balancing of assimilation systems on massively parallel computers.

Another priority area is research to ensure that the system state estimates are physically consistent without spurious shifts, whether made for ~200 years using only surface observations as observational input or for ~20 years using the full suite of observations: from space-based radar and GPS to satellite radiances to ground-based Doppler lidars and millimeter-wavelength radars.

Because the data assimilation systems use a comprehensive, numerical model-based first guess, and ensemble-based assimilation methods use on the order of 10 to 100 realizations of the complete state of that model, corresponding HPC needs are large and demanding. They cross-cut the issues of running large ensemble with atmospheric, oceanic, land, and cryospheric system modeling components. Reliable and fast CPUs and easily divisible model components are needed to integrate the model rapidly to form the first-guess states and then send those states to the processors to merge with the observations to form the reanalyzed states. The entire process is then repeated. Well-balanced I/O capabilities are essential for the efficient execution of reanalysis tasks. A more efficient workflow, including I/O, allows for higher-fidelity results using higher-resolution modeling frameworks because model resolution is often compromised in order for the analysis to take place within a reasonable length of time (typically measured in months).

3.2.1.2.3 Cross-Cutting Research Directions

The computing requirements for generating large ensembles using atmospheric, oceanic, and Earth system models on the order of 100 members cuts across many topic areas, including data assimilation.

The use of iterated data assimilation and short-term forecasts to identify and possibly correct model errors is another key cross-cutting research area. The data assimilation capability gives researchers the ability to perform such diagnoses directly with observations from short-term forecasts on the order of a few hours. This approach should help isolate errors before the nonlinear equations have taken an initial error in one variable, such as meridional wind, and rotated that error to another variable that affects processes, such as vertical wind affecting precipitation.

There is considerable overlap with the methods in uncertainty quantification, as such uncertainties must be accounted for in an optimal assimilation strategy.

We can envision a concerted research effort to use such error information from UQ or from short-term forecast error studies for maximal impact by converting the numerical model to a stochastic partial differential equation where the stochastic term comes from the statistics of the first-guess minus observations accumulated over many assimilation cycles.

Another cross-cutting research area is determining the difference between integrating from assimilated states and “initializing” a model, mainly in long timescale systems (e.g., ocean and land) where spin-up is more important. To make these two equivalent, researchers would need numerous assimilation cycles for long spin-up systems. Some investigation must be performed to determine the timescales of assimilation needed for appropriate coupled model initialization. An alternative strategy may be to relax the coupled model to some large-scale, reduced-space estimate of plausible trajectories.

Another important research idea is to determine the difference between the current testbed strategy of nudging or specifying from a separate initial state and then making 5-day forecasts using the atmospheric component of the coupled model to diagnose errors in the fast physics as compared to the alternative strategy of instead performing actual assimilation using the model to be evaluated for the initial states for those forecasts.

A challenge will be to maintain data provenance and security, both of the original ingested observations and the output reanalysis fields.

3.2.1.2.4 Computing Needs and Requirements

Large storage on the order of 5–10 PB with convenient access for outside users is needed to maximize the utility of any reanalysis data set produced at exascale. Because of the diversity of users, the data set has to be duplicated, with one storage strategy being the “timeseries” of single variables at one or a few levels organized by years and members, while the other is to allow access to the complete hypercube of reanalysis states for individual members for a short time period, such as a month. Having the highest resolution output available, as well as time averages, such as monthly means, is important for increasing the use cases and making the output as convenient as possible for a wide variety of users.

As an example of the computational workflow, for the cycling of the Ensemble Kalman Filter to be used in 20CR version 4, a researcher must integrate 80 ensembles simultaneously using thousands of cores (each ensemble member can be on one or two nodes, up to the number of nodes that have strong scaling) for a few minutes, gather the members, and combine with observations with the EnKF algorithm again using the same number of cores for a few minutes. This procedure is then repeated. It can be performed in parallel for years separated by decades, so that many years can be reanalyzed simultaneously and tens to hundreds of thousands of cores can be utilized.

After reanalysis data are generated, these data must be provisioned both to project and outside users in various configurations for study. These provision requirements mean that the data set should be essentially duplicated for key variables so that some users who need full three-dimensional states for a few members and a short period of time are equally served with those who require a few variables on selected levels for the entire period of the reanalysis.

As an example of the size, the current 20CR version 2c stored at NERSC — spanning the years 1851 to 2014 at ~200-km latitude-by-longitude resolution and 28 levels in the vertical — takes up 60 TB on spinning disk of the NERSC Science Portal (portal.nersc.gov) and 400 TB on the science Tape Portal (http://portal.nersc.gov/archive/home/projects/incite11/www/20C_Reanalysis_version_2c).

To produce this reanalysis output took 2 million CPU hours per month for 6 wallclock months. Computation using the version 4 that is expected to be delivered in the timeframe of this report, at a 25-km latitude-by-longitude resolution and 91 levels in the vertical, is possible only with exascale capabilities, will need storage that is about 100 times larger than present levels, and would be expected to need 250 million CPU hours per month for 6 wallclock months. An accompanying oceanic reanalysis or coupled reanalysis would be expected to require even more hours and about double the storage.

Another important example use case of computing needs will come from the ARM site. An integrated approach is planned that will build a “4-D data cube” of reanalysis states that can provide a more complete picture of the ARM region’s atmosphere over time. Incorporating atmospheric modeling into ARM’s observational strategy will increase computing needs. The anticipated data rate of the LES models is more than 1 PB per year. LES modeling will also significantly increase the computational processing requirements for the program. One of the goals of a pilot project currently under way is to examine the costs and benefits of how much model output to store and what model settings (e.g., resolution) to use to run the model. Current estimates of the LES domain to be used are a 25-km square with 100-m grid spacing. An ensemble of roughly a half dozen simulations would be run for each shallow cloud event, with different forcing conditions used for each ensemble member. Turn-around time is approximately 1.6 times the simulated time when using 512 cores on the NERSC Edison computer with the Weather Research and Forecasting (WRF) model and a bulk microphysics representation of the clouds. In addition, plans include running one simulation per case with a spectral bin representation of the clouds, which is approximately a magnitude more expensive for a given simulation. Based on reasonable assumptions, the cost of running this simulation would be approximately 17–20 million core hours per year. Using a more generalized LES configuration with nested boundary conditions to capture spatiotemporally varying boundary conditions more accurately would raise the cost to 65–80 million core hours per year.

3.2.2 Terrestrial and Subsurface Research

3.2.2.1 Scientific Challenges and Opportunities

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. Even at that future target resolution, with the global land surface resolved as hundreds of millions of grid elements, knowledge of land processes mainly resides at subgrid scales. Process-resolving scales depend on the region of interest and the science questions being addressed. Some examples include: hillslope hydrology representing lateral surface and subsurface flows at scales of meters to tens of meters; surface inundation and associated biogeochemical dynamics connected to microtopographic variation in flat and gently sloping landscapes at scales of centimeters to meters; interactions among plants and microbial communities occurring in the rhizosphere at scales of millimeters to tens of centimeters; interactions among microbial communities and the soil physical and chemical environment localized on mineral surfaces at scales of microns to millimeters; and a host of biological processes operating at the cellular and subcellular scale in plants and microbes that interact directly with the physical and chemical environment and with significant impacts at all larger scales up to the entire globe.

It is necessary for the purposes of future Earth system prediction to take account of all of these processes across the entire global land domain. Casting forward to a computational and analytical capacity 100 times greater than present-day capacities, it is still not feasible to represent even a minimal list of the most critical land surface and subsurface processes at their native spatial scales for the global domain. Both computational capability and empirical constraints limit our ability to fully resolve the global land system in a coupled predictive framework. We foresee that many of the most critical land processes will need to be parameterized and represented implicitly in exascale computational and analytical environments. At the moment, we lack comprehensive and demonstrably robust theories allowing process knowledge to migrate effectively up in scale from process-resolving to process-parameterized. We consider this gap to be both a major scientific challenge and an important research opportunity.

One avenue along which progress is being made is to develop and exercise highly resolved multiprocess models over limited spatial and temporal domains. These models explicitly represent our finest-grained process knowledge, and provide a direct pathway for simulation initialization, calibration, and evaluation using empirical data. The highly resolved models also provide a foundation from which accurate and unbiased models may be developed at larger spatial scales, collapsing mechanistic details as reduced-form parameterizations. Although various approaches for cross-scale knowledge migration in multiscale modeling frameworks are being explored at present, generalized methods, robust performance, and quality metrics are not yet established. The implementation of process-resolving models in many geographic and functional spaces and the development of rigorous approaches for up-scale knowledge migration founded on fine-scale models represent important challenges and opportunities relevant to exascale systems.

Pushing deeper into the numerical methods underlying our most detailed process-resolving models, another emerging challenge is the implicit solvers issue.

The representation of humans, our actions, and the built environment in coupled Earth system prediction frameworks is addressed in Section 3.2.4 of this report, but we note here that terrestrial and subsurface processes are primary points of contact between humans, climate, and environmental systems. Land and human system modeling components will need to evolve together as both move toward more explicit and detailed representation.

In all of these areas, integration of multiple heterogeneous observational data streams with multiscale and multiprocess modeling frameworks to generate analysis and insight and to propose and test new hypotheses presents major challenges. The lack of comprehensively coordinated data and computational and analytical workflows poses a significant challenge to progress in terrestrial and subsurface science. While we look toward exascale computational environments, we cannot ignore the need for workflow systems well suited to the new science challenges.

3.2.2.2 Priority Research Directions

Priority research directions for terrestrial and subsurface research include these:

- Improved predictive understanding of integrated hydrology/biogeochemistry, surface/subsurface, terrestrial/aquatic systems.
 - Microbially mediated biogeochemical cycling from pore to watershed to continental scales, including reactions in streams, estuaries, and the coastal ocean.
 - Dynamic vegetation, extreme events (hurricanes, wild fires).
 - Plant-soil-microbe interactions, moving toward a more mechanistic understanding and representation of plant and microbial physiology/biology. This effort includes a focus on dynamic trait expression as a result of biology–physical system interactions.
- Coupled human-Earth system interactions, building on agent-based and goal-seeking approaches, and introducing dynamic urban and infrastructure components. The focus here is on policy relevance (high resolution, predictive impacts).

3.2.2.3 Cross-Cutting Research Directions

Cross-cutting research directions to be pursued potentially with other BSSD and CESD groups include these:

- Migration of knowledge across scales
 - Approaches to the closure of complex equation sets
 - Surrogate modeling and UQ
 - Agent-based modeling
 - Subgrid parameterization
- Frameworks/interfaces
 - Enabling of sequential, operator split, and implicit multiphysics coupling
 - Sequential and concurrent coupling across scales
- Scalable solvers for exascale architectures
 - Nonlinear and linear solvers for systems of equations arising from implicit time discretization of systems of PDEs and stochastic PDEs (SPDEs) (implicit/explicit)
 - Large, stiff systems of ODEs/differential algebraic equation (DAEs), maybe 0D, but complex dynamic coupling (e.g., biology)
- Formulations of existing algorithms that permit scalable implementations and abstract interfaces
- Model-data integration workflows and tools
 - Multiple data types (e.g., urban data)
 - Multiple scenario simulations, reduced-order modeling
- I/O and data portability — to balance *in-situ* analysis and archived simulation output for post processing

3.2.2.4 Computing Needs and Requirements

Breakout participants identified the following areas as the most significant computing needs for exascale computing:

- Workforce development/training and staffing issues.
 - Developers need a viable percentage of time and training in computer science and engineering (CSE).
 - Reward system should recognize quality CSE.
- Scalable I/O libraries that are officially supported and maintained within the DOE complex, with examples, and best practices.
- Heterogeneous systems supporting mixed compute-intensive and data-intensive computing.
- Practical assessment of speed-up and prioritization of supporting ASCR research (e.g., implicit time-evolution of stiff systems).
- A software and model development environment on HPC systems that is/feels like a laptop (docker both ways).
- Ability to “steer” the simulation mid-stream.
- Continuous integration testing/and testing at scale.
- Access to external databases.

3.2.3 Oceans and Cryospheric Research

3.2.3.1 Scientific Challenges and Opportunities

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. Over the twentieth century, the ocean system has absorbed approximately 90% of the heat trapped by greenhouse gases. In addition, the ocean has absorbed more than one-third of all anthropogenic carbon emissions. Gaining understanding and the ability to predict potential changes in the rate of ocean uptake of heat and carbon in the twenty-first century remains a grand challenge. The ocean also plays a leading role in driving changes in regional precipitation patterns through changes in structure and variability of ocean sea-surface temperature. The change in regional precipitation has a large influence on the availability of water for agriculture, cities, and energy production. Finally, the ocean contributes a long-term modulation to the Earth system, through low-frequency variability of the global ocean thermohaline circulation (THC) in response to climate perturbations (Rahmstorf 2006).

The cryosphere is undergoing the most rapid anthropogenic-driven changes within the entire Earth system. This change is occurring particularly in the Arctic, where a transition toward a summertime sea-ice-free condition is under way (Jeffries et al. 2015; Kinnard et al. 2011). The large amplitude changes in the seasonality and extent of sea ice have profound impacts on: (1) the Earth's radiation budget through the ice-albedo effect, (2) the stratification and circulation of the Arctic Ocean and their effect on the Greenland Ice Sheet and THC, and (3) the marine ecosystems within and beyond the Arctic. Arguably the most pressing challenge in appraising the impact of climate variability is the prediction of sea-level rise (SLR) and, in particular, the likelihood of abrupt SLR during the twenty-first century and beyond. It is generally expected that if abrupt SLR rise occurs in the twenty-first century, it will emanate from ocean/land-ice interaction around Antarctica. Here the relatively warm Circumpolar Deep Water enters ice cavities beneath ice shelves that exist in, for example, the Ross, Weddell, and Amundsen seas. Any increase in the rate that these warm subsurface waters come into contact with the ice shelves poses a risk to dramatically increase the rate of melting of land ice, which, in turn, could cause sea levels to rise beyond our current upper-bound estimates. Both process-based studies and the geometry of these 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. Exascale computing, along with the multiresolution modeling capability that is available with ACME, should allow for an accurate simulation of this ocean-ice interaction.

The challenge of accurately predicting the impacts of SLR has three components: (1) accurately simulating the source of additional ocean volume, (2) simulating the modes of climate variability that redistribute this additional volume in space and time, and (3) simulating the extreme events that bring the rising ocean waters onto land, where impacts on a wide range of human and ecological systems can occur.

Enhanced melting will instigate a retreat of the ice grounding line which, in turn, will in some places initiate a positive feedback leading to further increases in melting. Accurately simulating the movement of the grounding line not only requires a subkilometer level of resolution but also a realistic representation of basal hydrology at the ice-bedrock interface. The rate at which the land ice can flow into the ocean is controlled, in large part, by the friction at the land-ice/bedrock interface. This interface is lubricated by liquid water (i.e., basal hydrology) that flows through an ever-changing network of water channels. Confidence in our estimates of sea-level rise depends not only on a robust simulation of basal hydrology but also on access to observational data sets to validate the simulations.

Exacerbating the challenge imposed by the need for high-resolution outputs are the timescales involved on the ocean/land-ice interaction. Circumpolar Deep Water is formed by the mixing of North Atlantic Deep Water, originating from dense waters overflowing Greenland-Iceland-Scotland Ridge and Labrador Sea Water, with the Antarctic Bottom Water formed in and around Antarctica. The timescales involved in the formation and mixing of these water masses is hundreds to thousands of years. Within the land-ice system, the ice temperature and velocity still “remember” the Last Glacial Maxima from 25,000 years ago. A research priority focused on data assimilation is required to meet the challenge of predicting sources of 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. Within the atmosphere, routine simulation of hurricanes will be readily attainable with exascale computing resources; however, obtaining accurate simulation of hurricane frequency, track, and amplitude will likely remain a challenge. It is critical to understand that the damage produced by, for example, a Category 4 hurricane in 2030 is, in large part, a function of how the health of ecosystems at the terrestrial-aquatic interface evolves over the intervening 15 years. Changes in the patterns and frequency of inundation affect the health and resilience of landforms along the coast line. These landforms provide the primary protection to inland infrastructure and agricultural systems. An additional benefit of enhancing the modeling capability at the terrestrial-aquatic interface is the increased understanding of the carbon cycle in the coastal zone. The coastal zone stores approximately 30 Pg² of carbon in the top 1 m and processes all of the carbon moved from the land surface into the ocean.

Perhaps the easiest high-impact win that will come with access to exascale computing resources will be the routine simulation of ocean mesoscale eddies. Ocean mesoscale eddies are the ocean equivalent of weather in the atmosphere. Just like weather in the atmosphere, ocean mesoscale eddies are largely responsible for the observed ocean climate through their transport and mixing of buoyancy and momentum, and they locally control the ocean’s energetics (Hogg et al. 2015). Throughout the history of Earth system modeling, we have depended on coarse parameterizations of ocean mesoscale eddies to describe their role in the climate system. The opportunity to directly simulate mesoscale eddies should lead to greater confidence in our estimates of ocean heat and carbon uptakes and how that uptake is likely to change during this century. The opportunity to simulate energetic motions at and near the ocean model grid scale brings with it a substantial challenge in the representation of long-timescale processes, such as ocean vertical stratification. At present, ocean models that simulate mesoscale eddies typically produce excessive vertical mixing within the ocean column that results in large errors in stratification. A research priority accompanying the simulation of mesoscale eddies is an improved representation of vertical mixing in the ocean.

The ocean mixed layer (OML) is a region of well-mixed water, typically on the order of 100 m thick, that sits directly underneath the ocean-atmosphere interface (de Boyer Montegut et al. 2004). The OML mediates the transport of heat, carbon, and other tracer constituents across the ocean-atmosphere interface. Our ability to accurately simulate observed patterns and variability of OML remains a challenge. As a result, our ability to confidently assess anthropogenic-driven changes in ocean heat and carbon uptake remains somewhat dubious. A host of dynamical processes, such as mesoscale eddies, submesoscale instabilities, and Langmuir circulations, are ubiquitous in the OML and drive a significant fraction of the vertical mixing and uptake. While the scales of some of these processes, such as submesoscale and Langmuir instabilities, will not be tractable in global, uniform-resolution ocean models even with exascale computing, new multiresolution and/or dynamically

² Pg = petagram = 10¹⁵ g = 1 billion tons.

adaptive approaches available with maturing ocean models will allow for the exploration of these processes within a global Earth system modeling framework.

While resolution is important in the simulation of the OML, it may not be the only key roadblock. Possibly the likely key missing process is having an accurate representation of the broad spectrum of ocean surface waves. These surface waves drive and nonlinearly interact with the fluid instabilities that exist in the OML. On the other side of the ocean-atmosphere interface, these surface waves eject aerosols into the atmosphere boundary layer that serve as precursors to cloud formation. Full-spectrum surface wave models are mature and are a part of routine forecasting systems at use within the National Centers for Environmental Prediction (NCEP) and the Office of Naval Research (ONR) (e.g., Tolman 2002). The issue is that these models are typically more computationally expensive than the entire global ocean model that sits underneath these surface wave models. This example is yet another of the tension that exists on how to optimally allocate additional computer resources, that is, whether between increased complexity or increased resolution. Beyond driving the uptake of heat and carbon, a “residual” consequence of the set of complex processes that control the OML is the sea-surface temperature. As discussed above, changes in patterns of precipitation are strongly driven by sea-surface temperature. So improved fidelity on OML processes should not only improve our predictions of ocean uptake of heat and carbon but also our predictions of water availability.

While we fully expect that the increased model resolution afforded by exascale computing will yield substantial improvement in simulation fidelity, increased resolution in the simulation of sea-ice processes is particularly challenging. Since their creation, sea-ice models have solved a set of PDEs that rely on the assumption that sea ice behaves like a continuum fluid similar to air or water (Hunke et al. 2010). Only at the broadest spatial scales of 100 km and larger does it appear that this approximation is valid. Approaching the scales presently resolved in sea ice (~ 5 km) and definitely the scales we expect to resolve with exascale computing (~ 1 km), the continuum approximation will clearly not be valid. Patches or floes of sea ice behave much more like floating plates that coalesce and fracture as discrete elements, rather than deforming smoothly based on a continuous stress-strain relationship (Hopkins 2004; Baohui et al. 2014). Moreover, frictional loss, which typically accounts for ~95% of energy consumed in the pack during ridge building and rafting, needs a more realistic representation based on fundamental physics principles. The resulting sea ice morphology, including the ridge/keel distribution and characteristics, strongly affects sea-ice thickness, form drag, and drift, which, in turn, control dynamical coupling of the sea ice with the atmosphere and ocean. Possibly the most pressing near-term priority within the entire ocean-cryosphere enterprise is the reformulation, redesign, and rebuilding of the underlying sea-ice dynamical core based on extensive observational data, such that scales of motion from 100 m to basin-wide are faithfully represented.

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.

3.2.3.2 Priority Research Directions

Improving fidelity and confidence in rates of sea-level rise, including the likelihood of abrupt events, should be a leading research priority within the ocean-cryosphere community. This activity will have to address all sources of ocean volume, including in Antarctica, Greenland, ocean thermal expansion, melting of mountain glaciers, and extraction of nonrecharging aquifers. The priorities need to include mechanisms for the redistribution of ocean waters resulting from changes in the local gravitational force and ocean circulation, as well as from dominant modes of

climate variability. Efforts to gain a better understanding of sources of SLR and its redistribution are well under way. Increased focus on the impacts of SLR is required. The impacts of SLR will be primarily realized during extreme weather events. These extreme events need to be accurately simulated in the atmosphere and ocean systems. Exascale computing, along with multiresolution numerical methods, should enable rapid progress in our confidence in predictions of SLR. Many aspects of the SLR problem require resolutions well beyond what is tractable on today's LCFs. This situation is particularly true with respect to modeling the sources of SLR and its coastal impacts.

A foundational recasting of the governing equations for sea-ice dynamics is urgently needed. Without near-term progress in this area, the community will likely find it has computational resources to resolve sea-ice processes at subkilometer resolution but without a physically valid numerical model to carry out such simulations. Nascent ideas around discrete elements and granular particle dynamics show promise; however, more conservative lines of research need to be explored in parallel. Because of their quasi-Lagrangian nature, sea-ice models based on discrete elements should have numerous opportunities to exploit heterogeneous computing architectures. A substantial effort will be required to compare and contrast these new approaches to more traditional methods in order to fully describe the strengths and weaknesses of these next-generation models. In parallel, extensive and detailed observations are needed to advance knowledge and guide development of realistic representation of modeled floe-floe and multifloe interactions.

Activities centered around the very basics of how we produce and analyze simulation ensembles should be a research priority. While it has long been accepted that only the statistical properties of climate and climate effects are predictable, the heart of climate simulation is still built around the notion of conducting a single realization which, in and of itself, has little value in characterizing the climate or climate variability. In order to more accurately represent the statistical properties of the climate, we sometimes repeat these single simulations to build an ensemble. This naive approach has many shortcomings, not the least of which is that it runs counter to our present understanding of how to efficiently utilize exascale architectures.

At present, each ensemble member is simulated independently. Each member continuously writes its state variables to disk. Climate statistics can be produced only afterward by reading in these massive data sets to build probability distribution functions. A far more scientifically engaging approach is to embed the ensemble members into the very data structures of the Earth system model components. This approach is compatible with exascale architectures given that the climate statistics can be produced “on the fly” without need for write/read to/from disk. Having multiple realizations within a single executable could also serve as an application strategy for recovering from a certain class of system faults. Just as important, this recasting of the term “simulation” to include ensemble size would encourage a new class of *in-situ* simulation analyses, such as the identification of high-risk, low-probability climate trajectories and the breeding of particular climate pathways in order to more fully explore certain parts of the climate phase space.

Even after an exascale-compatible ensemble system is put into place, accurate prediction of probability distribution functions will require a data assimilation capability for the ocean, sea-ice, and land-ice systems. While many of these systems are less well observed than the atmosphere, recently introduced observational platforms (e.g., Argo, CryoSat-2, and SMOS) have dramatically increased both the resolution and breadth of observational data that can be used to estimate the states of the ocean and sea-ice systems. Broadly speaking, the ability to assimilate these observational data sets through ocean-cryosphere reanalysis is lacking. We should not expect to have skill in decadal climate prediction until these data assimilation capabilities are in place. Before tackling the problem of coupled system data assimilation, the capability to assimilate observational data into the individual component should be a research priority.

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).

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.

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.

3.2.4.2 Priority Research Directions

Earth system models continue to require research for long-standing problems in a number of areas, with clouds and cloud/aerosol interactions an ongoing focus. A number of other biases remain stubbornly persistent even though models have been improved (the double Intertropical Convergence Zone, Madden/Julian oscillation, El Nino/Southern oscillation, boundary current separation, etc.).

The coastal environment and terrestrial-aquatic interface present particular exascale computing challenges of high societal relevance. Coastal zones typically contain large populations and extensive infrastructure, as well as important ecosystem services, and couple terrestrial and ocean biogeochemical cycling. They are also strongly affected by future sea-level rise and storm activity, as well as continued human development. They are geographically concentrated, requiring high regional refinement (potentially down to 100 meters in some locations) to resolve important processes. Many processes relevant to the coastal zone are not yet implemented in ESMs, such as wave dynamics, shoreline erosion, and the transport of nutrients from rivers to ocean ecosystems.

Model testing and validation can pose significant challenges, particularly at high resolution or on newer architectures where bit-for-bit reproducibility may no longer be assured. Alternative testing paradigms and error checking will be needed. Ensembles can be used to generate statistical results and mitigate error, with short-term forecast ensembles an option at the highest resolution. Other techniques for propagating parameter error as part of model integration are also being developed and can be explored.

An increasing trend in Earth system modeling is the use of ensembles for computing climate statistics and exploration of changes in weather extremes. To date, independent ensemble members have been used with post-processing to analyze statistics. As data management becomes more difficult and as exascale resources permit more integrated ensembles, we will need to better explore how to perform ensemble simulations and how to utilize ensemble information to guide the simulation as it progresses. Much larger ensembles are needed, but ensembles of the highest-resolution simulations will continue to be infeasible, even with exascale computing.

Uncertainty quantification is important across almost all modeling activities, and for many applications, the target output of a simulation or set of simulations should be a probability distribution. An open question is what mix of model fidelities and simulation strategies will best achieve mission goals. Even with exascale computing, it will not be possible to run large simulation ensembles at the highest possible resolution. This is true almost by definition, because resolution can always be increased until it is no longer feasible to perform more than a few simulations. Although model improvements generally improve the model fidelity, they can also lead to increased uncertainty, given that increased complexity typically introduces more variables. High-resolution models can add complexity and introduce more uncertainties; however, they can also remove the need for approximate parameterizations in favor of directly resolving processes, and so eliminate uncertainties. Nevertheless, whatever uncertainties remain will be hard to characterize without low-fidelity coupled or high-fidelity stand-alone component models. New uncertainty quantification 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. New methods are also needed to accelerate or reduce the need for long control integrations after models are perturbed. On a related note, exascale computing can perform the very long time integrations necessary to understand past climate effects over thousands of years as an analogue for future change, to distinguish natural and human-forced variability, and to generate initial conditions for slowly varying subsystems, such as ice sheets. Another open question is how to characterize uncertainties in model structure resulting from developer choices in numerics, approximations, resolution, etc. These uncertainties are difficult to explore by traditional methods because it is currently necessary to manually change the model code to explore each new structural choice.

3.2.4.3 Cross-Cutting Research Directions

An ongoing question is how to determine when models should be coupled. Some climate science and impacts applications require high-resolution coupled modeling, whereas for other applications, off-line/stand-alone, single-component simulations would suffice. Simulations of stand-alone

climate impacts can sometimes be performed with boundary conditions derived from global models, if feedbacks to the Earth system are negligible; and stand-alone, ultra-high-resolution models, such as LES and DNS, can inform the development of scale-aware ESM parameterizations or can be used to train surrogates. In other cases, LES models might be included directly as super-parameterizations. As we continue to utilize a hierarchy of models, including global high-resolution, regional-focused, limited-domain LES/DNS, stochastic, and surrogate models, we will need a better understanding of how best to integrate these components to meet the science goals.

When the decision is made to fully couple models, we require a more rigorous understanding of that coupling. Even in current coupled models, implications of model coupling choices (e.g., lagging variables in time) are not well understood and have resulted in instabilities and other artifacts (this topic on coupling ESMs is also discussed in *Oceans and Cryospheric Research* in Section 3.2.3.3, *Cross-Cutting Research Directions*, page 79). Improved analysis and better algorithms for model coupling are required. Variable resolution and high-resolution models will require new formulations, particularly for sea ice where current continuum formulations are no longer valid. Given the expense of very-high-resolution models, it may be possible to combine low- and high-resolution configurations to address different aspects of the problem (e.g., ensembles of lower resolution for variance but high resolution to address biases).

Improved techniques for initializing models are also needed. For components like biogeochemicals, ice sheets, and the deep ocean, long timescales either require a very long spin-up (infeasible at high spatial resolutions) or data-assimilative techniques. Algorithms for accelerating spin-up or assimilating appropriate data for time periods of interest are still needed.

In addition to model spin-up, time integration is a bottleneck at high spatial resolution. As vendors look to computer parallelism for performance improvement and per-core performance stagnates, it becomes difficult to overcome explicit time-step constraints without making significant changes in algorithms for time integration.

Uncertainty quantification cuts across all aspects of model development and science. Surrogate models trained to simulation output provide a means to approximately predict the results of new simulations that are too expensive to run. Surrogate techniques could also be employed to learn scale-aware statistical parameterizations from high-resolution simulation data to embed within ESMs themselves. These may lead to inherently stochastic versions of ESMs.

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.

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.

3.2.5.2 Priority Research Directions

Session participants discussed a number of use cases in which interactions among human and Earth system processes warrant new coupled modeling applications and/or the integration of heterogeneous data across a range of spatial and temporal scales. Herein, we describe three of these broad use cases: the water-energy-land nexus, urban systems, and coastal systems.

3.2.5.2.1 Water-Energy-Land Systems

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 (DOE 2014). Climate variability is expected to reduce snowpack, increase stream temperature, alter the magnitude and seasonal water availability for hydropower production, and increase irrigation demands as air temperatures rise. Meanwhile, measures taken to reduce greenhouse gas emissions in the energy sector, such as the deployment of renewable energy technologies and biofuels, have systemwide implications for water demands and uses. Land plays a critical role in the water-energy nexus as well: irrigation is a primary water demand in many regions, and land-based climate mitigation strategies such as biofuels and afforestation compete for land resources and have potentially significant implications for irrigation demand.

3.2.5.2.2 Urban Systems

Most of the world’s population lives in cities. Thus, understanding how climate variability will affect the urban environment and its infrastructure systems is critical to developing effective strategies for water and energy resilience, as well as for predicting climatic conditions at decision-relevant scales to support a wide range of impacts, adaptation, and vulnerability assessments related to air quality and public health, emergency preparedness, and flood risk, etc. Changes in urban natural resource demands also feed up to larger scales through energy, water, and food markets. Thus, multiple scales of interaction — ranging from neighborhoods to the globe — must potentially be considered to understand how cities can and will respond to climate effects and what the broader implications of these responses will be.

3.2.5.2.3 Coastal Systems

Like cities, coastal environments are the locus of multiscale interactions among a number of physical, ecological, and human systems. Not only are many cities themselves located on coasts, but so are important infrastructure assets, such as power plants, pipelines, ports, and other transportation infrastructure. Near-coastal ocean dynamics drive mesoscale climate phenomena such as fog and sea breeze, and the dynamics of wave propagation in the coastal environment determines how severe the impacts of storm surge, tsunamis, and sea-level rise will be for human infrastructure. Estuaries and wetlands supply critical ecological services such as wave attenuation and food. Adaptive measures such as the construction of seawalls and other protective barriers can interact with coastal waves to enhance or mitigate flood risk in surrounding areas that are not protected by such structures.

3.2.5.3 Cross-Cutting Research Directions

Several cross-cutting research themes emerged related to the coupling of human and Earth system models across multiple scales of analysis.

3.2.5.3.1 Theme 1: What Is the Required Level of Detail?

For a number of applications, it is not clear what level of fine-scale heterogeneity and interactions must be represented to adequately address the science and management questions at hand. The use of statistical emulators and simplified process representations holds promise for reducing computational costs and permitting multiscale uncertainty quantification. However, for some applications, a more resolved approach may be required. As an example, some integrated assessment models represent multiple sources of heterogeneity within markets (e.g., variations in costs, preferences, policies, etc.) using a generic logit formulation (Clarke and Edmonds 1993). However, if the structure of this heterogeneity would change in response to a particular set of drivers (e.g., climate impacts, new policies), it may be necessary to explicitly represent this heterogeneity and the dynamics that influence it.

3.2.5.3.2 Theme 2: What Is the Required Degree of Model Coupling?

To address the interconnected science challenges at the interface of human and Earth systems described previously, heterogeneous and multiscale data and models must be brought together in unique configurations. The levels of coupling and scales of analysis required are open scientific questions and will most certainly differ according to the specific application. Thus, a highly flexible and modular framework for bringing together diverse data and models would efficiently serve a number of science applications and facilitate careful hypothesis testing that controls for multiple interaction terms in a systematic fashion.

However, while it is widely understood that different kinds of model coupling — in terms of coupling frequency, which variables are coupled, two-way vs. one-way coupling, etc. — are appropriate for different applications, the mapping of applications to coupling configurations is not well understood. The establishment of flexible and general frameworks for model coupling, tools for translating data across scales and among unstructured geographic representations, and the construction of modular component models that can easily interface with general coupling frameworks will facilitate this line of analysis.

3.2.5.3.3 Theme 3: The Value of Prototyping

Assessing what level of model detail, what degree of model coupling, and what kind of coupling framework is best suited to various science applications can be challenging because the construction of highly detailed models and coupling interfaces represents a significant investment of resources. Moreover, even if it is determined that statistical emulators are well suited for a certain class of problems, their development first requires the development of detailed physics-based models to generate training data for the emulators. In order to deploy resources efficiently, then, it will be important to prototype a variety of models and coupling frameworks, starting with more loose couplings, to establish their value before refining further.

3.2.5.4 Computing Needs and Requirements

The heterogeneous set of questions, processes, scales, and models used in integrated assessment and multimodel coupled assessment presents interesting and, at times, unique challenges.

3.2.5.4.1 Computing

Coupled human-Earth system models could involve both capability and capacity computing. The former is required when one of the models included is operating at the exascale (e.g., a high-resolution Earth system model). The latter is required in the quantification of uncertainty, as this may require thousands of coarse-resolution model simulations operating in parallel. The use of reduced-order models, statistical emulators, and mesh refinement are essential tools for reducing unnecessary computational complexity and enabling a larger number of scenarios to be explored efficiently.

3.2.5.4.2 Software

As the scientific challenge outlined in Section 3.2.5.1.1 involves not just multiple models but different combinations of models for different questions, a flexible, modular coupling framework is necessary. Such a framework may include new standards and interfaces establishing the means of interacting between many heterogeneous components. In addition, this framework involves both model development and coordination across modeling communities. A system for rapid prototyping and testing would also be beneficial.

3.2.5.4.3 Workflow

The move to larger-scale and more formal uncertainty quantification requires changes in workflow. Today, uncertainty quantification in integrated assessment is largely a manual effort. Scientists initiate suites of parallel simulations, choose which variables to save *ex ante*, and analyze results as part of a post-processing effort. A more integrated scenario management system is needed as the scale of these analyses expands. Such a system could help with ensemble setup and synthesize results, perhaps generating probability distribution functions automatically. Tools for adaptively searching scenario space would facilitate data assimilation and refinement of particular portions of the space where interactions affect decision-relevant outcomes. This latter concept would be the equivalent of adaptive mesh refinement within scenario space.

3.2.5.4.4 Algorithms

As these models expand both in resolution and in the number of simulations, improved analysis tools will be necessary. Algorithms that can identify the interesting aspects of outputs will be critical, as examining all results manually will no longer be tractable for scientists. Better visualization techniques could also facilitate analysis of results. Including these analysis and visualization techniques in real time could enable adaptive adjustment of the simulations as interesting features arise (akin to adaptive mesh refinement).

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.

3.2.5.4.6 Workforce

Finally, the coupled human-Earth system modeling enterprise places challenges on workforce development. Software engineers are needed to develop and maintain such a coupling infrastructure.

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.

3.2.6.1.2 *Priority Research Directions*

As described in Section 3.2.6.1.1, no matter what the future computational capabilities are, there are strong needs or requirements to establish consistent:

- Parallel *in situ* visualization and analysis tools (exploiting parallel I/O);
- Secure access of managed data and compute resources;
- Local and distance streaming of visualization and data;
- Comparative visualization and statistical analyses;
- Robust tools for regridding, reprojection, and aggregation;
- Support for unstructured grids and nongridded observational data, including geospatial formats often used for observational data sets;
- Workflow analysis and provenance management;
- Security and trust to ensure the integrity of data, workflow, and provenance; and
- Training for a scientific workforce on use of the new tools and systems to enhance productivity and therefore to effect science outcomes.

3.2.6.1.3 *Cross-Cutting Research Directions*

The climate science community has made large investments in existing extensible tools. Integrating new capabilities into tools with familiar interfaces will reduce the barriers to adoption. In addition, relatively simple interfaces are needed for other target audiences (adaptation and mitigation researchers and decision makers). One approach for simplifying data management, access, analysis, and visualization is to provide a workflow analysis and provenance environment that captures common workflows for standard data products. Such an environment will provide easy access to popular data products and document exactly how these products are produced. The entire process can also be optimized to improve the overall efficiency and reduce the data transmission load on the climate data infrastructure. With this type of environment, the broad community of researchers and decision makers (including nonresearchers) will be able to access popular data products in a highly transparent manner.

No matter what the future computational capabilities are, there are strong needs and requirements to establish consistency among the elements listed in Section 3.2.6.1.2.

3.2.6.1.4 *Computing Needs and Requirements*

Although facilities may offer different levels and types of resources, they often face common challenges. Therefore, we recommend several strategies for addressing hardware, network, and software requirements and managing facility resources for the overall good of the global computing ecosystem. It also includes implementation requirements for primary computing/data centers. Data centers within the global computing ecosystem will be categorized as either Tier 1 or Tier 2. Tier 1 sites represent the major computing and data facilities in the world. Tier 2 sites are all other facilities with lesser capabilities. Together, the Tier 1 and 2 sites throughout the global computing ecosystem make up a much larger federation.

Tier 1 sites are centers that can allocate the necessary hardware, network, and staff resources to act as a major continental hub for computing and distributing data holdings. Tier 1 sites are expected to run the full suite of required services for computing, data, and user management, which can be used to support their own activities and those of Tier 2 sites. Tier 1 sites have these minimal resource requirements:

- **Storage:** a multi-petabyte storage system to serve the external community;
- **Front-end server(s):** a minimum number of platforms running designated operating systems to interface with the user community;
- **Metadata index:** a set of index engines and corresponding back-end indexes to publish and search the data, including a master/slave configuration for local holdings, plus replica shards for all other nodes in the federation;
- **Compute:** shared compute clusters/servers for remote data reduction, derived data, large-scale data manipulation and visualization, and data streaming; also connects to HPCs to run Earth system models;
- **Identity provider:** secure Single Sign-On access control attributes and user authentication X.509 certificates;
- **Node manager:** a “super-node” instance of the node manager (peer-to-peer deployment for resiliency);
- **Network:** connects to the wide-area network at futuristic gigabits per second with a minimum of four designated GridFTP/Globus-enabled servers, and a perfSONAR instance for network performance measurement in a Science DMZ-like infrastructure (Globus 2017; Hanemann et al. 2005; Dart et al. 2013);
- **Software:** a stack that uses the latest software (today, the software includes Postgres, Apache httpd, Tomcat, Solr, THREDDS Data Server [TDS]), LAS, and UV-CDAT); and
- **Staff:** a designated, full-time administrator to help ensure that the system is always up.

Tier 2 sites typically have fewer physical or staff resources that can be allocated to interactions than Tier 1 sites; however, they still need to distribute a certain (possibly significant) amount of data to the climate community. Tier 2 sites are encouraged to leverage some of the services supported by Tier 1 sites, such as a metadata index and identity provider, and focus instead on supporting local services for data download and possibly analysis. Tier 2 sites have these minimal resources requirements:

- **Storage:** adequate data storage to serve their data to their external climate community;
- **Front-end Linux server(s):** a platform running the designated operating system for front-end user interfacing;
- **Software:** The implementation of the latest software stack (see Tier 1); and
- **Staff:** a designated, part-time ESGF administrator.

Adequate hardware, software, facilities, and materials must be provided to Tier 1 and Tier 2 sites for simple installation and community project use. Moreover, the basic aspects/foundation of the enterprise system — composed of hardware, networks, and software — work together harmoniously to form the fundamental overall success of the global computing ecosystem and the projects that it supports. Choosing the appropriate hardware and operating system is key for realizing consistent and sustainable operations.

3.2.6.2 Observational Data Processing: Retrieval Algorithms, Instrument Network Simulation

3.2.6.2.1 Scientific Challenges and Opportunities

BER’s Earth system modeling activities collaborate strongly with its Terrestrial Ecosystem Science (TES) and Atmospheric Radiation Measurement programs to build high-fidelity representations of physical and biogeochemical processes in global modeling frameworks. The process-resolving

observational and modeling activities are essential to the improvement of the broader global simulation capabilities. These observational programs — such as the TES program’s SPRUCE (Spruce and Peatland Responses Under Climatic and Environmental Change), Next-Generation Ecosystem Experiments (NGEE)-Arctic, NGEE-Tropics, and Ameriflux, along with the ARM in-situ and remote sensing observations of clouds, aerosols, radiation, and atmospheric boundary-layer processes — continue to reduce uncertainty of key physical and biophysical components of the Earth system model.

The majority of the world’s population lives in urbanized centers, near major ports of entry (airports, harbors) and developed coastal areas. It is important that in these largely heterogeneous areas, we understand the land-atmospheric energy and moisture exchanges, the urban boundary layer turbulent flow, its mixing with the free atmosphere, and other processes that control the meteorology and climate of these locations. Such areas are future examples of where siting of integrated high-resolution field observations is envisioned (BERAC 2013). These field laboratories and their counterpart modeling frameworks are required to: (1) advance our current understanding about the structure, variability, and interactions of complex physical processes that act across a wide range of scales; and (2) improve our ability to predict changes and evaluate the impact of decision making on human-nature interactions.

In climate research, the DOE Atmospheric Radiation Measurement program’s Southern Great Plains (SGP) megasite provides a good example of an integrated field laboratory. The DOE ARM Climate Research Facility provides the climate research community with strategically located *in situ* and remote sensing observatories designed to improve the understanding and representation, in Earth system models, of clouds and aerosols as well as their interactions and coupling with the Earth’s surface. Past observational strategies at the ARM facilities were limited to a profiling view of the atmosphere with a set of sensors that were capable of observing only parts of the atmospheric system. The new ARM radar facilities are capable of providing holistic, multidimensional observations of aerosols, radiations, clouds, and precipitation.

The new ARM program facilities operate radar sensors at five different frequencies covering a wide range of scattering mechanisms, thus improving the information content of collocated, multiwavelength observations. This improvement is particularly true in the vertical column, providing enhanced, calibrated, multiparametric measurements of clouds and precipitation. In addition, the ARM radar facilities feature scanning polarimetric Doppler radar observations over a wide range of climatological conditions. The main objective is to provide a revolutionary characterization of a complete volume of the cloudy atmosphere, over a long period of time, to act as a natural laboratory for the modeling community for the testing of their models and for improving the parameterizations of clouds in Earth system models. The configurable ARM radar facilities employ adaptive scanning strategies that enable focused experiments to study critical aspects of the water cycle at a range of spatial scales from the inner scale (30–50 m) to the outer scale (50–100 km).

In parallel with preparations for the SGP megasite, the ARM program recently kicked off a pilot study (LES ARM Symbiotic Simulation and Observation [LASSO]) that aims to provide a framework for performing routine high-resolution modeling and model-data integration at the SGP site. In addition to the SGP, it is conceivable that similar activities will take place at other ARM sites in the future. Once developed for this application, it is anticipated that the modeling framework will be applied to other, more complex scientific problems.

3.2.6.2.2 Priority Research Directions

Three areas of research have been identified as critical in order to achieve the science mission of observational data processing and its use for improving high-resolution model development across

the land-atmosphere and cloud-environment interface. These are: development of an instrument network simulator, use of data assimilation techniques, and development of machine learning algorithms to exploit the information content of the observations.

The data cubes from observational megasites are generated from a wide range of scanning sensors (radars in particular). The technical specifications of these sensors impose constraints with respect to their ability to sample the temporal/spatial structure of the atmospheric state within the cube and also with respect to their ability to detect (sensitivity constraint). It is important that these constraints are also applied to the high-resolution model output. In addition, converting the model output to synthetic (forward) observations (e.g., radar reflectivity) is often regarded as a less uncertain path for comparing observations and models as compared to the option of performing highly unconstrained retrievals (inversions). This step is crucial in our effort to combine high-resolution observations and models and thus compare apples with apples. This “forward modeling” is performed using instrument simulators that are forward models in that they transform numerical model output to the observation space and thus offer an alternative path to retrievals for comparing numerical model output and observations. Although not new, instrument-forward simulators have recently attracted interest as a promising method to compare model output and observations. In particular, a comprehensive radar simulator is needed that is able to emulate all types of ARM radars and is capable of interfacing with cloud-resolving models (CRMs) and LESs that employ a wide range of microphysical schemes is needed. Research should be toward the development of radar/lidar simulators that can emulate the instrument network configuration at observational sites.

Data assimilation has traditionally been performed at operational weather forecasting centers as a way to conduct model initialization. The value of data assimilation, however, goes well beyond that. Data assimilation can be considered as an interpolation technique to blend sparse observations with model results. There is a lot of experience in the community with assimilating heterogeneous data into coupled atmosphere-ocean-land models. The one consistent thing that has emerged from these experiences is the huge demand that data assimilation places on computing resources. Some of the challenges include the following: (1) organization of data structures from different sources and for processes across vastly different temporal and spatial scales with rapid throughput, (2) interfaces for models to take observations (forward operators–simulators), and (3) algorithms to conduct data assimilation with coupled systems of vastly different scales. The community will need to be engaged to identify the values of ESM data assimilation, the end users of the products, the availability and desirability of data, and the algorithms to assimilate data in coupled models as complex as ESM. Data assimilation techniques for assimilating at high resolution and physical properties that traditionally have not been considered before (e.g., cloud field properties, mass flux) need to be explored.

Finally, development of new algorithms is recommended to unlock the scientific information content of the multidimensional data sets. In particular, the use of machine learning algorithms in the form of applied neural networks has already demonstrated the ability to extract information from radar Doppler spectra. In addition, on the machine learning front, we see potential in evolutionary computing (i.e., genetic algorithms, genetic programming, etc.), an area of continuing active research, to lead to discovery of new atmospheric process parameterizations that will be valuable to Earth system modeling in an exascale computing ecosystem. Devising data synthesis techniques is another area where new algorithm capabilities are needed. Tens of instruments observing different, limited parts of the atmospheric state generate the need to merge their information into a best 3D/4D estimate of the atmospheric state.

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.

3.2.6.2.4 *Computing Needs and Requirements*

The sheer volume of observations from sensor networks and model output and the need to derive understanding from these sources suggest that extant computing resources are required. A number of new instruments or upgraded instruments were added to the facilities during the American Recovery and Reinvestment Act; these instruments provide new information on some key gaps in our measurements of cloud microphysical processes and aerosol composition. However, many of these instruments (scanning and profiling multiparametric radars, high spectral resolution radiometers, and aerosol-observing system measurements) produce very large and complex data streams. The volume of the data recorded by the new facilities will be 10 times higher (2.5 PB/yr) than the volume of the data recorded in 2010 (e.g., 250 TB/yr). If we include the high-resolution model output and higher-order, value-added products, the volume of data (raw, value-added, and model output) could reach 5 PB per year.

An intelligent data inventory tool that marks/flags the observational data not only for data quality but also for atmospheric scene content should be considered. This tool should reach beyond the development of static plots and perhaps needs to operate at the moment of data collection (e.g., tagging photos/images of the sky, hydrometeors). This capability will help the user community access and process the data that are of most relevance to their applications.

A concurrent, significant increase in the I/O compared to today’s models is required (~2 orders of magnitude). Already there are examples across the ARM program (e.g., processing of ARM radar Doppler spectra) where GPU processing is capable of processing 1 TB of observations per hour, and the I/O is the real bottleneck. The I/O needs to evolve (develop) to enable processing of larger volumes of data.

In addition to data storage and I/O requirements, there is a need for developing new statistical and analytic methods for handling these large data volumes and for extracting useful information from them. While additional research is required to determine the class of algorithms needed (e.g., machine learning, genetic algorithms, neural networks) to tackle the observational data

sets, it is clear that these efforts will benefit from collaboration between data scientists, computer scientists, and others. Partnering with other communities with relevant expertise in data analytics — for example, possibly other DOE offices, DOE institutional computing centers, DOE labs with data analytics capabilities, or outside organizations like the National Optical Astronomy Observatory (NOAO) — should be pursued.

The academic and broader research community needs to have access to data and data diagnostics. Traditionally, data access implies having the ability to download the data; however, as the data set volumes grow, the data sets are becoming too large for science end users to download easily. Furthermore, algorithm development and verification are slowed by the amount of time it takes to transfer large data sets to external machines for processing. In several cases, computational facilities sitting next to large data storage facilities have been put together to enable users to log in and perform their computations locally. However, the computing environment, software stack, and accessibility of compute resources are often a barrier to performing server-side data analysis. It is clear that a different data service model needs to be explored. Lessons learned from other communities such as high-energy physics, astrophysics, and genomics that have already addressed such issues should be incorporated in the proposed paradigm shift.

One approach is to enable users to bring their codes to the data rather than having the data shipped to the users. It is desirable that the codes/operators used by principal investigators (PIs) be available to the whole user community. The inclusion of PI algorithms at the data centers enhances the ability of the broader user community to access and derive analytics from the large data sets; however, several issues will need to be resolved. University scientists tend to resist sharing their algorithms. Fostering change in this area will require a large shift from the current paradigm. There is a need for a reward system for people who develop algorithms, etc., and then provide them to the community. Furthermore, intellectual property and licensing issues will need to be resolved. Finally, coding standards will need to be developed (see ARM’s Application Programming Interface [DOE-ARM 2016]).

3.2.6.3 Uncertainty Quantification and Analytics of Observations and Model Results: Uncertainty Quantification, Statistics, Emulators, and Analytics of Model-Observation Comparisons

3.2.6.3.1 Scientific Challenges and Opportunities

The need for uncertainty quantification (UQ) and rigorous model-observation comparison manifests in many aspects of Earth system modeling. On a high level, this is driven by the need to assess the predictive fidelity of ESMs as well as the use of scenario analysis to inform decision makers in the area of 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.

Particular challenges in UQ analyses today are the high dimensionality of the uncertainty input spaces and the associated cost of running perturbed parameter ensembles that adequately cover this uncertainty input space. Further challenges are the need to characterize the uncertainty properly in the available observational data and the question of where observational data are most useful for model calibration. Challenges in emerging UQ activities are the need to assess uncertainties in coupled models effectively and to address the presence and impact of structural uncertainties (model error) in ESM model components and couplings.

An overarching demand of UQ is to provide an “uncertainty management” framework. In particular, climate scientists are looking for ways to assess what the key uncertainties are, and which of these uncertainties can be reduced through more observational data or more refined modeling efforts. A closely related question is where these observational and modeling efforts

should be focused to have the highest impact. On the flip side, it is important for decision makers to assess which uncertainties are very difficult to reduce, and to provide accurate assessments for the impact of these uncertainties on the model predictions or scenario analysis outcomes.

Many of the challenges faced by UQ are tightly connected to the high computational cost of UQ analyses, in particular, the cost of running large ESM or IAM ensembles. The availability of exascale computing resources will help to offset this cost, but only partly so. A one hundred- or even one thousandfold increase in computational power by itself will not be enough for uncertainty analyses of fully coupled systems. Moreover, the expected increase in system faults as computational component voltages are reduced, the need for very high on-node parallel efficiency, and the high cost of data movement and storage will complicate the use of brute-force ensembles for UQ. More effective approaches for running ensembles will be required.

3.2.6.3.2 Priority Research Directions

Many approaches are available for making existing UQ analyses more effective, so they can be applied to larger-scale problems. Also, many algorithmic developments are enabling new types of UQ analyses that will deliver significantly enhanced information about the predictive fidelity of ESMs. Key research directions are as follows and are further detailed in this section:

1. Optimal choice of ensemble members.
2. Effective use of exascale computing resources for ensemble runs.
3. Comprehensive, multi-fidelity, hierarchical management of UQ workflow.
4. Determination of structural uncertainty and mesh discretization errors.

A first approach to keep the cost of ensembles at bay is to optimally choose the number and the configuration of ensemble members. Very often, ensembles serve to build computationally cheap surrogates that can be used in lieu of the full model in a variety of analyses such as sensitivity analysis or calibration. There is especially a need for more effective ways to build multivariate, temporal surrogate models. Given the high dimensionality of the uncertainty input space, effective ways to construct high-dimensional surrogates are needed. Adaptive sampling approaches are particularly relevant here, as well as methods to identify effective, reduced dimensional manifolds or subspaces within the high-dimensional computational data. Machine learning methods can be brought to bear on this problem as well. Further multifidelity methods can reduce the computational cost by using an appropriate mix of low-, medium-, and high-fidelity runs to create effective multilevel surrogates. On a higher level of abstraction, adaptive methods that incorporate model-to-observation comparisons can be effective in identifying key areas in the parameter space where surrogates need to be most accurate.

Better ways to extract information from each ensemble member are also needed. For example, a mini ensemble of model runs with perturbed initial conditions can be more effective for reducing the natural climate variability in the model outputs than one very long run. Novel and faster approaches are needed to determine the change in equilibrium model outputs in response to changes in model parameters or boundary conditions.

A second approach is the more effective use of exascale computing resources for ensemble simulations using algorithmic enhancements in the forward models. In the context of optimization and calibration efforts, this approach includes the use of adjoints or Hessians, more advanced parallel-chain Markov Chain Monte Carlo (MCMC) methods, or more effective ensemble Kalman filters. In the context of regular ensembles, embedded ensembles, which use a more intrusive or hybrid (rather than a purely nonintrusive sampling) approach to ensembles, can dramatically improve the effective use of exascale computing resources. Both the addition of adjoints/Hessians

and the use of embedded ensembles require enhancements to the forward models, but the payoff of this higher algorithmic complexity is a much better use of on-node, fine-grained parallelism on exascale architectures.

While advances in the first two research directions can deliver tremendous gains in computational efficiency, a more comprehensive approach to manage the computational workflow for UQ is required to obtain optimal use from the available computational resources. This usage is particularly important in the context of fully coupled model analyses. Novel approaches are needed to adaptively determine the appropriate mix of high- and low-fidelity model runs, as well as single or coupled model runs, to arrive at an integrated, hierarchical workflow for UQ that makes optimal use of the available computational resources and observational data and best informs the fully coupled models. Multiple elements will be required to enable this capability, such as approaches to fuse information from different sources, ways to relate single-component runs to coupled model predictions, methods to improve comparisons between model outputs and observational data, well-characterized uncertainty information in the observational data, effective ways to identify the key contributors to the predictive fidelity of each model output, and optimal design-of-experiment (DoE) methods to determine where additional computational resources or observational data are most effective in reducing the key sources of uncertainty.

To comprehensively incorporate all relevant sources of uncertainty, the fourth research direction is to develop effective methods for assessing the contributions of structural uncertainty (model error) and mesh discretization errors to the model predictive fidelity. While very challenging to assess, these sources of uncertainty often are at least as important as parametric uncertainty. Methods to address structural uncertainty and mesh discretization errors are synergistic with multifidelity, multimodel approaches. In fact, these methods can provide the glue to take a collection of runs of various fidelities and different model formulations, of both a single-component and coupled nature, and relate these runs to each other to inform a comprehensive hierarchical UQ workflow, as discussed in the third research direction.

Methods for assessing structural uncertainty and mesh discretization errors are at the bleeding edge, but promising methods are emerging that use randomness in select model parameters to capture structural uncertainty, or stochastic processes to capture the effect of mesh discretization errors. Closely related to this development is the use of reduced-order models with free parameters to represent structural differences between models.

3.2.6.3.3 Cross-Cutting Research Directions

Uncertainty quantification as a whole is a cross-cutting research activity.

3.2.6.3.4 Computing Needs and Requirements

A variety of computational ecosystem aspects can assist in achieving effective UQ workflows. Based on the preceding sections, examples include:

- Support for effectively running large ensembles of runs in an automated fashion (job submission and monitoring, automated data extraction, and processing).
- Task-based parallelism frameworks for making effective use of exascale computing architectures.
- Methods to provide resilience to exascale computing runs with minimal user intervention: either make the runs oblivious to hard and soft faults or provide automated ways to monitor and resubmit ensemble members that have been affected by system faults.

3.2.6.4 Model Development and Testbeds

3.2.6.4.1 Scientific Challenges and Opportunities

A model development testbed is a systematic, automatic framework involving a combination of model and observations and used to understand physical processes and to evaluate and identify sources of error in a model during its development (DOE-CESD 2014). 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.

Unlike during the production or application phases (performing climate predictions or scenario simulations), an Earth system model undergoes constant evaluations and modifications throughout its development phase. Therefore, HPC resources are especially crucial in this phase. Although each model testbed has its own unique features and targeted science questions, the common goal is to have a fast turnaround cycle for a typical simulation test case, which includes designing model experiments, configuring model parameters/codes, performing simulations, and evaluating the results. The same workflow will repeat many times until a model version is deemed satisfactory by the model developers. An HPC resource of increased size will accelerate the process of performing simulations and evaluating results, which is always a desirable demand for the developers.

For current DOE-funded model testbeds — the Cloud-Associated Parameterizations Testbed (CAPT) (Phillips et al. 2004; Ma et al. 2015), the Regionally Refined Model (RRM) (Guba et al. 2014; Zarzycki et al. 2015), and the Aerosol Modeling Testbed (AMT) (Fast et al. 2011) — the major focuses are on evaluating and improving the representations of cloud and aerosol processes in the models. The unique feature of CAPT is to use a numerical weather prediction (NWP) method to evaluate cloud-associated processes (cloud, precipitation, or radiation) in Earth system models by performing series of computationally inexpensive, short-term (2- to 5-day-long) hindcasts. The unique feature of RRM is to use a global model with a regional high-resolution refinement to study fine-scale cloud features for a particular area without the same computational burden from a globally uniform, high-resolution model. The unique feature of AMT is to evaluate aerosol process modules in the models with field measurements using a regional model (weather research and forecasting [WRF]).

The common challenge of these testbeds in terms of computational needs is the large amount of simulation output with very high temporal frequencies in order to study detailed cloud-associated processes. Another challenge is the efficiency of post-processing procedures for model output, as well as application of proper metrics and diagnostics packages. Individual testbeds may have additional challenges. For example, CAPT requires additional procedures (e.g., nudging or data assimilation) to generate initial conditions for hindcast experiments. Application of data assimilation to CAPT to generate initial conditions usually requires substantial computational resources, which is usually impractical for routine parameterizations testing with the current HPC resources.

Unique opportunities that would come about with increased computational capability for current DOE model testbeds include efforts to:

1. Study detailed cloud and aerosol processes or extreme precipitation events with high-resolution (globally $\frac{1}{4}$ - or $\frac{1}{8}$ -degree) CAPT simulations or with RRM under various parameterization configurations.
2. Study climate model variability attributable to cloud processes with multiyear hindcasts.

3. Test computationally expensive parameterizations (e.g., bin microphysics) in high-resolution model simulations or test unified or scale-aware parameterizations with the current RRM or a future nonhydrostatic version of RRM.
4. Incorporate more frequent use of instrument simulators.

3.2.6.4.2 *Priority Research Directions*

The following three areas of priority research directions will be most crucial for improving model development capability for the next generation of DOE’s ACME model.

- The first area is to develop the capability of initialized and coupled hindcasts for cloud parameterization testing or coupled model bias studies in high-resolution simulations for CAPT. Because cloud parameterizations must perform well in a fully coupled climate system, it is logical and desirable to perform parameterization testing with a coupled Earth system model.
- The second area is to apply the nonhydrostatic RRM framework for very-high-resolution cloud processes, cloud-aerosol interactions, and extreme events studies for a targeted region, such as the continental United States domain given that the global cloud-resolving model will still be too expensive computationally for climate studies.
- The third area is to apply UQ techniques for routine and systematic model physics parameters estimates or tuning in all model testbeds.

All priority research directions require substantial computation resources and storage space to store for high-resolution model output. A well-developed workflow and efficient post-processing or even in-line processing programs for providing metrics or diagnostics are also necessary in order to meet the requirement of efficient and fast-turnaround model development cycles.

3.2.6.4.3 *Cross-Cutting Research Directions*

Uncertainty quantification and data assimilation are two useful techniques for climate applications, especially for model testbeds.

As indicated in the priority research directions, ACME model development will benefit from development of the capability to perform initialized, coupled model hindcasts for cloud parameterization testing or coupled model bias studies in high-resolution simulations for CAPT. The current method of initialization for CAPT (Ma et al. 2015) does not involve data assimilation capability. For fully coupled hindcast studies, a data assimilation system (e.g., ensemble Kalman filter) is necessary, especially for modeling initial oceanic conditions.

Another relevant application is to use UQ for model tuning or accounting for structural and discretization error. Model tuning or exploring parameter sensitivity is a common practice during the model development phase. The traditional way of tuning free parameters in a model is usually time consuming because developers often have to change each parameter one at a time. UQ provides several techniques to find an ultimate set of parameters in a systematic way. Many studies have demonstrated that UQ is an innovative way to obtain a set of free parameters for a model’s physical parameterizations for performing better Earth system simulations than the traditional method (e.g., Qian et al. 2015; Boyle et al. 2015).

3.2.6.4.4 *Computing Needs and Requirements*

All priority research directions mentioned in Section 3.2.6.4.3 require substantial computational resources. We expect that it is essential to have more than tens of thousands of CPU or GPU cores of computational resources, and storage space with hundreds or thousands of terabytes for scratch memory or long-term storage for high-resolution model output.

Second, a well-developed workflow and efficient post-processing programs or in-line processing of performance metrics and diagnostics are also necessary in order to meet the requirement of efficient and fast-turnaround model development cycles.

Finally, dedicated and reliable computational resources for model testbeds are key to successful and efficient model development and evaluation.

3.2.7 Transforming Science through Exascale Capabilities: Algorithms and Computational Science

3.2.7.1 Algorithms

3.2.7.1.1 Scientific Challenges and Opportunities

The challenges facing Earth system models as they approach exascale are broad and diverse, and so the necessary algorithms to enable their use are as broad and diverse as the problems they address. Scientists have outlined a suite of requirements for Earth system models that should be addressed with new and expanded algorithms; and these requirements span model testing, integration with multiple scales, coupling, analysis, and understanding.

Model testing at the exascale is ever more critical because stakeholders need strong confidence that the models are fully tested and evaluated for accuracy and relevance for the problem(s) being solved. Algorithms for testing will extend testing from the current methodology of running benchmark cases to running ensembles and generating statistics to measure diversions from acceptable levels both as part of test cases and during production. In the case of production-level model execution, we will need to deploy algorithms capable of identifying issues and attributes of model behavior *in situ*. Because exascale machines are expected to experience more frequent faults, this is a key capability for Earth system models.

Models will have increased resolution and depth of feature representation, with orders of magnitude more degrees of freedom. The growth will come from increased resolution in space and time (in terms of data creation) in each component; new parameterizations will have new variables and moments to treat subgrid detail. In addition, higher-order and stochastic methods may be included, and they will create more data to represent a given feature. Ensembles are widely expected to be used more frequently for testing and analysis, both during and after model execution, which will create many more points to represent an event. Algorithms will need to be developed to handle these extra representations and provide a robust and timely solution.

Connected to the increase in degrees of freedom for Earth system models is the increase in the range of resolved scales and complexity in Earth system models that will be needed to solve the next generation of climate research problems. As a result, in some cases, new and more complex equations are needed to represent climate behavior at the more refined scales. For example, we will use equations to solve both nonhydrostatic flow for the atmosphere and to solve Stokes flow for ice sheets, at least in some regions. In some cases, new representations of the finer resolutions are needed, for example, using Lagrangian flow to represent features in sea ice and stochastic representations of cloud physics rather than columnar representation. With more resolution and more complete equations, more climate behavior at small but critical scales for understanding impacts and vulnerabilities can be resolved, while still maintaining a good representation of global and century-scale behavior; thus, algorithms that can capture and solve for wider scales and complexity — and equation types — are needed.

In some cases, scientists will need to execute new models and perform novel analysis and will require algorithms that can handle these additional requirements. Two examples include machine learning to allow scientists to flag individual events and behaviors within large data sets and scale-aware algorithms of an as-yet-to-be-determined type that can handle “semi-resolved” behaviors.

3.2.7.1.2 Priority Research Directions

Researchers have identified algorithm development in a number of areas as being critical to achieving their science goals in the exascale era, and in most cases, these algorithms span what the Earth system modeling community is doing already. However, they could be developed further to

cross more scales and parameter space and provide more resilience and performance across multiple climate components when applied to increasingly complex, stiff, and coupled equations. These needs were identified as a main priority because of the expense and complexity of the resolved flow field of all the main climate components (i.e., atmosphere, surface and subsurface land, ocean, sea, and land ice). The desire to perform more robust and converged coupling of intra- and intermodel coupling was also identified as a key goal; thus, algorithms that can handle coupling for large problem sizes, scales, and diverse mappings are needed. Performing analyses of model simulation data — both *in situ* and as a post-processing step — was also highlighted as a key direction for exascale-supported climate understanding. Data assimilation and the design of methods to process data to match observed data (e.g., satellite swaths) are also key drivers for new algorithm work.

The methods to address these problems include the time integration and steady-state solution of PDEs using nonlinear and linear methods; linear solutions (of linear problems and as updates to nonlinear iterations) using matrix operations, multilevel methods, and eigensolvers; methods for interpolation for purposes of data assimilation; remapping; initialization; sampling methods to capture events and conditions of climate behavior and create more resolution-aware parameterizations; and optimization methods to quantify uncertainties and define sensitivities. These methods are challenging when implemented on existing architectures, and so extending to a diversity of as-yet-to-be-defined exascale architectures will require some testbed development of the algorithms applied to specific Earth system model examples. Access is needed to computing that enables algorithms that provide efficient execution and, separately (and concurrently), analysis of model output.

In some cases, specific types of new algorithms have been called out as areas of research needed to achieve exascale simulation. With nonhydrostatic equations in the dynamical core as a target, implicit algorithms in the vertical and possibly the horizontal will be needed. Explicit methods that enable larger time steps are also relevant for these applications. Preconditioners to the implicit algorithms and accelerants (e.g., Andersen Acceleration) that are specifically tailored for the configuration should be defined at runtime based on the configuration of the model. As for algorithms that address couplings between Earth system components (and also within components), a simple Picard and then more advanced implicit algorithms are needed. The development, implementation, and proper configuration on exascale machines are areas of research to pursue. As for interoperability, all algorithm development would benefit from the ability to cross languages to take advantage of libraries and other language-specific tasks, like directives.

3.2.7.1.3 Cross-Cutting Research Directions

The components of DOE's ACME Earth system model each solve different equations and perform suites of analysis relevant for their efforts, and the coupling of these components and the corresponding analysis is an additional mode of solution needed to address the large-scale climate system questions. Coordinating these efforts as a community will allow the developments to be leveraged, and the interaction with the facilities can benefit from coordinated algorithm development, as well. Using algorithms to uncover undesirable behavior in the model, either during post-processing or model execution, is needed across all aspects of Earth system modeling and in all components. These tests need to account for the potential lack of reproducibility of next-generation and exascale systems.

3.2.7.1.4 Computing Needs and Requirements

Given the new developments required for Earth system models in the next 5–10 years, identifying and in some cases developing new algorithms will be necessary. These new algorithms need to have sufficient and demonstrated accuracy and to be accessible through quality, flexible, and modular

software that allows easy, continued improvement and enables multiple languages to interoperate seamlessly across data types. The algorithms need to provide portable, resilient, and adaptable behavior, enabling use of the machines with maximum effectiveness on a diversity of exascale architectures. Maximum effectiveness is defined here as allowing the fastest and most robust speed to solution by using dense data structures, providing a maximum reuse of data, and requiring minimal communication, although algorithms with the most flexibility to adjust these parameters across architecture types are the priority. Dedicated staff with the expertise across all of these aspects are needed for successful development and implementation of these algorithms within Earth system models.

Figure 3-5 shows the key positioning of algorithms along the computing ecosystem trajectory. Algorithms, when accessed through calls to libraries within Earth system models, can be selected and optimized at runtime and take advantage of ongoing hardware developments.

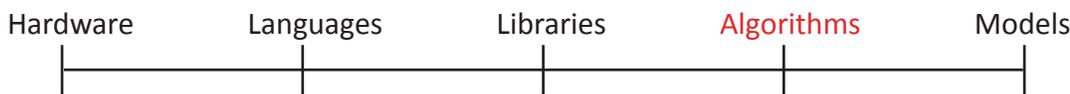


Figure 3-5. Schematic showing the interaction of algorithms among models in which they are used and the facilities ("hardware" above), languages, and libraries that define how they are designed and optimized.

In addition to climate-specific algorithms to produce relevant tests and performance metrics, exascale facilities should provide generic algorithms optimized for the target machines to enable testing to capture machine issues during execution. Some algorithms, specifically those to be used for analysis, such as to perform *in-situ* data reduction to lower output, will need bursts of power and memory at some stages of the calculation.

If Earth system modelers can successfully implement the required algorithms, using advances from exascale facilities, Earth system models will run more efficiently and with better fidelity, which, in turn, enables more science and more interest in computer allocations by analysis teams. Thus, it will be necessary to provide optimal algorithms to handle coupled, high-resolution models executed in ensemble mode. If successful, these models will require at least an order of magnitude of computer time for execution beyond what is needed for development.

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.

Table 3-2 highlights the advantages and disadvantages of the various programming models.

| Table 3-2. Advantages and Disadvantages of the Various Programming Models | | |
|---|--|--|
| Programming Model | Advantages | Disadvantages |
| MPI+OpenMP (omp) | <ul style="list-style-type: none"> ■ Portable across all CPU architectures ■ Graceful transition (tuning) between processes and threads on a node ■ New SIMD support helps vectorization ■ Support provided for tree parallelism (tasks) | <ul style="list-style-type: none"> ■ Imperative, user-intensive approach; must explicitly thread everything to avoid Amdahl bottlenecks ■ Not portable to GPUs (requires disjoint set of omp target/map clauses to use GPUs) ■ Anything more than BLAS1 requires tuning ■ Quality of OpenMP runtime can vary greatly among compilers (tasks, nested, overheads) |
| MPI+OpenACC | <ul style="list-style-type: none"> ■ Runs on GPUs ■ Ideally declarative (#pragma acc kernels) | <ul style="list-style-type: none"> ■ In reality, imperative, user-intensive approach that requires substantial tuning ■ Only recently portable to CPUs ■ Requires OpenMP to exploit both CPUs and GPUs ■ Standard lags vendors (CUDA) by a few years; compilers lag the standard; not all SC centers have a compliant compiler ■ No support provided for tree parallelism ■ Strongly motivates one process per GPU |
| Domain-Specific Embedded Languages (DSeL) | <ul style="list-style-type: none"> ■ Embedded DSL == legal code that can be compiled with any C++ compiler ■ Can convey high-level knowledge (stencils, meshes, matrices) into code generation for optimization | <ul style="list-style-type: none"> ■ Requires source-to-source compiler support for optimization, threading, GPU-acceleration (e.g., includes domain-specific knowledge in ROSE) |
| C++ Templates (e.g., Kokkos) | <ul style="list-style-type: none"> ■ Declarative approach (parallelism is specified) ■ Kokkos hides complexity of OpenMP or CUDA behind parallel abstractions ■ May eventually be standardized in C++ and supplant OpenMP, OpenACC, and Kokkos-specific constructs | <ul style="list-style-type: none"> ■ Kokkos has been demonstrated on coarse-grained, bandwidth-intensive operations (easy to amortize overheads/inefficiencies behind the memory wall) ■ User-intensive, explicit management for hierarchical memory (vendors are exploring automated solutions for hardware/software [HW/SW]) |
| Partitioned Global Address Space (PGAS) (e.g., UPC++, CAF, etc.) | <ul style="list-style-type: none"> ■ Language-based instead of library-based like MPI-3 ■ Natural solution for global data structures with irregular access patterns ■ More productive solution for irregular data reorganizations (AMR, couplers, etc.) ■ Interoperable with OpenMP and MPI ■ Active Message/Remote Procedure Call support | <ul style="list-style-type: none"> ■ Synchronization and PGAS consistency can be tricky to master in a scalable fashion ■ Little/no benefit for ghost exchanges in PDEs (asymptotically two-sided) |
| Task-based (e.g., Legion) | <ul style="list-style-type: none"> ■ Provides standardized mechanisms for distributed, task-based execution (alternate to static MPI decomposition) ■ Potentially expresses more parallelism | <ul style="list-style-type: none"> ■ Does not solve problems for users; gives them tools they can use to solve problem (e.g., user must write a mapper, users must still write OpenACC leaf kernels) ■ Performance for operations with complex, fine-grained data locality challenges has not been demonstrated |

3.2.7.2.2 *Factors Affecting Mutually Beneficial Vendor-ASCR-BER Collaborations*

Given the uncertainty in language (C++, Fortran), on-node programming model (OpenMP, OpenACC, Kokkos, etc.), and distributed-memory programming model (MPI, PGAS, task-based), there is a strong desire in the modeling community to evaluate them all to determine which will be appropriate. Unfortunately, targeting full applications or even components is an extremely time-consuming activity. In some cases, vendors have simply refused to port, analyze, optimize, and simulate such large pieces of code. They, like those developing new programming models, prefer smaller, more tractable, yet representative proxies that can be used to drive the development of architecture, tools, compilers, and programming models. The development of climate and biological proxy apps will facilitate collaboration between architects and researchers from vendors, ASCR, and BER. Such proxy apps will set realistic performance expectations and help determine which architectural or programming model constructs are viable before substantial effort is expended porting large applications.

3.2.7.3 **Software Engineering for Portability**

3.2.7.3.1 *Technical Challenges and Opportunities*

The system architectures for the DOE exascale systems are still to be determined. Because of uncertainty in the architectural directions of the future, it will be necessary for the code to be portable across two or more different exascale systems.

DOE is currently investigating two different pre-exascale architectures: (1) a modest number of compute nodes with multiple multicore CPUs and multiple accelerators (e.g., general purpose graphical processing units [GPGPUs]) per node; and (2) a large number of compute nodes, each with many-core CPUs. These are represented by the Summit system to be located at OLCF and the Aurora system to be located at ALCF, respectively, with the Cori system at NERSC also representing the many-core system architecture.

These two architecture “swim lanes” do share many characteristics — characteristics that are highly likely to be shared by future exascale architectures. One example is a deep memory hierarchy that will require explicit attention in algorithms and application implementations in order to preserve even current performance levels, much less to take advantage of the 100 times increase over current performance capability that will characterize the exascale systems. Present architectures also do not have as great an increase in either network or I/O performance (latency and bandwidth) as compared to the increase in theoretical compute performance.

However, these two swim lanes do differ significantly, and each will have different requirements for achieving good computer performance. These, and the ultimate exascale architectures, also differ from the systems upon which much application development is and will be taking place.

Here, portability has many aspects. First is the ability even to run an application and/or science case across the different systems. Achieving portability between compilers is often a difficult task even today, and with the differences in the underlying target hardware, the back ends of the compilers could be very different. The same required elements (libraries, programming models, etc.) also need to be available on all systems.

Second is performance portability — it should be feasible to (re)optimize the application codes for each platform while still having a testable, maintainable, and development-friendly code base that is “essentially” single source. In a worst-case scenario, multiple exascale architectures will require different code bases and possibly divergence between programming models, algorithms, or even model formulations.

A third, somewhat unique, aspect of portability for next-generation and future exascale architectures is that portability is not guaranteed on a single system for a single application if the application allows different science cases and different levels of parallel granularity. For example, in Earth system models, there are typically a number of fixed grid resolutions, number of tracers, etc., each to address a given class of science questions. With the expected sensitivity to the deep memory hierarchy, the change in memory requirements and memory access patterns could require significant re-optimization for different configurations of the application. This sensitivity also has implications on testing, requiring much more work to provide significant testing coverage.

A fourth aspect of portability is scientific portability. Climate, for example, is a chaotic system, and differences in CPU architecture can lead to bitwise differences in model output at a given timestep. The simulated Earth system, however, as calculated from averages of the model's direct output, should be independent of hardware platform even in the exascale era as it has been so far.

Finally, the need for portability may affect the entire tool chain including analysis and pre and post-processing programs because the LCFs may not have a separate, more standard analysis system co-located with the exascale system.

3.2.7.3.2 Priority Research Directions

The priority research directions are also cross-cutting; see Section 3.2.7.3.3.

3.2.7.3.3 Cross-Cutting Research Directions

All of the HPC applications in BER will be concerned with portability across exascale systems. Methods for achieving portability may be of research interest to ASCR. An immediate concern is the portability of DOE application codes between the development systems — Summit, Aurora, and Cori — and the portability between development systems and the ultimate exascale architectures.

One way to enhance portability is to employ a single programming model or programming language that can address both systems; however, even the best individual programming models that will be appropriate for the exascale architecture(s) have yet to be completely determined, and currently appear to differ somewhat between the Summit and Aurora/Cori systems, at least as first delivered.

Identifying representative mini-apps to pass on to LCFs and vendors will be important to help in understanding the differences between future machine-specific programming models, and these differences could be minimized if the mini-apps are incorporated early in the design process. At the minimum, it will allow the Earth system modeling community to select from among the available programming models that modelers can focus on in future development and identify ways to support these programming models within a given application.

The current state-of-the-art is to continue with our procedural language legacy (Fortran/C) codes while adding OpenMP and OpenACC directives. This approach is not yet affecting readability but does affect modifiability if other developers are trying to change code targeted by the directives. Machine-specific versions of some routines are starting to appear. While OpenACC and OpenMP do not yet have identical capabilities, there is movement toward a unification. One capability likely to be critical is an effective, hierarchical tiling clause that will permit effective use of the memory hierarchy.

Even the adoption of a single programming model that captures the capabilities of both OpenMP and OpenACC will likely require customizations in the way that it is used when targeting the different architectures, for example, to capture the different sizes and performance capabilities of each level of the memory hierarchy between the two systems and to expose vectorization opportunities on one architecture as compared to GPU-like SIMD parallelism on another.

Auto-tuning approaches have been useful on small kernels but have yet to show applicability to large applications. However, auto-tuning could enhance portability by mitigating the need to develop and maintain multiple, manually optimized codes. However, auto-tuning is most effective when optimizing over runtime options. This circumstance will require that some of the architecture-specific optimizations — for example, the details of the memory hierarchy and the associated placement of application data — be runtime controllable.

Where multiple languages or programming models are necessary, the software development process to maintain them needs improvement. HPC-integrated development environments could help with maintenance, especially if multiple algorithms and solution methods are required. Domain-specific languages also show promise but, like new programming languages, require community buy-in which, in turn, requires some guarantees of future support and a way to add them incrementally.

Assuming a single new language or programming model (PGAS, C++ meta programming, etc.) provides the best path to portability, we will need methods to confidently rewrite our 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. The portability verification system may need to account for exascale systems that do not provide reproducibility or are prone to soft errors.

3.2.7.3.4 Computing Needs and Requirements

Portability does not by itself create additional demands on the size or scale of the future hardware in the global computing ecosystem. The needs and requirements are primarily on the system software and programming environment sides. However, access to small and medium-size systems of the same architectural types for development purposes will greatly aid in evaluating portability and preparing applications for transition to the large systems.

Today, portability is greatly aided by the existence of common languages, libraries, and development environments found across the computing landscape (e.g., Fortran/C, MPI, vi/emacs). As indicated above, there is still much uncertainty about what the programming language for exascale will be; however, portability will likely not be possible if the same programming model is not available on all platforms of the exascale computing ecosystem. For performance portability, we will need sufficient documentation for all performance-related options in the software stack. For scientific portability, highly volatile systems with potential for many soft errors will need to offer some control, perhaps at increased cost, to help verify that a system will produce a simulation that is comparable to known good solutions.

This page is intentionally left blank.

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.

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.

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.

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.

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.

4.5 Conclusion

The requirements have been identified that are crucial to the development of a computing ecosystem and that fall within the broad categories of method development, computational environment, data, and communication and community involvement. These areas provide many opportunities for structured, collaborative engagement so that BER and ASCR can address the increasing complexity of high-performance computing as a fundamental component of the scientific enterprise.

This page is intentionally left blank.

5 REFERENCES

- Allan, R., P. Brohan, G.P. Compo, R. Stone, J. Luterbacher, and S. Brönnimann, 2011, “The International Atmospheric Circulation Reconstructions over the Earth (ACRE) Initiative,” *Bull. Amer. Met. Soc.*, Vol. 92, pp. 1421–1425. DOI: 10.1175/2011BAMS3218.1.
- Baohui, Li, L. Hai, L. Yu, W. Anliang, and J. Shunying, 2014, “A Modified Discrete Element Model for Sea Ice Dynamics,” *Acta Ocean. Sinica*, Vol. 33, No. 1, pp. 56–63. DOI: 10.1007/s13131-014-0428-3.
- BERAC (Biological and Environmental Research Advisory Committee), 2013, *BER Virtual Laboratory: Innovative Framework for Biological and Environmental Grand Challenges*, U.S. Department of Energy, February, http://science.energy.gov/~media/ber/berac/pdf/20130221/BERACVirtualLaboratory_Feb-18-2013.pdf, accessed October 23, 2016.
- Boyle, J.S., S.A. Klein, D.D. Lucas, H.-Y. Ma, J. Tannahill, and S. Xie, 2015, “The Parametric Sensitivity of CAM5’s MJO,” *J. Geophys. Res. Atmos.*, Vol. 120, pp. 1424–1444. DOI: 10.1002/2014JD022507.
- Clarke, J.F., and J.A. Edmonds, 1993, “Modelling Energy Technologies in a Competitive Market,” *Energy Economics*, Vol. 15, No. 2, pp. 123–129.
- Clarke, L., K. Jiang, K. Akimoto, M. Babiker, G. Blanford, K. Fisher-Vanden, J.-C. Hourcade, V. Krey, E. Kriegler, A. Loschel, D. McCollum, S. Paltsev, S. Rose, P.R. Shukla, M. Tavoni, B. van der Zwaan, and D.P. van Vuuren, 2014, “Assessing Transformation Pathways,” *Climate Change 2014, Mitigation of Climate Change*, O. Edenhofer, R. Pichs-Madruga, Y. Sokona, et al., eds., Cambridge, United Kingdom, and New York, New York, United States: Cambridge University Press.
- Collins, W.D., A.P. Craig, J.E. Truesdale, A.V. Di Vittorio, A.D. Jones, B. Bond-Lamberty, K.V. Calvin, J.A. Edmonds, S.H. Kim, A.M. Thomson, P. Patel, Y. Zhou, J. Mao, X. Shi, P.E. Thornton, L.P. Chini, and G.C. Hurtt, 2015, “The Integrated Earth System Model Version 1: Formulation and Functionality,” *Geosci. Model Dev.*, Vol. 8, No. 7, pp. 2203–2219.
- Compo, G., J. Carton, X. Dong, A. Kumar, S. Saha, J.S. Woollen, L. Yu, and H.M. Archambault, 2016, *Report from the NOAA Climate Reanalysis Task Force Technical Workshop*, NOAA Technical Report OAR CPO-4, Silver Spring, Maryland. DOI: 10.7289/V53J39ZZ.
- Compo, G.P., J.S. Whitaker, et al., 2011, “The Twentieth Century Reanalysis Project,” *Q.J.R. Meteorol. Soc.*, Vol. 137, pp. 1–28. DOI: 10.1002/qj.776.
- Dart, E., L. Rotman, B. Tierney, M. Hester, and J. Zurawski, 2013, “The Science DMZ: A Network Design Pattern for Data Intensive Science,” in *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, SC ’13. New York, New York, United States: ACM, 2013, pp. 85:1–85:10, accessed February 24, 2017.
- Davison, B.H., J. Parks, M.F. Davis, and B.S. Donohoe, 2013, “Plant Cell Walls: Basics of Structure, Chemistry, Accessibility, and the Influence on Conversion,” *Aqueous Pretreatment of Plant Biomass for Biological and Chemical Conversion to Fuels and Chemicals*, First Edition, Charles E. Wyman (ed.), JohnWiley & Sons, Ltd., http://bioenergycenter.org/besc/publications/AqueousPretreatment_3.pdf, accessed February 13, 2017.

de Boyer Montegut, C., G. Madec, A.S. Fischer, A. Lazar, and D. Iudicone, 2004, “Mixed Layer Depth over the Global Ocean: An Examination of Profile Data and a Profile-based Climatology,” *J. Geophys. Res.*, Vol. 109, C12003, DOI:10.1029/2004JC002378.

DOE (U.S. Department of Energy), 2014, *The Water-Energy Nexus: Challenges and Opportunities*, June, <http://energy.gov/downloads/water-energy-nexus-challenges-and-opportunities>, accessed July 14, 2016.

DOE, 2015a, *Building Virtual Ecosystems: Computational Challenges for Mechanistic Modeling of Terrestrial Environments: Workshop Report*, DOE/SC-0171, Office of Science, Subsurface Biogeochemical Research Program, February, <http://doesbr.org/BuildingVirtualEcosystems/>, accessed December 21, 2016.

DOE, 2015b, *Building a Cyberinfrastructure for Environmental System Science: Modeling Frameworks, Data Management, and Scientific Workflows: Workshop Report*, DOE/SC-0178, Office of Science, Subsurface Biogeochemical Research Program, November, <http://doesbr.org/ESS-WorkingGroups/>, accessed December 21, 2016.

DOE, 2017, *IDEAS Project – Interoperable Design of Extreme-Scale Application Software (IDEAS)*, <https://ideas-productivity.org/>, accessed December 21, 2016.

DOE-ARM (DOE-Atmospheric Radiation Measurement), 2016, *Application Programming Interface*.

DOE-BER (DOE-Office of Biological and Environmental Research), 2015, *Biological Systems Science Division Strategic Plan*, Office of Science, October, <http://genomicscience.energy.gov/pubs/BSSDStrategicPlanOct2015.pdf>, accessed December 20, 2016.

DOE-CESD (DOE-Climate and Environmental Sciences Division), 2014, *Atmospheric Testbed Workshop*, 42 pp., DOE/SC-0163, U.S. Department of Energy, http://science.energy.gov/~media/ber/pdf/Brochures/CESD_Atmospheric_Testbed_Workshop_V9_LOW.pdf?id=97, accessed December 30, 2014.

Fan, J.W., Y.C. Liu, K.M. Xu, K. North, S. Collis, X.Q. Dong, G.J. Zhang, Q. Chen, P. Kollias, and S.J. Ghan, 2015, “Improving Representation of Convective Transport for Scale-Aware Parameterization: 1. Convection and Cloud Properties Simulated with Spectral Bin and Bulk Microphysics,” *J. Geophys. Res.*, Vol. 120, pp. 3485–3509. DOI:10.1002/2014jd022142.

Fast, J.D., W.I. Gustafson, Jr., E.G. Chapman, R.C. Easter, Jr., J.P. Rishel, R.A. Zaveri, G. Grell, and M. Barth, 2011, “The Aerosol Modeling Testbed: A Community Tool to Objectively Evaluate Aerosol Process Modules,” *Bulletin of the American Meteorological Society*, Vol. 92, No. 3, pp. 343–360. DOI:10.1175/2010BAMS2868.1.

Friedlingstein, P., M. Meinshausen, V.K. Arora, C.D. Jones, A. Anav, S.K. Liddicoat, and R. Knutti, 2014, “Uncertainties in CMIP5 Climate Projections Due to Carbon Cycle Feedbacks,” *J. of Climate*, Vol. 27, pp. 511–526. DOI: 10.1175/JCLI-D-12-00579.1.

Georganas, E., 2016, “Scalable Parallel Algorithms for Genome Analysis,” Ph.D. thesis, EECS Department, University of California, Berkeley, August.

Globus, 2017, “The Globus Project,” <https://www.globus.org/>, accessed March 6, 2017.

Guba, O., M.A. Taylor, P.A. Ullrich, J.R. Overfelt, and M.N. Levy, 2014, “The Spectral Element Method (SEM) on Variable-Resolution Grids: Evaluating Grid Sensitivity and Resolution-Aware Numerical Viscosity,” *Geosci. Model Dev.*, Vol. 7, pp. 2803–2816. DOI:10.5194/gmd-7-2803-2014.

Gustafson, W.I., and A.M. Vogelmann, 2015, *LES ARM Symbiotic Simulation and Observation (LASSO) Implementation Strategy*, 31 pp., DOE/SC-ARM-15-039, U.S. Department of Energy Atmospheric Radiation Measurement Climate Research Facility, <http://www.arm.gov/publications/programdocs/doe-sc-arm-15-039.pdf>, accessed October 3, 2015.

Hanemann, A., J.W. Boote, E.L. Boyd, J. Durand, L. Kudarimoti, R. Lapacz, D.M. Swamy, S. Trocha, and J. Zurawski, 2005, “Perfsonar: A Service Oriented Architecture for Multi-Domain Network Monitoring,” in *Proceedings of the Third International Conference on Service-Oriented Computing*, ICSOC’05, Berlin, Heidelberg: SpringerVerlag, pp. 241–254, http://dx.doi.org/10.1007/11596141_19.

Helland, B., 2016, *BER Requirements Review*, presentation at the Biological and Environmental Research Exascale Requirements Review Meeting, U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research, Rockville, Maryland, January 27.

Hogg, A. McC., M.P. Meredith, D.P. Chambers, E.P. Abrahamson, C.W. Hughes, and A.K. Morrison, 2015, “Recent Trends in the Southern Ocean Eddy Field,” *J. Geophys. Res. Oceans*, Vol. 120, pp. 257–267. DOI:10.1002/2014JC010470.

Hopkins, M.A., 2004, “A Discrete Element Lagrangian Sea Ice Model,” *Engineering Computations*, Vol. 21, Iss 2/3/4, pp. 409–421, <http://dx.doi.org/10.1108/02644400410519857>.

Hunke, E.C., W.H. Lipscomb, and A.K. Turner, 2010, “Sea-ice Models for Climate Study: Retrospective and New Directions,” *J. Glaciol.*, Vol. 56, No. 200, pp. 1162–1172(11).

Im, W., J. Liang, A. Olson, H.X. Zhou, S. Vajda, and I.A. Vakser, 2016, “Challenges in Structural Approaches to Cell Modeling,” *Journal of Molecular Biology*, Vol. 428, No. 15, pp. 2943–2964. DOI: 10.1016/j.jmb.2016.05.024.

Jeffries, M.O., J. Richter-Menge, and J.E. Overland (eds.), 2015, *Arctic Report Card: Tracking Recent Environmental Changes Relative to Historical Records*, Arctic Program, National Oceanic and Atmospheric Administration, <http://www.arctic.noaa.gov/reportcard>, accessed February 14, 2017.

Jung, J.-H., and A. Arakawa, 2014, “Modeling the Moist-Convective Atmosphere with a Quasi-3-D Multiscale Modeling Framework (Q3D MMF),” *J. Adv. Model. Earth Syst.*, Vol. 6, pp. 185–205. DOI:10.1002/2013ms000295.

Khairoutdinov, M.F., S.K. Krueger, C.H. Moeng, P.A. Bogenschutz, and D.A. Randall, 2009, “Large-Eddy Simulation of Maritime Deep Tropical Convection,” *J. Adv. Model. Earth Syst.*, Vol. 1. DOI:10.3894/james.2009.1.15.

Kinnard, C., C.M. Zdanowicz, D.A. Fisher, E. Isaksson, A. de Vernal, and L.G. Thompson, 2011, “Reconstructed Changes in Arctic Sea Ice over the Past 1,450 Years,” *Nature*, Vol. 479, pp. 509–512.

Klinker, E., and P.D. Sardeshmukh, 1992, “The Diagnosis of Mechanical Dissipation in the Atmosphere from Large-Scale Balance Requirement,” *J. Atmos. Sci.*, Vol. 49, pp. 608–627. DOI: 10.1175/1520-0469(1992)049<0608:TDOMDI>2.0.CO;2.

Kraucunas, I., L. Clarke, J. Dirks, J. Hathaway, M. Hejazi, K. Hibbard, M. Huang, C. Jin, M. Kintner-Meyer, K. van Dam, R. Leung, H.-Y. Li, R. Moss, M. Peterson, J. Rice, M. Scott, A. Thomson, N. Voisin, and T. West, 2015, “Investigating the Nexus of Climate, Energy, Water, and Land at Decision-Relevant Scales: the Platform for Regional Integrated Modeling and Analysis (PRIMA),” *Climatic Change*, Vol. 129, No. 3, pp. 573–588.

Ma, H.-Y., C.C. Chuang, S.A. Klein, M.-H. Lo, Y. Zhang, S. Xie, X. Zheng, P.-L. Ma, Y. Zhang, and T.J. Phillips, 2015, “An Improved Hindcast Approach for Evaluation and Diagnosis of Physical Processes in Global Climate Models,” *J. Adv. Model. Earth Syst.*, Vol. 7, pp. 1810–1827. DOI:10.1002/2015MS000490.

McCammon, J.A., and M. Karplus, 1980, “Simulation of Protein Dynamics,” *Annual Review of Physical Chemistry*, Vol. 31, No. 1, pp. 29–45.

Miyamoto, Y., Y. Kajikawa, R. Yoshida, T. Yamaura, H. Yashiro, and H. Tomita, 2013, “Deep Moist Atmospheric Convection in a Subkilometer Global Simulation,” *Geophys. Res. Lett.*, Vol. 40, pp. 49224926. DOI:10.1002/grl.50944.

Molins, S., D. Trebotich, L. Yang, J.B. Ajo-Franklin, T.J. Ligocki, C. Shen, and C.I. Steefel, 2014, “Pore-Scale Controls on Calcite Dissolution Rates from Flow-Through Laboratory and Numerical Experiments,” *Environmental Science and Technology*. DOI: 10.1021/es5013438.

NRC (National Research Council of the National Academies), 2009, *A New Biology for the 21st Century*, Committee on a New Biology for the 21st Century: Ensuring the United States Leads the Coming Biology Revolution, Board on Life Sciences, Division on Earth and Life Studies, Washington, D.C.: The National Academies Press, <https://www.nap.edu/read/12764>, accessed December 20, 2016. ISBN: 978-0-309-14488-9, DOI: 10.17226/12764.

Phillips, T.J., G.L. Potter, D.L. Williamson, R.T. Cederwall, J.S. Boyle, M. Fiorino, J.J. Hnilo, J.G. Olson, S. Xie, and J.J. Yio, 2004, “Evaluating Parameterizations in General Circulation Models Climate Simulation Meets Weather Prediction,” *Bull. Amer. Meteor. Soc.*, Vol. 85, pp. 1903–1915. DOI:10.1175/bams-85-12-1903.

Poli, P., et al., 2016, “ERA-20C: An Atmospheric Reanalysis of the 20th Century,” *J. Climate*. DOI: 10.1175/JCLI-D-15-0556.1, in press.

Putnam, N.H., B.L. O’Connell, J.C. Stites, B.J. Rice, M. Blanchette, R. Calef, C.J. Troll, A. Fields, P.D. Hartley, C.W. Sugnet, D. Haussler, D.S. Rokhsar, and R.E. Green, 2016, “Chromosome-Scale Shotgun Assembly Using an In Vitro Method for Long-Range Linkage,” *Genome Assembly*, Vol. 26, pp. 342–350. doi:10.1101/gr.193474.115.

Qian, Y., et al., 2015, “Parametric Sensitivity Analysis of Precipitation at Global and Local Scales in the Community Atmosphere Model CAM5,” *J. Adv. Model. Earth Syst.*, Vol. 7, pp. 382–411. DOI:10.1002/2014MS000354.

Rahmsdorf, S., 2006, “Thermohaline Ocean Circulation,” in *Encyclopedia of Quaternary Sciences*, S.A. Elias (ed.), Amsterdam: Elsevier.

Reed, K.A., C. Jablonowski, and M.A. Taylor, 2012, “Tropical Cyclones in the Spectral Element Configuration of the Community Atmosphere Model,” *Atmos. Sci. Lett.*, Vol. 13, pp. 303–310, DOI: 10.1002/asl.399.

Rodwell, M.J., and T.N. Palmer, 2007, “Using Numerical Weather Prediction to Assess Climate Models,” *Q.J.R. Meteorol. Soc.*, Vol. 133, pp. 129–146. doi:10.1002/qj.23.

Sardeshmukh, P.D., G.P. Compo, and C. Penland, 2015, “Need for Caution in Interpreting Extreme Weather Statistics,” *J. Climate*, Vol. 28, pp. 9166–9187. DOI: 10.1175/JCLI-D-15-0020.1.

Schalkwijk, J., H.J.J. Jonker, A.P. Siebesma, and E. Van Meijgaard, 2015, “Weather Forecasting Using GPU-Based Large-Eddy Simulations,” *Bull. Amer. Meteor. Soc.*, Vol. 96, pp. 715–724. DOI: 10.1175/bams-d-14-00114.1.

Sherwood, C.S., S. Bony, and J.-L. Dufresne, 2014, “Spread in Model Climate Sensitivity Traced to Atmospheric Convective Modeling,” *Nature*, Vol. 505, pp. 37–42. DOI: 10.1038/nature12829.

Shindell, D.T., J.-F. Lamarque, M. Schulz, M. Flanner, C. Jiao, M. Chin, P.J. Young, Y.H. Lee, L. Rotstayn, N. Mahowald, G. Milly, G. Faluvegi, Y. Balkanski, W.J. Collins, A.J. Conley, S. Dalsoren, R. Easter, S. Ghan, L. Horowitz, X. Liu, G. Myhre, T. Nagashima, V. Naik, S.T. Rumbold, R. Skeie, K. Sudo, S. Szopa, T. Takemura, A. Voulgarakis, J.H. Yoon, and F. Lo, 2013, “Radiative Forcing in the ACCMIP Historical and Future Climate Simulations,” *Atmos. Chem. Phys.*, Vol. 13, pp. 2939–2974. DOI: 10.5194/acp-13-2939-2013.

Smith, J.C., and B. Roux, 2013, “Eppur si muove: The 2013 Nobel Prize in Chemistry,” *Structure*, Vol. 21, No. 12, pp. 2102–2105.

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

Tolman, H.L., 2002, User Manual and System Documentation of WAVEWATCH-III version 2.22, National Oceanic and Atmospheric Administration, National Weather Service, National Centers for Environmental Prediction, OMB technical note 222, 133 pp.

Xie, S., R.T. Cederwall, and M. Zhang, 2004, “Developing Long-Term Single-Column Model Cloud System-Resolving Model Forcing Data Using Numerical Weather Prediction Products Constrained by Surface and Top of the Atmosphere Observations,” *J. Geophys. Res.*, Vol. 109, D01104. DOI: 10.1029/2003JD004045.

Zarzycki, C.M., C. Jablonowski, D.R. Thatcher, and M.A. Taylor, 2015, “Effects of Localized Grid Refinement on the General Circulation and Climatology in the Community Atmosphere Model,” *J. Climate*, Vol. 28, pp. 2777–2803. DOI: 10.1175/JCLI-D-14-00599.1, 2015.

Zhang, M.H., J.L. Lin, R.T. Cederwall, J.J. Yio, and S.C. Xie, 2001, “Objective Analysis of ARM IOP Data: Method and Sensitivity,” *Mon. Weather Rev.*, Vol. 129, pp. 295–311.

This page is intentionally left blank.

6 ACRONYMS AND ABBREVIATIONS

| | |
|------|--|
| ACME | Accelerated Climate Modeling for Energy (DOE) |
| AE | algebraic equation |
| ALCF | Argonne Leadership Computing Facility |
| AMT | Aerosol Modeling Testbed |
| API | application programming interface |
| ARM | Atmospheric Radiation Measurement (DOE) |
| ASCR | Advanced Scientific Computing Research (DOE) |
| ASIC | application-specific integrated circuit |
| | |
| BD | Brownian Dynamics |
| BER | Biological and Environmental Research |
| BGI | Beijing Genome Center |
| BSSD | Biological Systems Science Division |
| | |
| CAF | Coarray Fortran |
| CAPT | Cloud-Associated Parameterizations Testbed (DOE) |
| CESD | Climate and Environmental Sciences Division |
| CF | climate forecast |
| CFD | computational fluid dynamics |
| CG | coarse graining |
| CNN | convolutional neural network |
| CPU | central processing unit |
| CUDA | compute unified device architecture |
| | |
| DAE | differential algebraic equation |
| DFT | density functional theory |
| DNA | deoxyribonucleic acid |
| DNN | deep neural network |
| DNS | Direct Numerical Simulation |
| DoE | Design of Experiment |
| DOE | U.S. Department of Energy |
| DR | dimensionality reduction |
| DSeL | domain-specific embedded language |
| DSL | domain-specific language |

| | |
|---------|--|
| EA | Exploratory Analysis |
| EnKF | Ensemble Kalman Filter |
| eQTL | expression quantitative trait loci |
| ESGF | Earth System Grid Federation |
| ESM | Earth System Model |
| | |
| FFT | Fast Fourier Transform |
| | |
| GC | gas chromatography |
| GCRM | global cloud-resolving models |
| GPGPU | general purpose graphical processing unit |
| GPS | global positioning system |
| | |
| HPC | high-performance computing |
| | |
| ICA | Independent Components Analysis |
| I/O | input/output |
| IAM | integrated assessment modeling |
| iRF | iterative Random Forests |
| | |
| KNH | Knights Hill (Intel) |
| KNL | Knights Landing (Intel) |
| | |
| LASSO | LES ARM Symbiotic Simulation and Observation |
| LBNL | Lawrence Berkeley National Laboratory |
| LC | liquid chromatography |
| LCF | leadership computing facility |
| LES | large-eddy simulation |
| | |
| MCMC | Markov Chain Monte Carlo |
| MD | molecular dynamics |
| MIC | many integrated core |
| MILP | mixed-integer linear programming |
| MMF | multiscale modeling framework |
| MPI + x | Message Passing Interface Extension |
| MS | mass spectrometry |
| MSI | mass spectrometry imaging |

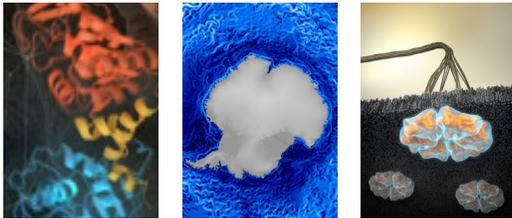
| | |
|-----------|--|
| NERSC | National Energy Research Scientific Computing Center |
| NIST | National Institute of Standards and |
| NLP | nonlinear programming |
| NOM | natural organic matter |
| NWP | numerical weather prediction |
| ODE | ordinary differential equation |
| OGC | Open Geospatial Consortium |
| OLCF | Oak Ridge Leadership Computing Facility |
| PCR | polymerase chain reaction |
| PDE | partial differential equation |
| PGAS | partitioned global address space |
| PME | Particle-Mesh-Ewald |
| PRD | priority research direction |
| Q3D-MMF | quasi-3D multiscale modeling framework |
| QM/MM | quantum mechanical/molecular mechanical |
| RNA | ribonucleic acid |
| RRM | Regionally Refined Model |
| SIMD | single instruction multiple data |
| SLR | sea level rise |
| SPDE | stochastic partial differential equation |
| SPMD | single-program multiple-data |
| UoI | Union of Intersections |
| UPC/UPC++ | unified parallel C |
| UQ | uncertainty quantification |
| UV-CDAT | Ultrascale Visualization Climate Data Analysis Toolkit |
| WPF | web processing services |
| WRF | weather research and forecasting |

Units of Measure

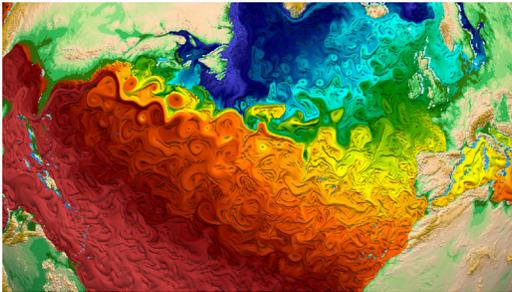
| | | | |
|----|---------------|-----|-------------|
| bp | base pair(s) | m | meter(s) |
| | | Mb | megabyte(s) |
| cm | centimeter(s) | | |
| | | PB | petabyte(s) |
| g | gram(s) | | |
| GB | gigabyte(s) | sec | second(s) |
| K | thousand(s) | TB | terabyte(s) |
| kb | kilobase(s) | | |
| km | kilometer(s) | μg | micron |
| | | yr | year |

BER

BIOLOGICAL AND ENVIRONMENTAL RESEARCH



APPENDICES: MEETING MATERIALS



An Office of Science review sponsored jointly by
Advanced Scientific Computing Research and
Biological and Environmental Research



This page is intentionally left blank.

APPENDIX A: BIOLOGICAL AND ENVIRONMENTAL RESEARCH ORGANIZING COMMITTEE AND REVIEW PARTICIPANTS

A.1 BER ORGANIZING COMMITTEE

Adam Arkin (aparkin@lbl.gov)
David C. Bader (bader2@llnl.gov)
Tjerk P. Straatsma (str@ornl.gov)
Jack C. Wells (wellsjc@ornl.gov)
James J. Hack (jhack@ornl.gov)
Richard Coffey (richardc@alcf.anl.gov)
Eli Dart (Dart@es.net)

A.2 BER SPONSORS AND REPRESENTATIVES

Dorothy Koch (Dorothy.Koch@science.doe.gov)
Ramana Madupu (Ramana.Madupu@science.doe.gov)
Carolyn Lauzon (carolyn.lauzon@science.doe.gov)

A.3 BER REVIEW PARTICIPANTS: BIOSYSTEMS SCIENCE DIVISION

| | |
|------------------------|--|
| Srinivas Aluru | Georgia Tech |
| Adam Arkin | Lawrence Berkeley National Laboratory |
| Nitin Baliga | Institute for Systems Biology |
| Rich Bonneau | New York University |
| Eoin Brodie | Lawrence Berkeley National Laboratory |
| James (Ben) Brown | Lawrence Berkeley National Laboratory |
| Aydin Buluc | Lawrence Berkeley National Laboratory |
| Bill Cannon | Pacific Northwest National Laboratory |
| Xiaolin Cheng | Oak Ridge National Laboratory |
| Margaret Cheung | University of Houston |
| Markus Covert | Stanford University |
| Mike Crowley | National Renewable Energy Laboratory |
| Ron Dror | Stanford University |
| Kjiersten Fagnan | National Energy Research Scientific Computing Center |
| Iddo Friedberg | Iowa State University |
| Gnana Gnanakaran | Los Alamos National Laboratory |
| Vassily Hatzimanikatis | École Polytechnique Fédérale de Lausanne |
| Dan Jacobson | Oak Ridge National Laboratory |

| | |
|-------------------------|---|
| Matt Jacobson | University of California-San Francisco |
| Ananth Kalyanaraman | Washington State University |
| Kerstin Kleese Van Dam | Brookhaven National Laboratory |
| Leslie Loew | University of Connecticut |
| Radhakrishnan Mahadevan | University of Toronto |
| Costas Maranas | Pennsylvania State University |
| Reed Maxwell | Colorado School of Mines |
| Lee Ann McCue | Pacific Northwest National Laboratory |
| David Moulton | Los Alamos National Laboratory |
| Chongle Pan | Oak Ridge National Laboratory |
| Loukas Petridis | Oak Ridge National Laboratory |
| Dan Rokhsar | Lawrence Berkeley National Laboratory/Joint Genome Institute |
| Ingo Ruczinski | Johns Hopkins University |
| Tim Scheibe | Pacific Northwest National Laboratory |
| Steven Singer | Lawrence Berkeley National Laboratory/Joint BioEnergy Institute |
| Jeremy Smith | Oak Ridge National Laboratory |
| Ilya Tagkopoulos | U.C. Davis |
| Steven Wiley | Pacific Northwest National Laboratory |
| Sergio Wong | Lawrence Livermore National Laboratory |
| Kathy Yelick | Lawrence Berkeley National Laboratory |
| Nir Yosef | University of California-Berkeley |
| Yang Zhang | University of Michigan |

A.4 BER REVIEW PARTICIPANTS: CLIMATE AND ENVIRONMENTAL SCIENCES DIVISION

| | |
|------------------|--|
| Val Anantharaj | Oak Ridge National Laboratory |
| Dave Bader | Lawrence Livermore National Laboratory |
| Susan Bates | National Center for Atmospheric Research |
| Susannah Burrows | Pacific Northwest National Laboratory |
| Kate Calvin | Pacific Northwest National Laboratory |
| Xinguan Chen | Pacific Northwest National Laboratory |
| Bill Collins | Lawrence Berkeley National Laboratory |
| Gil Compo | University of Colorado |
| Bert Debusschere | Sandia National Laboratories |
| Katherine Evans | Oak Ridge National Laboratory |
| Ian Foster | Argonne National Laboratory |

| | |
|------------------------|--|
| Jeremy Fyke | Los Alamos National Laboratory |
| Chris Golaz | Lawrence Livermore National Laboratory |
| Bill Gustafson | Pacific Northwest National Laboratory |
| Glenn Hammond | Sandia National Laboratories |
| Forrest Hoffman | Oak Ridge National Laboratory |
| Charles Jackson | University of Texas-Austin |
| Rob Jacob | Argonne National Laboratory |
| Jeff Johnson | Lawrence Berkeley National Laboratory |
| Andy Jones | Lawrence Berkeley National Laboratory |
| Phil Jones | Los Alamos National Laboratory |
| Kirsten Kleese van Dam | Brookhaven National Laboratory |
| Pavlos Kollias | Stony Brook University |
| Ruby Leung | Pacific Northwest National Laboratory |
| Robert Link | Pacific Northwest National Laboratory |
| Hsi-Yen Ma | Lawrence Livermore National Laboratory |
| Dan Martin | Lawrence Berkeley National Laboratory |
| Wieslaw Maslowski | Naval Postgraduate School |
| Reed Maxwell | Colorado School of Mines |
| Lois Curfman McInnes | Argonne National Laboratory |
| Dave Moulton | Los Alamos National Laboratory |
| Esmond Ng | Lawrence Berkeley National Laboratory |
| Matt Norman | Oak Ridge National Laboratory |
| Travis O'Brien | Lawrence Berkeley National Laboratory |
| Giri Palanisamy | Oak Ridge National Laboratory |
| George Pau | Lawrence Berkeley National Laboratory |
| Dave Randall | Colorado State University |
| Phil Rasch | Pacific Northwest National Laboratory |
| Todd Ringler | Los Alamos National Laboratory |
| Tim Scheibe | Pacific Northwest National Laboratory |
| Chitra Sivaraman | Pacific Northwest National Laboratory |
| Mark Taylor | Sandia National Laboratories |
| Peter Thornton | Oak Ridge National Laboratory |
| Adrian Turner | Los Alamos National Laboratory |
| Nathan Urban | Los Alamos National Laboratory |
| Dean Williams | Lawrence Livermore National Laboratory |
| Sam Williams | Lawrence Berkeley National Laboratory |

Pat Worley

Shaocheng Xie

Minghua Zhang

Oak Ridge National Laboratory

Lawrence Livermore National Laboratory

Stony Brook University

APPENDIX B: BIOLOGICAL AND ENVIRONMENTAL RESEARCH MEETING AGENDA

MONDAY, MARCH 28

19:00 Chairs and Leads Pre-Meeting

TUESDAY, MARCH 29

7:30 Registration

8:30 Welcome & Introduction
Steve Binkley, Advanced Scientific Computing Research (ASCR), and
Sharlene Weatherwax, Biological and Environmental Research (BER)

8:40 Genesis of this Meeting
Barb Helland, ASCR

9:15 Todd Anderson and Gary Geernaert, BER
View from Biological and Environmental Research
Dorothy Koch, Ramana Madupu, BER

9:50 View from ESnet – Eli Dart, ASCR

10:00 Exascale Review Introductions
Chairs and Session Leaders

10:30 Break

10:45 ASCR Computing Facilities Presentation
Katherine Riley, ALCF

12:00 Working Lunch, Chairs and Sessions Leads

- ASCR Computing Facilities Presentation
- Charge to Working Groups

13:00–16:00 Breakout Sessions:

Climate 13:00–14:00
2) Climate and Environmental Science in the Exascale era (Leung)
What are the grand challenges for climate and Environmental science in the 5-10 year timeframe? What are limiting factors to achieving results? How would a 100× improvement in computational capabilities enhance research? Plenary speakers: Dave Randall and Bill Collins

14:00–16:00
Component-specific science and challenges (Leung) (2 hours; 3 parallel breakouts).
For each component: What are the priorities and challenges in representing the system? What advances are needed or anticipated?
(3a) Atmospheric Research (Minghua Zhang, Bill Gustafson)
(3b) Terrestrial and Subsurface Research (Peter Thornton, Dave Moulton)
(3c) Oceans and Cryospheric Research (Todd Ringler, Wieslaw Maslowski)

| | |
|-------------|--|
| Biology | 13:00-16:00 Big Data for Knowledge Discovery in Biology (Kathy Yelick, Eoin Brodie) OR From Sequence to Models of Organisms and Communities. Leaders: Dan Rokhsar and Rich Bonneau |
| | 13:00-14:00 From Sequence to Molecular Structure and Function Dan Rokhsar — Genome sequencing and diversity Yang Zhang — Protein structure prediction and annotation Iddo Friedberg — Functional annotation Ingo Ruczinski — Frontiers in primary data bioinformatics |
| | 14:00-15:30 Integrative modeling of cells and networks Ilias Tagkopoulos — Multi-omics network models Marcus Covert — Whole cell models: what’s missing Vassily Hatzimanikatis — Metabolic modeling – a hybrid approach Rich Bonneau — Simultaneous estimation of networks and activities/function Kerstin Kleese Van Dam — Putting the genie back in the bottle |
| | 15:30 Discussions of sessions and white paper tuning |
| 16:00 | Break |
| 16:15 | Q&A Session with the BER and ASCR Associate Directors |
| 17:30–18:15 | Breakout leads present key questions and issues discussed in each session |
| 18:30 | Dinner on your own |

WEDNESDAY, MARCH 30

| | |
|------------|---|
| 8:00 | Check-in |
| 8:30 | Summary from the chairs, outline of report |
| 9:00–12:00 | Breakout Sessions |
| Climate | Coupled System Integration Challenges (Bill Collins) Working toward frameworks and hierarchies of modeling, e.g., ESM-IA-IAV parallel breakouts on: (4a) ESMs (Phil Jones, Nathan Urban) (4b) IA models (Kate Calvin, Andy Jones) |
| Biology | From Sequence to Models of Organisms and Communities (Dan Rokhsar, Rich Bonneau) OR Big Data for Knowledge Discovery in Biology (Kathy Yelick, Eoin Brodie) Big Data applications and challenges across scales 1. Assembly: Extreme scale metagenomic assembly (Buluc) 10 min 2. Multi-omics in complex environments: Metagenomics and proteomics of soil (Pan) – 10 min 3. Single cells- Big data: Understanding and utilizing cellular variability (Yosef) – 10 min 4. Discussion – 10 min |

Combinatorial Analysis

1. Integration of multi-omics and interaction data for predictive modeling and analytics (Tagkopolos) – 10 min
2. Combinatorial Applications for Systems Biology (Jacobson) – 10 min
3. BigData For Knowledge Discovery in Biology & Multiscale simulation of Biophysical Processes from Molecules to Biological Interfaces (Brown) – 10 min

| | |
|-------------|---|
| EMSL | 9:00–12:00 Joint Session (Lee Ann McCue) |
| 12:00 | Working Lunch |
| 13:00–16:00 | Breakout Sessions |
| Climate | Computational Issues for Climate and Environment (Mark Taylor) |

| Schedule | Breakout Room 1 | Breakout Room 2 | Breakout Room 3 |
|----------|--------------------|-------------------------------|-------------------|
| Hour 1 | Algorithms | Observational Data Processing | Data Management |
| Hour 2 | Programming Models | (no breakout) | Testbeds |
| Hour 3 | SE for portability | UQ | Data Assimilation |

5a) Climate and Environment with Exascale architectures

- 5a-1) Algorithms (Kate Evans)
- 5a-2) Programming Models (Esmond Ng)
- 5a-3) Software Engineering for Portability (Pat Worley, Rob Jacob)
- 5b) Unlocking Scientific Knowledge with Exabytes of Data
- 5b-1) Large-Scale Heterogeneous Data Management (Dean Williams)
- 5b-2) Observational Data Processing: Retrieval Algorithms, Instrument Network Simulation (Pavlos Kollias)
- 5b-3) Data Reduction and Analytics of Observations and Model Results: UQ, Statistics, Emulators, Analytics of Model-Observation (Bert Debusschere)
- 5c) Data-Model Fusion
- 5c-1) Model Development Test Beds (Hsi-Yen Ma)
- 5c-2) Data Assimilation, Model Initialization and Reanalysis (Gil Compo)

| | |
|-------------|--|
| Biology | Multiscale Simulation of Biological Processes from Molecules to Biological Interfaces (Jeremy Smith, Matt Jacobson) |
| 16:00 | Break |
| 16:20 | Reports on Wednesday Breakouts, Breakout Leads, 15 minutes each |
| 17:20–17:45 | Summary and Thanks from Chairs |

End for Most Participants

THURSDAY, MARCH 31

All Day: Co-chairs, Leads, Writers meet to continue working on report

This page is intentionally left blank.

APPENDIX C: BIOLOGICAL AND ENVIRONMENTAL RESEARCH (BER) WHITE PAPERS

The following white papers were submitted by the following authors in advance of the BER Exascale Requirements Review to guide both the agenda and review discussions.

C.1 White Papers Addressing Biological Research

Page No. *Multiscale Biophysical Simulation from Molecules to Cells (corresponds to Section 3.1.1)*

- C-5 Rich Bonneau (New York University) and Iddo Friedberg (Iowa State University)
- C-7 William Cannon and Senghwa Kang (Pacific Northwest National Laboratory)
- C-10 X. Cheng (Oak Ridge National Laboratory)
- C-13 M. Cheung (University of Houston)
- C-16 M. Crowley (National Renewable Energy Laboratory)
- C-19 Ron Dror (Stanford University)
- C-23 Gnana Gnanakaran (Los Alamos National Laboratory)
- C-26 Matt Jacobson (University of California-San Francisco)
- C-28 Leslie Loew (University of Connecticut School of Medicine) and Phillip Colella (Lawrence Berkeley National Laboratory)
- C-31 Loukas Petridis (Oak Ridge National Laboratory)
- C-34 Roland Schulz (Intel)
- C-35 Jeremy Smith (Oak Ridge National Laboratory)
- C-36 D.J. Tantillo (University of California-Davis)
- C-37 Marat Valiev, Niri Govind, Edoardo Aprá, and Karol Kowalski (all of EMSL-Pacific Northwest National Laboratory)

Page No. *Microbes to the Environment (corresponds to Section 3.1.3)*

- C-39 Amity Andersen, Ryan Renslow, Marat Valiev, Niri Govind, and Eric Bylaska (all of Pacific Northwest National Laboratory)
- C-42 Ananth Kalyanaraman (Washington State University)
- C-44 Timothy D. Scheibe (Pacific Northwest National Laboratory)
- C-48 Carl Steefel, David Trebotich, Eoin Brodie, Sergi Molins, and Eric King (Lawrence Berkeley National Laboratory)
- C-54 Will Wieder (NCAR/University of Colorado-Boulder), Mark A. Bradford (Yale University), Stuart Grandy (University of New Hampshire), and Jenifer Talbot (Boston University)

Page No. *Biological Big Data Challenges (corresponds to Section 3.1.4)*

- C-57 Ian Blaby, Crysten Blaby-Haas, and Shinjae Yoo (all of Brookhaven National Laboratory)
- C-59 Ben Bowen (Lawrence Berkeley National Laboratory)
- C-62 James (Ben) Brown (Lawrence Berkeley National Laboratory)
- C-65 Dan Jacobson (Oak Ridge National Laboratory)
- C-68 Sang-Yun Oh (University of California-Santa Barbara), Ariful Azad (Lawrence Berkeley National Laboratory), Aydin Buluc (Lawrence Berkeley National Laboratory), Penporn Koanantakool (University of California-Berkeley), Dmitriy Morozov (Lawrence Berkeley National Laboratory), Leonid Oliker (Lawrence Berkeley National Laboratory), and Katherine Yelick (Lawrence Berkeley National Laboratory)
- C-71 Allon Wagner and Nir Yosef (both of University of California-Berkeley)

C.2 White Papers Addressing Environmental Research

Page No. *Atmospheric Simulation and Data Assimilation within the Earth System (corresponds to Section 3.2.1)*

- C-75 William I. Gustafson, Jr. (Pacific Northwest National Laboratory)
- C-81 Yangang Liu, Shinjae Yoo, Nicholas D’Imperio (all of Brookhaven National Laboratory); and Xiaolin Li and Zheng Gao (both of Stony Brook University)
- C-83 David Randall (Colorado State University)
- C-84 Chitra Sivaraman, Laura Riihimaki, and Bill Gustafson (Pacific Northwest National Laboratory); and Giri Prakash (Oakridge National Laboratory)
- C-94 Shaocheng Xie (Lawrence Livermore National Laboratory) and Wuyin Lin (Brookhaven National Laboratory)

Page No. *Terrestrial and Subsurface Research (corresponds to Section 3.2.2)*

- C-97 Lois Curfman McInnes (Argonne National Laboratory), J. David Moulton (Los Alamos National Laboratory), David Bernholdt (Oak Ridge National Laboratory), Michael Heroux (Sandia National Laboratories), and Hans Johansen (Lawrence Berkeley National Laboratory)
- C-101 Glenn Hammond (Sandia National Laboratories) and Jeff Johnson (Lawrence Berkeley National Laboratory)
- C-104 Carl Steefel, William Collins, David Trebotich, and Hans Johansen (Lawrence Berkeley National Laboratory) and Tim Scheibe, L. Ruby Leung, Xingyuan Chen, and Alex Tartakovsky (Pacific Northwest National Laboratory)

Page No. *Oceans and Cryospheric Research (corresponds to Section 3.2.3)*

- C-109 Jeremy Fyke (Los Alamos National Laboratory)
- C-112 Daniel Martin (Lawrence Berkeley National Laboratory), Charles Jackson (University of Texas-Austin), Esmond Ng (Lawrence Berkeley National Laboratory), Stephen Price (Los Alamos National Laboratory), and Andrew Salinger (Sandia National Laboratories)
- C-116 Wieslaw Maslowski, Andrew Roberts, and Frank Giraldo (Naval Postgraduate School, Monterey, California), Elizabeth Hunke (Los Alamos National Laboratory, Los Alamos, New Mexico), and Michal Kopera (University of California-Santa Cruz)
- C-119 Title: Ocean Mesoscale Eddies

Page No. *Coupled System Integration – Earth System Models (corresponds to Section 3.2.4)*

C-121 Susan Bates (National Center for Atmospheric Research)

C-130 Phil Jones (Los Alamos National Laboratory)

C-132 Minghua Zhang (Stony Brook University)

Page No. *Integrated Assessment Modeling (corresponds to Section 3.2.5)*

C-133 Kate Calvin and Robert Link (both of Pacific Northwest National Laboratory)

C-135 Robert Jacob (Argonne National Laboratory), Andrew Jones (Lawrence Berkeley National Laboratory),
and Rao Kotamarthi (Argonne National Laboratory)C-138 Andrew Jones (Lawrence Berkeley National Laboratory), Robert Jacob (Argonne National Laboratory),
Haruko Wainwright (Lawrence Berkeley National Laboratory), Zenaida Mourao (University of Cambridge),
and Rebecca Pass (Lawrence Berkeley National Laboratory)**Page No. *Transforming Science through Exascale Capabilities: Model-Data Fusion and Testbeds
(corresponds to Section 3.2.6)****Large-Scale Heterogeneous Data Management*C-141 Forrest M. Hoffman, Jitendra Kumar, and Sarat Sreepathi (Oak Ridge National Laboratory);
William W. Hargrove (USDA Forest Service); and Richard T. Mills (Intel Corp.)*Observational Data Processing: Retrieval Algorithms and Instrument Network Simulation*C-144 Pavlos Kollias (Stony Brook University and Brookhaven National Laboratory) and Edward Luke (Brookhaven
National Laboratory)*Uncertainty Quantification and Analytics of Observations and Model Results: Uncertainty Quantification,
Statistics, Emulators, and Analytics of Model-Observation Comparisons*

C-146 Bert Debusschere, Habib Najm, Khachik Sargsyan, and Kenny Chowdhary (Sandia National Laboratories)

C-149 Charles Jackson (University of Texas at Austin)

C-152 George S. H. Pau (Lawrence Berkeley National Laboratory)

C-157 Nathan Urban (Los Alamos National Laboratory)

Model Development Testbeds

C-161 Hsi-Yen Ma (Lawrence Livermore National Laboratory)

**Page No. *Transforming Science through Exascale Capabilities: Algorithms and Computational Science
(corresponds to Section 3.2.7)****Software Engineering for Portability*

C-163 Matt Norman (Oak Ridge National Laboratory)

C-165 Matt Norman (Oak Ridge National Laboratory)

This page is intentionally left blank.

C.1 White Papers Addressing Biological Research

Multiscale Biophysical Simulations from Molecules to Cells (corresponds to Section 3.1.1)

BER White Paper

Rich Bonneau (New York University) and Iddo Friedberg (Iowa State University)

1. Please specify the current science drivers for your field of research.

Our work focuses on modeling and simulating at multiple scales, ranging from the molecular to cellular and microbe community levels. To achieve DOE's sustainable biofuel production and contaminant bioremediation mission, we wish to predict the consumption and production of substrates and metabolic products at the microbe community level. For example, the software and data platform Kbase leverages multiple data resources from genomic or metagenomic data. The results are dynamic models of cellular metabolism both in isolation and in the context of a community.

An essential activity that remains highly neglected is *in silico* protein function prediction. Function prediction must include function in the context of biochemical reactions, that is, metabolism, signal transduction, the role in cellular regulation in response to external stimuli, protein/protein interactions and the function of membrane proteins in terms of pore transport, cell-cell interactions, and the dynamics of heteromer complexes at the membrane.

Function prediction methods based on sequence similarity, operon context, and protein structure similarity can work well. But fundamentally many of these methods leverage known information to transfer annotation to query proteins. Many query protein sequences have low sequence similarity to well-annotated proteins. And it is precisely those different, organism-specific proteins that can provide information about novel functionality. Thus protein function prediction is a key aspect of our work.

2. Describe the science challenges expected to be solved in the 2020–2025 time frame using extant computing ecosystems.

Exaflop computing ecosystems will make it possible to use multi-scale modeling to simulate molecular, pathway, regulatory, cellular, community, and ecosystem processes. Simulations and modeling at this level will also guide experimental efforts to test for unknown protein functions.

Exaflop computing will allow the simulation of larger quantum regions within enzymes, capturing of protein conformational changes associated with protein function, and simulation of heteromer protein complexes in the cytosol and the membrane. In addition, the transformation of chemical energy to mechanical motion/force would be key.

Together, the advances listed above will allow predictions of cellular response to external stimuli, for example, chemotaxis, quorum sensing, endocytosis.

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

Reliable *de novo* protein folding as part of a function prediction pipeline will likely remain unsolved. Furthermore, the predictive design of cell membrane transporters that intake raw materials or efflux desirable products will likely remain an open question. In addition, systematic computation of enzymatic kinetic rate constants necessary for metabolic models may remain unfeasible.

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

| Accelerate | Why? |
|-------------------------|---|
| 1.Models and algorithms | We need algorithms that can scale well to exaflop computing. |
| 2. Hardware resources | Output of all the computing will require disk storage or at the very least temporary storage in memory. Furthermore, I/O and node intercommunication can be a significant bottleneck. |
| 3.Application codes | Codes widely used by the community need to employ algorithms that scale well into the exoscale regime so researchers can reap the benefits of exascale computing. |

| Impede | Why? |
|----------------------------------|---|
| 1.Compiler/library development | Compilers and internode communication libraries need to be compatible with legacy code and compatible with modern programming language features. |
| 2. Sharing and archiving of data | As data set sizes increase by 100- to 1000-fold, transfer, sharing, and archiving of the data will be increasingly difficult. A centralized repository does not necessarily solve the problem as data transfer over the internet will be impractical. |

Exascale Computing to Enable Predictive Models of Emergence and Self-Organization: From Metabolism to Ecosystems

William Cannon and Senghwa Kang
Pacific Northwest National Laboratory, Richland, Washington

Emergence and self-organization

In the post-genomic era, it has been recognized that genome sequences are only just the beginning of the grand challenge of determining function from sequence. It is now widely recognized that the higher level functions and phenotypes of cells emerge from the collective dynamics of smaller-scale phenomena. From collections of enzymes emerges a functional pathway, and the transient fluxes through these pathways ebb and flow until a stable steady state emerges from the collective reactions. The synchronous reaction dynamics of a cell alter its local environment and thereby other cells, until individual cells themselves synchronize through the exchange of metabolites. This hierarchy of synchronization leads to self-organization on many levels, from metabolic pathways to large-scale correlations such as the collective swarming of microbes, birds, and fish. Traditional models of biological systems are for the most part unable to provide an understanding of emergent properties and self-organization. Too often, models are focused on empirically modeling the emergent function through parameterization rather than understanding the dynamics leading to the emergent function. In order to understand emergence and self-organization, multi-scale physics-based models are needed.

Today's models

Currently, physics-based models are not frequently used for modeling microbes and plants. There are, of course, exceptions, but for the most part the models in the literature consist of either high-level, empirical models based on parameterized Monod-type kinetics or constraint-based flux models. In the former, the rate parameter determining system dynamics is obtained by fitting to data. In the latter, experimentally determined boundary fluxes are used as constraints, and the set of feasible fluxes of internal pathways is found by optimization based on fit to an empirical goal/objective function. These approaches have been very successful for understanding biological behavior under the conditions in which the measurements were made, but provide little insight into emergence and self-organization.

Physics can be loosely brought into constraint-based models as additional constraints. However, the thermodynamic constraints are based on rough assumptions about feasible concentrations in the cell. These concentrations are based on experimental assays, and the assays measure total cell concentrations—chemical species in solution and also bound to enzymes. Yet the thermodynamic constraints pertain to the overall reactions in solution only. Consequently, the thermodynamic constraints determined from total cell concentrations are usually crude estimates. Predicting dynamical behavior of organisms in this way is extremely challenging even when constraining models with experimental data.

Future models

In the future (2020–2025), these mostly empirical models of internal cell dynamics can be scaled to larger time and space dimensions and coupled in a hierarchical manner with other models.

Promising approaches include those used to build a hierarchical model of mycobacterium tuberculosis [1]. At other times, because of the complexity of the system, it may be necessary to represent the microbes or other species using summary descriptions instead of detailed representations of the biology internal to each cell. Because the details of the cell are only roughly represented, the model is not a high-fidelity model and is appropriately described as a functional guild in which the rough overall function of the organism is represented. Such models can be coupled to the environment using a mean field representation of the environment, including representation of average chemical gradients based on a continuum model. If sufficient effort is invested in the model development, the environment may even be represented as a physical structure with laminar flow of material through the environment to the organisms. Already, Biocellion [2], Morpheus [3], and CompuCell3D [4] allow for multi-cellular simulation in which the environment can be represented in detail if the effort is put into building the appropriate models.

On the current path, the ability to physically model the underlying dynamics that lead to emergent functions and self-organization will continue to be an underdeveloped area for biology. Bringing more physics into the models is especially critical for engineering organisms to produce biofuels, where it is critical to predict whether engineered pathways lead to the emergent function of over-producing target molecules. Frequently, engineered pathways do not result in the desired function because of the complexity and ability of the system to adapt. Detailed multi-scale physical models enabled by high-performance computing (HPC) would result in better predictions because this complexity and adaptivity are also emergent functions of lower-scale dynamics. While ODE/kinetic simulations based on rate parameters will incrementally expand for model organisms such as *E. coli* and *Saccharomyces*, new methods are needed that do not depend on rate parameters. Future investments should encourage the use of more detailed, physics-based models.

Vision for tomorrow: Physics-based predictive models from molecules to ecosystems enabled by exascale computing.

The development of agent-based simulation systems such as Biocellion [2], Morpheus [3], and CompuCell3D [4] suggest that the field of computational biology is ready to take the next step to a broader range of multi-scale simulations that include detailed, physics-based models. These models should address scales from molecules to cells to cultures to communities and ecosystems. Rather than empirical modeling, models that relate dynamics, energetics, and material flow are needed, while at the same time addressing complexity, self-organization, and emergence of function. Statistical thermodynamics is the theory for understanding how complexity at one scale leads to emergence of function on higher scales, and it is widely used to understand higher-order properties, from local correlated motions to swarming behaviors. While these methods were sometimes deemed to be too compute intensive for modeling more than one or two orders of magnitude in scale, exascale computing will enable these methods to span a range of scales.

A new generation of computing power can enable a significant increase in predictive power by allowing researchers to relate flux to energy and material gradients in a quantitative manner and even predict concentrations, which are the experimentally observable quantities. Already, statistical thermodynamic models have been developed to model reaction pathways [5] and gene and protein expression [6], and these models can be scaled up to model cellular-level processes such as full metabolism and gene and protein expression of individual organisms. When used as agents in programs such as Biocellion and Morpheus, which can represent the environment in detail, researchers will be able to model orders of magnitude more than merely the billions of microbes growing in only 1 mL of growth culture. More extensive systems with precise detail and predictive power will be possible if exascale computing comes to biology.

HPC-enabled, physics-based predictive models are especially needed for biosystem design and engineering. Currently, engineering of cellular processes for specific functions often produces unpredictable results. Frequently the unintended behavior is conjectured to be due to issues of regulation. But, regardless of how much regulation is engineered into a system, organisms cannot be designed that disobey the complex, adaptive dynamics of the cell. While this statement seems obvious, current modeling methods do not readily quantify the thermodynamic profiles of biological pathways that emerge because of the adaptive nature of the systems. Because cells are complex and adaptive, unintended consequences of engineering are the rule and not the exception. New models are needed that can handle the complexity and the ability of the cell to adapt. Doing so at the cellular level is an HPC problem. Modeling adaptivity at the community or ecosystem level is an exascale problem.

Of particular importance for modeling natural environments such as soils is the trade-off between the dynamics of energy dissipation and growth. While the latter leads to high carbon utilization efficiencies (CUEs) in natural systems, the former results in significant increases in production of gaseous by-products such as CO₂ and methane. These trade-offs are important to understand for the productivity of natural systems. Current models require researchers to specify the CUE rather than the model predicting the CUE. CUEs are a function of the thermodynamic gradients of the systems and are highly dependent on the final electron acceptors used by the soil microbes. Thermodynamic gradients in the environment also synchronize the metabolism between organisms—a form of self-organization. While the emergence of persistent synchronization across cells, such as circadian rhythms, has been entrained by natural selection, transient synchronizations due to energy gradients also exist and are important in the robustness of microbial and plant communities.

In fact, the exchange of essential nutrients and metabolites is driven by the need to minimize the energetic cost of growth of a microbial or plant community. By developing specialists whose role in the community is to produce a costly cofactor or metabolite, the specialists enable the rest of the community to focus on increasing productivity at lower energy costs [7]. These are all thermodynamic principles, the consequences of which we need to be able to predict. Finally, and perhaps most importantly, we need a better grasp on how energy gradients in the environment drive natural selection and adaptation. New predictive, physics-based models enabled by exascale computing will inform policies on ecosystem management to prevent ecosystem collapse using strategic interventions.

From the perspective of ‘omics data analytics, thermodynamics is the original data science in which the mathematics were developed to understand experimental observations. With the introduction of Shannon entropy, thermodynamics became the basis for information theory as well. Information entropy and thermodynamic entropy are identical when the same probability distribution is used. Most recently, information theory has become a foundational approach for data analysis using deep learning. In fact, exascale computing may enable researchers to integrate detailed physical models with machine learning, and this would accelerate technology development and discovery in many areas of biology.

What will not be possible even with exascale are all-atom simulations of cells.

The *top three* computing ecosystem aspects that will accelerate or impede progress in the next 5–10 years are (in order of precedence) (1) application codes (implementation, development, portability, etc.), (2) models and algorithms, and (3) workforce development (current training for biologists does not prepare them to understand the science from a fundamental level). Moreover, the common factor for success in all three aspects is that the funding models need to change to regard those biologists and computational scientists who have significant training in both quantitative science and biology as not only service providers but also biological domain experts. It is the slow churning of the culture of traditional biology that prevents rapid progress.

References

1. Karr JR, et al. (2012) A whole-cell computational model predicts phenotype from genotype. *Cell* 150(2):389–401.
2. Kang S, Kahan S, McDermott J, Flann N, and Shmulevich I (2014) Biocellion: accelerating computer simulation of multicellular biological system models. *Bioinformatics*.
3. Starruss J, de Back W, Bruschi L, and Deutsch A (2014) Morpheus: a user-friendly modeling environment for multiscale and multicellular systems biology. *Bioinformatics* 30(9):1331–1332.
4. Swat MH, et al. (2012) Multi-scale modeling of tissues using CompuCell3D. *Methods Cell Biol* 110:325–366.
5. Thomas DG, Jaramillo-Rivera S, Baxter DJ, and Cannon WR (2014) Comparison of Optimal Thermodynamic Models of the Tricarboxylic Acid Cycle from Heterotrophs, Cyanobacteria, and Green Sulfur Bacteria. *J Phys Chem B*.
6. Rydenfelt M, Cox RS, Garcia H, and Phillips R (2014) Statistical mechanical model of coupled transcription from multiple promoters due to transcription factor titration. *Phys Rev E* 89(1).
7. Tasoff J, Mee MT, and Wang HH (2015) An Economic Framework of Microbial Trade. *PLoS One* 10(7):e0132907.

BER White Paper
X. Cheng, Oak Ridge National Laboratory

1. Please specify the current science drivers for your field of research.

Where we are today? Be sure to include broad impact, DOE interest, ties between experiment/theory, etc.

My current research is computational biochemistry and biophysics, with application to bioenergy (biofuel production from lignocellulosic biomass, photo system I for light to electricity conversion), and carbon fixation (bacterial microcompartment, i.e., carboxysome self-assembly), etc. We have run petaflop capability molecular dynamics simulations of lignocellulosic biomass materials (to tackle the biomass recalcitrance problem) and complex biomembrane systems (to interpret neutron scattering experiments), and nearly 1,000 free energy perturbation simulations (to determine the ligand-receptor binding affinities). The former two systems typically consist of several million to a few tens of million atoms, and the simulations last for 1–2 microseconds. The latter systems typically consist of a few hundred thousands of atoms, and each simulation lasts for up to 100 ns.

2. Describe the science challenges expected to be solved in the 2020–2025 time frame using extant computing ecosystems.

What will probably be solved in the next 5–10 years? Why is this important to the field?

In the next 5–10 years, multiscale physical simulations of a whole cell or part of a cell, such as bacterial cell membrane (envelop), plant cell wall, and key components of a cellular pathway, will become feasible. These simulations will shed significant light on how the plant cell wall can be degraded more efficiently for biofuel production, and how cell membranes can be adapted for coping with environmental stress, etc. In the 2020–2025 timeframe, I envision computational systems biology will rely more and more on physical (spatial, temporal, and reaction) descriptions of the underlying systems. The traditional computational systems biology models will be bridged with atom/molecular level physics-based models, such as molecular dynamics and Brownian dynamics simulations, so as to become less dependent on parameters, and thus more predictive.

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

What might not necessarily be solved in the next 5–10 years? Again, what is the importance?

In the next 5–10 years, a full particle-based and functional model of an entire cell will still be impossible. Or even if possible in some special cases, the models will still be very much limited by the timescale that can be accessed. Physical simulations, such as molecular dynamics and Brownian dynamics simulations, will not be able to make predictions about the function of the cell or how a mutation may perturb the function/behavior of the cell, questions that are of direct interest to experimental systems biologists and microbiologists.

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

- a. Application codes (implementation, development, portability, etc.)
- b. Models and algorithms
- c. Hardware resources (at all scales) including I/O, memory, etc.
- d. Data workflow (including sharing, transmitting, archiving, etc.)
- e. Visualization and analysis resources
- f. Internal/external libraries/frameworks
- g. Workforce development
- h. Other

| Accelerate | Why? |
|---|---|
| 1. Hardware resources (at all scales), including I/O, memory, etc. | Hardware resources are the basis. The extreme-scale computer systems will lay a foundation to enable physical simulations at greater spatial and timescales. |
| 2. Models and algorithms | The availability of extreme-scale capability computer systems, together with the goal of solving more complex scientific challenges, will require a paradigm shift in computational models and algorithms, including multiscale, multiphysics models, and more scalable algorithms. |
| 3. Application codes (implementation, development, portability, etc.) | Application codes are key to accelerating our progress because they are the tools we use to produce results and make predictions. The implementation makes a great difference for the performance of even the same model/algorithm. |

| Impede | Why? |
|--------------------------|--|
| 1. Workforce development | The emerging area of computational biophysics requires the next-generation workforce to possess truly interdisciplinary knowledge in biology, chemistry, and physics, as well as a deep understanding of computer architectures, programming models, and simulation algorithms. The lack of a qualified workforce in this area will impede our progress. |

5. (Optional) Characterize the data ecosystem aspects if the primary drivers for your field of research involve the transmission, analysis (including real-time analysis), or processing of data.

For example:

- Classify the data as simulation, experimental, both, or something else.
Simulation data.
- Characterize these data using the 3 V's (Velocity, Volume, and Variety).
Coordinates and velocities of particles.
How much data will be generated/stored within the next 5–10 years?
100 terabytes.
What will the data rates be?
Typically a few gigabytes per hour.
Will the data be multi-modal somehow?
No.
- Describe current or planned solutions for any of these data challenges, e.g., in situ analysis.
In situ analysis, hierarchical storage, and data compression.
- Note any particular data security or privacy requirements.
N/A

6. References (please keep the reference length to no more than 10 critical references)

Hu X, Hong L, Smith MD, Neusius T, Cheng X, and Smith JC, 2016, "The Dynamics of Single Protein Molecules Is Non-equilibrium and Self-similar over Thirteen Decades in Time," *Nature Phys.* 12, 171–174.

Smith MD, Mostofian B, Cheng X, Petridis L, Cai CM, Wyman CE, and Smith JC, 2016, "Cosolvent Pretreatment in Cellulosic Biofuel Production: Effect of Tetrahydrofuran-water on Lignin Structure and Dynamics," *Green Chem.*, Advance Article.

Nickels JD, Cheng X, Mostofian B, Stanley C, Lindner B, Heberle FA, Perticaroli S, Feygenson M, Egami T, Standaert RF, Smith JC, Myles DA, Ohl M, and Katsaras J, 2015, "Mechanical Properties of Nanoscopic Lipid Domains," *J Am Chem Soc.* 137(50):15772–15780.

Mahalik JP, Brown K, Cheng X, and Fuentes-Cabrera MA, 2016, "Theoretical Study of the Initial Stages of Self-assembly of a Carboxysome's Facet," *ACS Nano*, In press.

7. (Optional) Images

Consider submitting one or two already published high-resolution images suitable for inclusion in the report. Please provide the reference(s). Submit these separately from the two-page report; they will not count against the page limit.

**BER White Paper:
M.S. Cheung, University of Houston**

1. Please specify the current science drivers for your field of research.

Where we are today? Be sure to include broad impact, DOE interest, ties between experiment/theory, etc.

The research of my group focuses on the dynamics of bio macromolecules in living systems [1–3]. The interior of a cell is a highly crowded space occupied by diverse kinds of macromolecules. However, it is still unclear as to how information passes through intracellular space, or how the transmission of information between the exterior and interior of a cell guarantees fidelity. These are the key questions underlying a grand “signal transduction” in biology. One of the challenges is that we do not completely know either experimentally or theoretically about the composition of a cell at various stages of its survival. However, with the development of theory and computation, we can provide useful insights into the molecular mechanisms of the biomolecular networks that can be tested experimentally.

Another important part of my research focuses on the solar energy conversion of artificial photosynthetic materials (APM) [4–6]. We investigate the structure-function relationship between the conformations of APMs and the efficiency of charge-transfer in various types of solvent and liquid-solid interfaces. We gain valuable insights from these simulations about materials design in heterogeneous environments.

2. Describe the science challenges expected to be solved in the 2020–2025 time frame using extant computing ecosystems.

What will probably be solved in the next 5–10 years? Why is this important to the field?

I believe that in the next 5–10 years, we will be able to simulate the transport of “some” information in a cell, such as the relation between motor proteins and vesicular dynamics, or synaptic plasticity in dendrite spines of neurons. These subcellular dynamics rely heavily on information at the molecular level – which we can simulate very well nowadays.

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

What might not necessarily be solved in the next 5–10 years? Again, what is the importance?

In the next 5–10 years, we probably still cannot simulate the development of a cancerous cell that requires a comprehensive understanding of “cellular wiring” that starts from mutations of DNAs.

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

- a. Application codes (implementation, development, portability, etc.)
- b. Models and algorithms
- c. Hardware resources (at all scales) including I/O, memory, etc.
- d. Data workflow (including sharing, transmitting, archiving, etc.)
- e. Visualization and analysis resources
- f. Internal/external libraries/frameworks
- g. Workforce development
- h. Other

| Accelerate | Why? |
|---|---|
| 1. Hardware resources (at all scales) including I/O, memory, etc. | Able to scale up the size of the simulations. |
| 2. Models and algorithms | Able to scale up the size of the simulations. |
| 3. Workforce development | Able to move the project forward. |

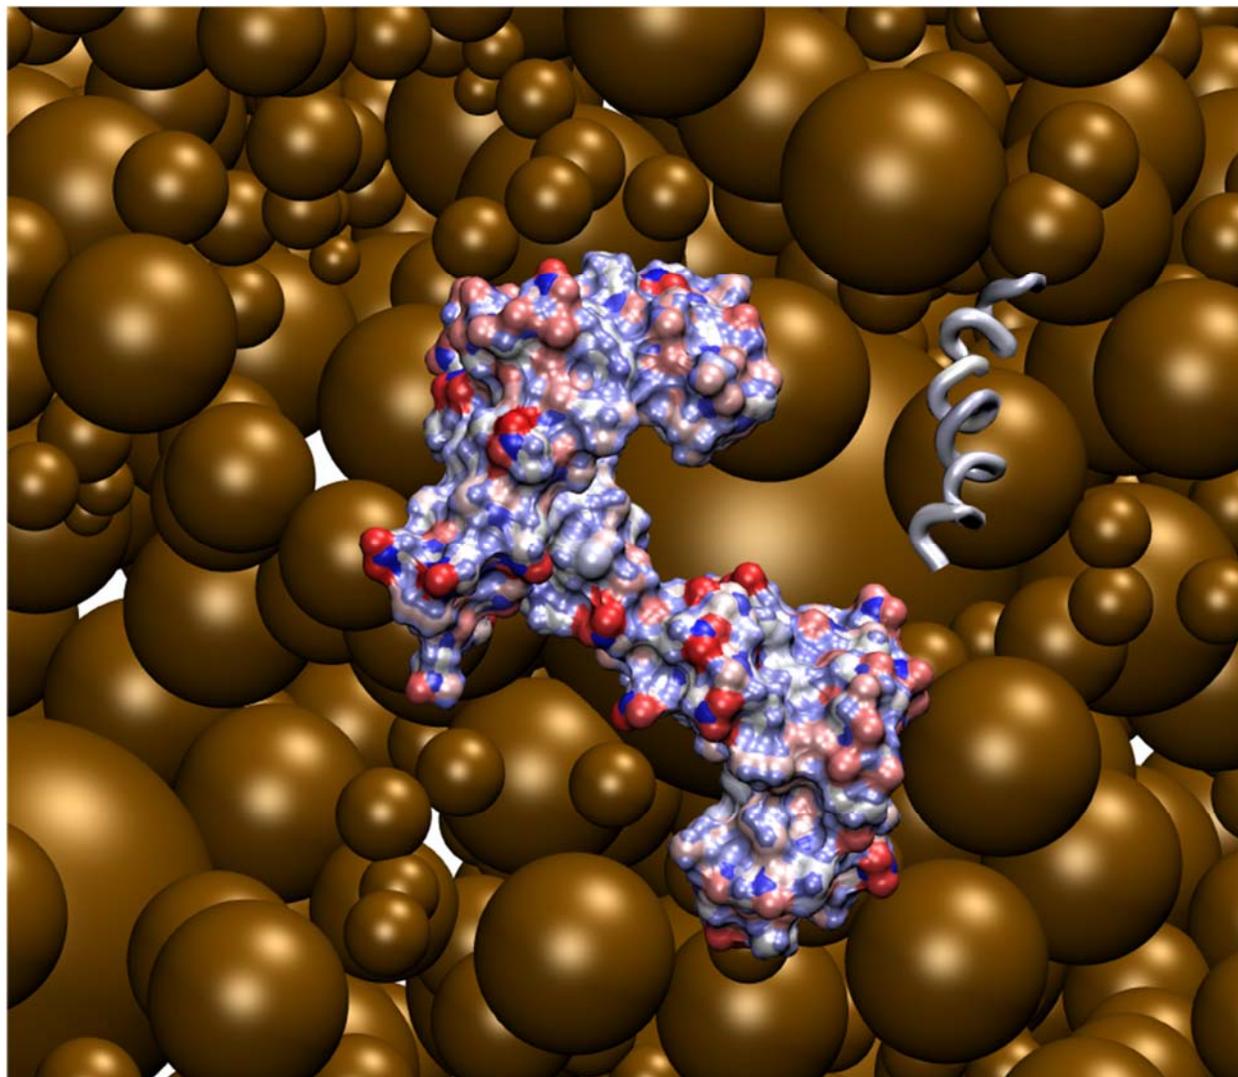
| Impede | Why? |
|---|---|
| 1. Internal/external libraries/frameworks | Lots of back compatibility issues. |
| 2. Data workflow (including sharing, transmitting, archiving, etc.) | The size of the data is very large; storage and transfer are problematic. |
| 3. Application codes | Portability to accelerators is a hurdle. |

1. References (please keep the reference length to no more than 10 critical references)

- Dhar, A., A. Samiotakis, S. Ebbinghaus, L. Nienhaus, D. Homouz, M. Gruebele, and M.S. Cheung, 2010, "Structure, Function and Folding of Phosphoglycerate Kinase Are Strongly Perturbed by Macromolecular Crowding," *Proc. Natl. Acad. Sci. U.S.A.*, 107:17586–17591.
- Wang, Q., and M.S. Cheung, 2012, "A Physics-based Approach of Coarse-graining the Cytoplasm of E. coli (CGCYTO)," *Biophys. J.*, 102:2353–2361.
- Wang, Q., P. Zhang, S. Tripathi, L. Hoffman, L. Yin, M.N. Waxham, and M.S. Cheung, 2013, "Protein Recognition and Selection through Conformational and Mutually Induced Fit," *Proc. Natl. Acad. Sci. U.S.A.*, 110:20545–20550.
- Su, G., A. Czader, D. Homouz, G. Bernardes, S. Mateen, and M.S. Cheung, 2012, "Multiscale Simulation on a Light-harvesting Molecular Triad," *J. Phys. Chem. B.*, 116:8460–8473.
- Balamurugan, D., A.J.A. Aquino, F. De Dios, L. Flores, Jr., H. Lischka, and M.S. Cheung, 2013, "Multiscale Simulation of the Ground and Photo-induced Charge-separated States of Molecular Triad in Polar Organic Solvent: Exploring the Conformations, Fluctuations and the Free Energy Landscapes," *J. Phys. Chem. B.*, 117:12065–12075.
- Manna, A.K., D. Balamurugan, M.S. Cheung, and B.D. Dunietz, 2015, "Unraveling the Mechanism of Photo-induced Charge-transfer in Carotenoid-porphyrin-C60 Molecular Triad," *J. Phys. Chem. Lett.*, 6: 1231–1237.

2. (Optional) Images

Consider submitting one or two already published high-resolution images suitable for inclusion in the report. Please provide the reference(s). Submit these separately from the two-page report; they will not count against the page limit.



**BER White Paper:
M. Crowley, National Renewable Energy Laboratory**

1. Please specify the current science drivers for your field of research.

Where we are today?

Biofuels research is currently focused on using biological approaches to producing fuels and materials from renewable sources. This has been accomplished thus far in creating successful industries for first-generation (corn ethanol) and second-generation (ligno-cellulosic ethanol) fuels. Current research is aimed at advanced drop-in fuels (hydrocarbons) from plants, which are predominantly from plant cell walls. A large amount of research has been conducted in all aspects of this conversion from cell walls to fuels and materials, including genetic design of plant morphology, pretreatment, breaking down cell wall polymers to monomers, and upgrading monomers to valuable fuels, chemicals, and materials. The role of computational research is to understand the cell wall, predict useful changes to the wall, model the breakdown of components, and model the metabolic processes in the upgrading processes. All efforts are ultimately targeting high carbon yields and titers, and lower overall costs of renewable energy and materials.

The research is highly tied to experiment in both the discovery and predictive roles. Experiments that are highly coupled to the modeling and simulation are cell wall imaging (TEM, SEM, TIRF, AFM, and many others), x-ray and neutron diffraction (wide and small angle), NMR, crystallography, proteomics, transcriptomics, and all manner of expression and evaluation of predictions. The field of molecular modeling has been reaching the time and length scales of many imaging technologies, and this connection is essential to validation and iteration of both experiment and simulation for new understandings and testing of hypotheses.

2. Describe the science challenges expected to be solved in the 2020–2025 time frame using extant computing ecosystems.

What will probably be solved in the next 5–10 years? Why is this important to the field?

The structure and nature of plant cell walls, and likely yeast, bacterial, and fungal cell walls, are the most likely problems to be solved by the combination of computing and experiment. The complex structure of the plant cell wall is extremely difficult to determine with its intertwining of multiple polymer components: cellulose, hemicellulose, lignin, pectin, and glycoproteins. This understanding can lead to understanding the role of each polymer in the wall and how it can be modified for better energy and material goals without compromising the function of the cell wall in a viable plant.

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

What might not necessarily be solved in the next 5–10 years? Again, what is the importance?

It is unlikely that the understanding of how a cell wall is constructed (i.e., the mechanism or process) in a plant will be attained in the next 5 years with the current computing ecosystems. This is an essential part of using plant material for renewable purposes, as we have to deconstruct the cell walls to produce other products. Deconstruction is a major obstacle to economic production of plant-based products, and understanding the detailed structure is the key to learning how to economically deconstruct. Until now, successes have been mainly from blindly following the lead of fungi and bacteria in their enzymatic approaches to deconstruction. Understanding the wall and producing viable computational models are key to engineering more efficient processes to deconstruction and engineering plants to build more easily deconstructed and more highly tuned cell walls.

This process involves time and length scales that are not accessible in the current environment. For cellulose alone, it includes understanding the cellulose synthase mechanism, the complex for producing fibers, the mechanism of crystallization of polymers into cellulose-I rather than the more stable II and III

versions, and the depositing of nanofibrils into the primary and secondary walls. Further, an understanding of how growing cells expand and how hemicellulose is deposited in the primary and then secondary walls is required, and finally, how lignin is deposited and polymerized in a forming or formed wall. The function and structural nature of the other polymers of pectin and glycoproteins are part of the whole mix as well.

Further, harnessing the power of organelles is unlikely to be attained without an understanding of their mechanistic processes, especially compartmentalization. This applies for golgi apparatus and endoplasmic reticulum processes in which different parts of the creation of hemicellulose and lignin are sequestered to different parts of ER and GA, and transported back and forth across the membranes in concerted and sequential fashion. The development of approaches has other far-reaching benefits of understanding light harvesting through chloroplasts, hydrogen production, chemical energy transduction in mitochondria, and active transport and trafficking of raw material and products in and out of industrial microbes.

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

| Accelerate | Why? |
|---|--|
| 1. Quantum Codes: made parallel and more efficient MD Codes: new sampling methods implemented, multiple time-scale and length-scale interfaces implemented, portable to exascale hardware QM/MM efficiency and accuracy improvement to remove limitations of poor scalability. | Current QM codes only scale to tiny numbers of processors and hold up most modeling regardless of available compute resource. This holds back both pure QM and QM/MM studies and keeps extremely important problems from harnessing the compute resources even currently available. |
| 2. Coarse-graining improved, and to include validity checking (CG only valid within certain limits). Multiscale algorithms designed for crossing scale boundaries, especially from longer to shorter. | Current CG is mostly an art and done correctly by very few of the most experienced and knowledgeable. When put in the hands of the masses, it is used incorrectly and outside the bounds of the validity due to approximations used in the CG process. There are many approaches to CG, and it is essential for multiscale, or at least for larger time- and length-scale modeling and simulation. It is essential to push this art into a more routine and robust implementation which will include preventing misuse of CG for scientific prediction and hypothesis testing. Crossing boundaries during coupled high- and low-scale simulations is also a frontier that has not been addressed systematically and with robust solutions. It is not as hard to take small-scale and apply the CG mapping to the large-scale models. It is a MUCH harder task to map back in the one-to-many direction of either atoms or ensembles, yet it is crucial for a true multiscale model of the future. |
| 3. Data Analysis: Development of new approaches to the vast amount of data coming out of larger compute resources will greatly accelerate the usefulness of the data. One approach is to do much of the data analysis as the data are created rather than store the data and analyze later. Buyer beware: this might bias the results to what the modeler is looking for and miss what is really interesting. | Currently, most modeling and simulation data are stored for analysis later. This storage process is expensive and gets prohibitive as we move to high-petascale and exascale. In analyzing on the fly, we save only what is interesting and important, or gather the statistics, or guide the simulation, much as you would in any experiment. Thus much of the analysis, which can take as much time and compute power as the simulation itself, is already done when the simulation finishes and the uninteresting data are either not stored or not even simulated. |

| Impede | Why? |
|--|--|
| <p>1. Lack of software development by research teams, lack of support and lack of qualified individuals with both computer programming and domain-specific expertise such as physics or chemistry.</p> | <p>Funding is largely available for science only, not for development of new approaches, algorithms, performance, and scaling. Existing approaches do not work well on current and future architectures except a few specific scientific problems. Those problems seem to get the support, and many hugely important problems do not get access to the HPC hardware because they do not scale. Resources need to be put into the problems that do not scale, yet.</p> <p>Without domain scientists who are trained and supported to improve and create scientific HPC software, the toolsets remain the same and new problems do not get solved.</p> |

Building predictive spatiotemporal models of cellular processes from fluorescence microscopy data

Ron Dror, Stanford University (with input from Mark von Zastrow, UCSF)

Introduction

The spatial organization of molecules within a cell is of fundamental importance to most physiological processes, yet it remains poorly understood [1]. Advances in structural biology have yielded tens of thousands of atomic-level structures of proteins and other biomolecules, leading to over a dozen Nobel Prizes; the 2013 Chemistry Nobel Prize recognized computational work in this area. By contrast, revealing how these molecules are organized within a cell has proven much more difficult, in part because many of the cell's components are highly mobile.

Substantial recent progress in understanding cellular organization has resulted from dramatic advances in fluorescence microscopy, a family of techniques in which molecules of interest are labeled by attachment of fluorescent probes and then imaged under an optical microscope to reveal their spatial distribution [2]. Thousands of experimental labs are now generating large volumes of fluorescence microscopy data, which can reveal both static spatial distributions of molecules and trajectories of individual molecules. The amount of information obtainable in a single microscopy study is limited, however; one cannot simultaneously observe all the molecules involved in a complex cellular signaling process with high spatial and temporal resolution. Moreover, microscopy alone is purely descriptive; it does not capture underlying mechanisms or predict how the spatial organization and dynamics of the cell would change in response to myriad perturbations.

We plan to develop a theoretical framework for accurate mathematical modeling of cellular processes—in space as well as time—on the basis of fluorescence microscopy data. This framework will allow us to combine the results of many microscopy experiments into a single, cohesive description of the relevant molecular-level spatial organization, dynamics, and interactions. Most importantly, it will allow us to predict the effects on cellular processes of drugs, disease-causing mutations, oxidative and metabolic stress, and cell maturation. Our theoretical framework will thus provide a guide for further experimental investigation of the cell's spatial organization.

Adrenergic signaling and endocytosis

Although we intend to develop broadly applicable methodologies, we will focus our initial efforts around the cellular signaling pathway mediated by the β_2

Key Points

- The recent explosion of fluorescence microscopy approaches has generated a wealth of spatiotemporal data, but creating quantitative, predictive models of cellular processes from these data will require new theoretical approaches.
- Developing theoretical approaches to characterize the spatial organization of GPCR signaling is particularly timely, because experimental evidence that contradicts the textbook role of receptor internalization is bringing about a sea change in the GPCR field and raising crucial but difficult questions about the role of spatial organization in cellular signaling pathways.
- The β_2 -adrenergic receptor (β_2 AR) serves as an ideal GPCR for this purpose, because experimental evidence demonstrates a critical role for subcellular organization in β_2 AR-mediated signaling and because particularly powerful experimental tools are available for testing model predictions in this signaling pathway.

adrenergic receptor (β_2 AR) (Fig. 1). This receptor, which is a target of beta blockers and beta agonists, has long served as an archetypal G-protein-coupled receptor (GPCR). The GPCRs represent both the largest family of human membrane proteins and the largest family of drug targets and are thus of great scientific and pharmaceutical interest.

Recent experimental work on this pathway has suggested a paradigm-changing idea in biology: that internalization of cell-surface receptors does not serve to silence signaling, as had long been assumed, but that a crucial phase of receptor signaling in fact takes place

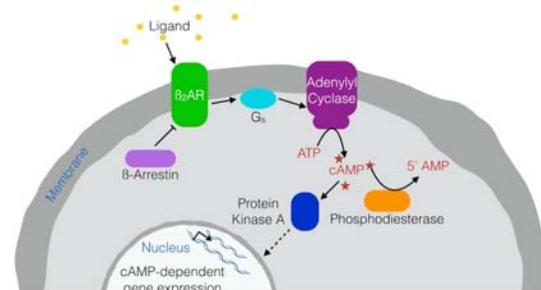


Figure 1. A cellular signaling pathway mediated by the β_2 -adrenergic receptor (β_2 AR). A number of cellular processes, including gene expression, depend on the spatial distribution of cAMP, as well as its overall concentration.

after the receptors are internalized into the cell within endosomes [3, 4]. Other studies have shown that cyclic adenosine monophosphate (cAMP)—the primary second messenger molecule modulated by β -adrenergic receptor signaling—is non-uniformly distributed within the cell, and that its distribution is critical to the signaling processes it controls [5, 6]. We have found that signaling from endosomes affects the spatial distribution of cAMP (Fig. 2) and has a downstream functional impact on gene regulation.

These findings raise a host of critical but difficult-to-answer questions. Is receptor-mediated endocytosis related to the spatial distribution of activated receptors? How is the spatial distribution of cAMP determined, and what role does endocytosis play? How do the many drugs that target β -adrenergic receptors alter these processes? These drugs tend to affect both cAMP production and receptor endocytosis, but not always in the same way.

We will address these questions by building spatiotemporal models for this signaling process on the basis of fluorescence microscopy data collected in the von Zastrow lab and in the labs of additional collaborators. These data include confocal microscopy, total internal reflectance fluorescence (TIRF) microscopy, and single-particle tracking, which collectively capture motions and interactions of key molecules involved in the signaling pathway of interest: β_2 AR, G_s , β -arrestin, adenylyl cyclase, cAMP, phosphodiesterase, and protein kinase A.

We will use our models to predict the behavior of the signaling pathway in response to various perturbations, thus guiding further fluorescence microscopy studies that will serve both to validate the

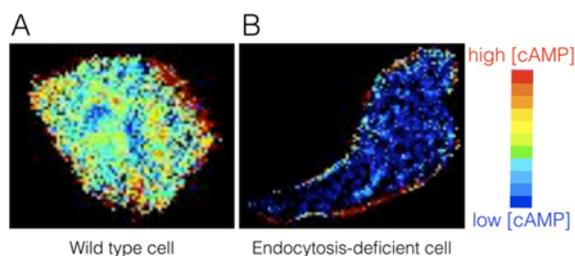


Figure 2. Endocytosis affects cAMP distribution. The spatial distribution of cAMP differs substantially between (A) a wild-type HEK293 cell and (B) a HEK293 cell with a mutant (K44E) dynamin that prevents endocytosis (unpublished results from the von Zastrow lab). In additional unpublished work using a light-activated adenylyl cyclase, the von Zastrow lab has found that production of cAMP at endosomes has a substantially greater effect on cAMP-dependent gene expression than production of cAMP at the plasma membrane.

models and to deepen our scientific understanding of this important biological system. Indeed, β_2 AR-mediated signaling constitutes an unusually well-studied experimental system, enabling us to effect many carefully controlled perturbations in a wet lab setting in order to test model predictions. For example, we can apply a wide variety of drugs with diverse pharmacological effects on the receptor (including full agonists, partial agonists, and inverse agonists, as well as “biased” ligands that differentially affect G-protein signaling and endocytosis [7]). We can create a number of well-characterized receptor mutants, including one that prevents receptor internalization but preserves the receptor’s ability to activate G proteins [8]. We can genetically manipulate other proteins in the signaling pathway and various components of the machinery involved in receptor internalization and sorting [9]. We anticipate that our models will also drive experimentalists to explore quantitative aspects of signaling pathways that will provide insight into the conditions that are required for a cell to perform a specific function or take on a particular role.

Modeling approach

Our modeling approach is inspired by reaction-diffusion simulations, in which molecules diffuse about the cell—alone or in complex with others—and pairs of molecules that come into close proximity have some probability of reacting with one another. Although such simulations (implemented with either a particle-based method or a finite-volume method) have been used to study biological systems for some time [10–15], their accuracy and utility have in many cases been limited by uncertainty in the underlying model assumptions [16]. In the absence of a method to determine from experimental data the initial spatial distribution of the various types of molecules involved, the dynamic processes characterizing their motion, and the rules governing their interactions, model assumptions are often oversimplified, leading to discrepancies between simulation results and biological reality.

We will use fluorescence microscopy data to build more general and more accurate models (Fig. 3). First, we will determine the simulation’s initial conditions by using microscopy data to determine realistic spatial distributions of each relevant molecular species. Second, instead of assuming that these molecules move by simple diffusion, we will characterize their motions statistically using single-particle tracking data. Third, by comparing the evolution of simulated systems to time series of fluorescence microscopy images, we will tune simulation parameters, including those governing interactions between molecules.

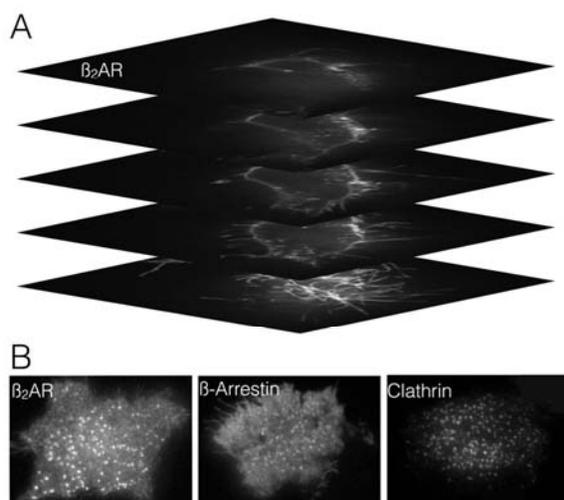


Figure 3. Samples of fluorescence microscopy data that we intend to use for modeling purposes. (A) A stack of confocal fluorescence microscopy images with labeled β_2 AR and (B) TIRF microscopy images with labeled β_2 AR, β -arrestin, and clathrin. All images are of HEK293 cells.

Accomplishing these tasks will require us to address several mathematical modeling challenges. We must determine appropriate functional forms and corresponding fitting methods, balancing the level of biological complexity our models can capture and the amount of data required to fit them. We need to develop statistical methods for comparing simulation state to fluorescence microscopy images; our simulations will be stochastic, so we cannot expect pixel-for-pixel matches. We must combine microscopy data from many different cells, each of which will have a different shape and a different internal structure, even if they are genetically and functionally identical. In addition, we must account for the fact that different proteins are overexpressed in different experiments, as most fluorescence microscopy experiments involve overexpression of fluorescently labeled proteins.

Computational requirements

Particle-based reaction-diffusion simulations are computationally expensive. In MCell [12, 15]; for example, a test simulation with 200,000 particles took about 24 hours per simulated minute on a modern CPU. We would like to do simulations of cells whose volumes are one hundred to one thousand times larger, with ten times higher particle density, and for approximately ten times longer. This requires scaling up the computation by 10^5 to 10^6 (scaling with particle density is quadratic). In addition, because these simulations are stochastic in nature, we need to run each simulation many times.

This will require both supercomputing resources and new, parallel codes (or substantially rewritten versions of existing codes). Additional work is necessary to determine optimal hardware architectures.

References

- Harold, F.M., *Molecules into cells: specifying spatial architecture*. Microbiol Mol Biol Rev, 2005. **69**(4): p. 544–564.
1. Sigrist, S.J., and B.L. Sabatini, *Optical super-resolution microscopy in neurobiology*. Curr Opin Neurobiol, 2012. **22**(1): p. 86–93.
 2. Irannejad, R., et al., *Conformational biosensors reveal GPCR signalling from endosomes*. Nature, 2013. **495**(7442): p. 534–538.
 3. Vilardaga, J.P., F.G. Jean-Alphonse, and T.J. Gardella, *Endosomal generation of cAMP in GPCR signaling*. Nat Chem Biol, 2014. **10**(9): p. 700–706.
 4. Nikolaev, V.O., et al., *Beta2-adrenergic receptor redistribution in heart failure changes cAMP compartmentation*. Science, 2010. **327**(5973): p. 1653–1657.
 5. Rich, T.C., et al., *Cyclic nucleotide-gated channels colocalize with adenylyl cyclase in regions of restricted cAMP diffusion*. J Gen Physiol, 2000. **116**(2): p. 147–61.
 6. Rosenbaum, D.M., S.G. Rasmussen, and B.K. Kobilka, *The structure and function of G-protein-coupled receptors*. Nature, 2009. **459**(7245): p. 356–363.
 7. Hausdorff, W.P., et al., *A small region of the beta-adrenergic receptor is selectively involved in its rapid regulation*. Proc Natl Acad Sci USA, 1991. **88**(8): p. 2979–2983.
 8. Di Fiore, P.P., and M. von Zastrow, *Endocytosis, signaling, and beyond*. Cold Spring Harb Perspect Biol, 2014. **6**(8).
 9. Andrews, S.S., *Spatial and stochastic cellular modeling with the Smoldyn simulator*. Methods Mol Biol, 2012. **804**: p. 519–542.
 10. Angermann, B.R., et al., *Computational modeling of cellular signaling processes embedded into dynamic spatial contexts*. Nat Methods, 2012. **9**(3): p. 283–289.
 11. Coggan, J.S., et al., *Evidence for ectopic neurotransmission at a neuronal synapse*. Science, 2005. **309**(5733): p. 446–451.
 12. Cowan, A.E., et al., *Spatial modeling of cell signaling networks*. Methods Cell Biol, 2012. **110**: p. 195–221.
 13. Roberts, E., et al., *Noise contributions in an inducible genetic switch: a whole-cell simulation study*. PLoS Comput Biol, 2011. **7**(3): p. e1002010.

14. Stiles, J.R., and T.M. Bartol, *Monte Carlo methods for simulating realistic synaptic microphysiology using MCell*, in *Computational Neuroscience: Realistic Modeling for Experimentalists*, E.D. Schutter, Editor. 2001, CRC Press: Boca Raton. p. 87–127.
15. Ridgway, D., G. Broderick, and M.J. Ellison, *Accommodating space, time and randomness in network simulation*. *Curr Opin Biotechnol*, 2006. **17**(5): p. 493–498

BER White Paper
Multiscale Biology, Gnana Gnanakaran, Los Alamos National Laboratory

1. Please specify the current science drivers for your field of research.

GENERAL STATEMENT:

The current state-of-the-art externally links models at many scales, producing “linked multi uni-scale models.”

MY SPECIFIC FIELD OF RESEARCH:

- Modeling and simulations of *model* membranes of organelles, plants, mammals, and bacteria.
 - Currently, all-atom MD simulations can handle only a few lipid mixtures/compositions. Often, coarse-grained approaches such as MARTINI force fields are used to address more complicated lipid mixtures without proper calibration using all-atom models and experiments.
- Mechanistic details of *individual* molecular assemblies or molecular machines such as efflux pumps, ribosome, viral capsids, etc.
 - Currently, only short-time dynamics of large molecular assemblies are possible. Often approaches such as targeted/steered MD are used to drive the system to gain biologically relevant insights.

2. Describe the science challenges expected to be solved in the 2020–2025 time frame using extant computing ecosystems.

GENERAL STATEMENT:

Multi-scale models support and enable knowledge discovery from *Sequence to Models of Organisms and Communities*. Whereas advances in computing will increase the size of the systems (length scales) and time scales for the simulations, the real benefit of ASCR will be enabling direct interconnections of multiple model scales, and relating the output to observations. This will request coupling to heterogeneous data at multiple scales. It will prompt the development of mathematics, machine-learning approaches, and computational tools. Coupling multi-scale simulations to observations enables researchers to solve inverse problems. As for uncertainty quantification (UQ), fitting computer models to data will require multiple runs of the code, sometimes using different models.

MY SPECIFIC FIELD OF RESEARCH:

- Modeling and simulations of *biologically meaningful* membranes of organelles, plants, mammals, and gram-negative bacteria.
 - Ability to characterize trans-membrane systems with complex lipid mixtures with all-atom MD simulations. Extreme length scales are treated with coarse-grained simulations based on all-atom simulations of very large-length systems.
- Mechanistic details of *multiple* molecular assemblies or molecular machines working together in an organism and integration of scales.
 - For example, lack of molecular-level understanding of the operating principles governing the transport properties of antibiotics across bacterial membrane has limited the full potential of therapeutic efforts. Characterization of intrinsic resistance mechanisms serving as barriers to transporting and keeping antibiotics inside the bacterium requires the simultaneous consideration of many components (LPS, peptidoglycan, etc.), multiple molecular assemblies (porins, channels, etc.), and machines (transporters, efflux pumps) (see figure).

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

One challenge is achieving high-throughput, multi-scale MD simulations of dynamical systems that can couple to high-throughput characterization of genomic, metabolic, and proteomic data in real time. Ab-initio structure prediction from folding of large, complex multi-domain proteins still remains a challenge.

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

- a. Application codes (implementation, development, portability, etc.)
- b. Models and algorithms
- c. Hardware resources (at all scales) including I/O, memory, etc.
- d. Data workflow (including sharing, transmitting, archiving, etc.)
- e. Visualization and analysis resources
- f. Internal/external libraries/frameworks
- g. Workforce development
- h. Other

1. Computing platform

For MD simulations, there is currently a healthy tension between general-purpose hardware platforms and specialized computing platforms (like Anton). Is it possible for these platforms to converge over the next 10 years? Depending on the answer to this question, one may need to consider investing in special-purpose architectures.

2. Application Code

Molecular dynamics application packages such as GROMACS need to go beyond the traditional parallelization of intra-node and inter-node. This will require complicated parallel software architecture that has been adapted for heterogeneous computing environments. Enhanced sampling approaches that rely on exchange between multiple simulations need to have automated runtime support for fail-proof production runs. In addition, better implementation of fast multipole methods for long-range interactions is desired.

3. Post-processing of Data

While an increase in computational power has the potential to reduce the need for saving results from simulations (as the results from the simulation can be readily be reproduced at low cost), there are benefits to archiving individual simulations to allow post-processing. Already, examples of such post-processing have suggested ways to couple the simulations across scales. We envision that for research purposes, researchers will save and analyze large simulation outputs using machine learning, statistics, and other modern information processing techniques. Big data challenges are likely to arise as this line of research takes off. From an infrastructure perspective, there will be a need to process large distributed datasets.

5. (Optional) Characterize the data ecosystem aspects if the primary drivers for your field of research involve the transmission, analysis (including real-time analysis), or processing of data.

The challenge of big data is the ability to triage information and organize it into hierarchies of value of information. This is likely to require the manipulation of complex relational objects (networks/graphs), and the development of problem-specific filters. The former will require computations that are memory intensive; the latter will increase the storage needs, possibly by an order of magnitude. An optimal computational ecosystem will need to strike a balance between these two requirements, while managing I/O needs and data transfer issues. Distributed systems may be a reasonable solution, provided fast and reliable interconnects are available.

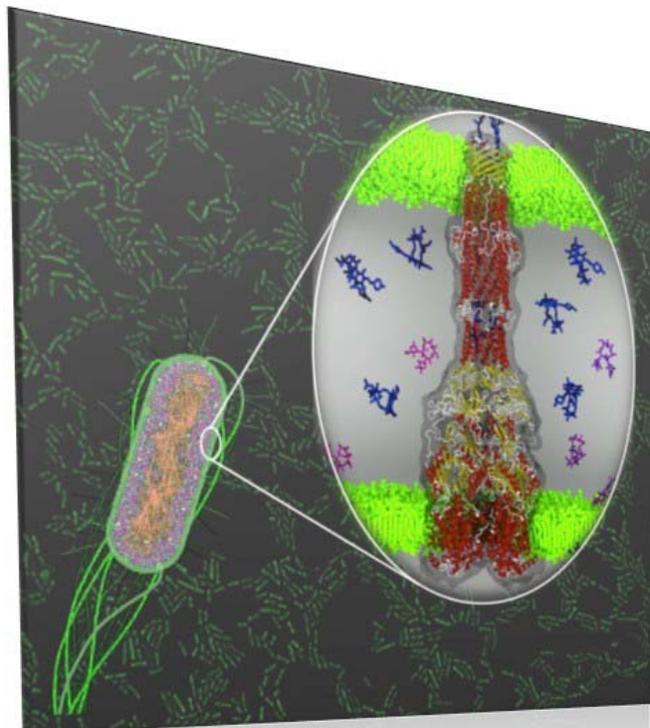
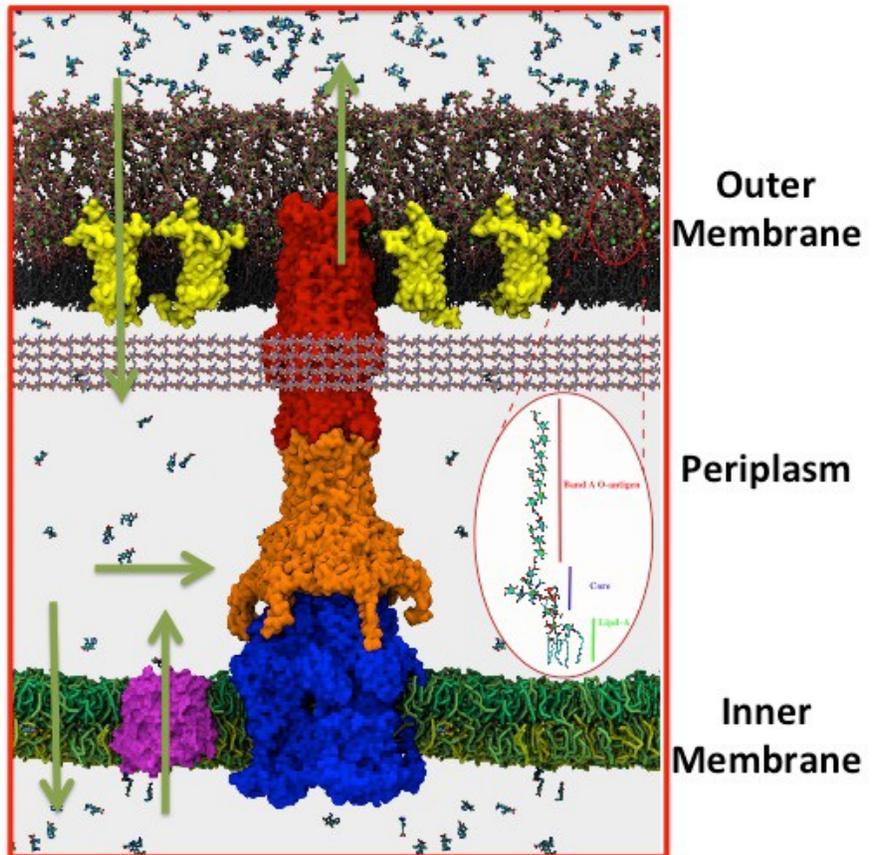
6. (Optional) Images

Molecular Model of Gram-Negative Cell Envelope

Translocation of Antibiotics

Proteins: Structure, Function, and Bioinformatics, [PMID: 24957790] (2015)

Permeability barrier of gram-negative cell envelopes and approaches to bypass it. ACS Infect. Dis. 2015. 1 (11), pp 512–522



Predicting Enzyme Function on a Genome-Scale
Matt Jacobson, University of California-San Francisco

1. Please specify the current science drivers for your field of research.

Broad impact: Genome sequencing efforts have greatly outstripped our ability to predict the functions of millions of uncharacterized gene products. Predicting the activities of uncharacterized, putative enzymes is of particular interest to define the biochemical capabilities of individual organisms and communities, and to identify enzymes of potential interest for synthetic biology applications. Finally, the same approaches can also help to meet the inverse challenge, designing enzymes with desired catalytic capabilities.

DOE interest: This work contributes very directly to the following goal excerpted from the BER mission statement: “The program seeks to understand how genomic information is translated to functional capabilities, enabling more confident redesign of microbes and plants for sustainable biofuel production, improved carbon storage, or contaminant bioremediation.” The work is also catalyzed by DOE’s investments in supporting structural biology (*vide infra*).

Overall strategy: Tackling such a challenging problem requires the synthesis of many computational methods in combination with diverse experimental data/methods, including gene sequences, protein structures, enzymology, and microbiology. The bioinformatics challenges will be discussed in other white papers. The core challenge from the standpoint of computational biophysics is predicting an enzyme’s properties — most notably, the chemical transformations it can perform — utilizing available structural information. For an overview, see M.P. Jacobson, C. Kalyanaraman, S. Zhao, and B. Tian, “Leveraging Structure for Enzyme Function Prediction: Methods, Opportunities and Challenges,” *Trends in Biochemical Sciences* 39 (2014), 363–371.

This challenge cannot be solved using any one computational method, but rather a hierarchy of methods addressing different aspects of the problem.

| Challenge | Current and emerging approaches |
|--|---|
| Predicting the structure of an enzyme from its amino acid sequence. | Comparative protein modeling is in certain respects mature, but the CASP contests have demonstrated continuing challenges which may be addressed in part by high-end computing (e.g., using extensive molecular dynamics sampling to improve models). |
| Predicting what metabolites could be substrates (or products) of the enzyme. | Metabolite docking utilizes large in silico ligand libraries and algorithms adapted from computer-aided drug design to identify plausible substrates. |
| Predicting what chemical transformation, if any, an enzyme could perform on a given hypothetical substrate. | Quantum mechanics and QM/MM (mixed quantum-classical) simulations provide the necessary framework, but are currently deployed almost exclusively to characterize mechanistic details of well-studied enzymes. |
| Integrating computational and experimental data concerning multiple enzymes in a given organism to predict biochemical pathways. | Relatively little work has been dedicated toward bridging the gap between computational studies of specific enzymes and “systems-level” modeling, e.g., genome-scale metabolic models. This will require deploying the above approaches on a genome scale, and methods for analyzing the large resulting data sets. |

2. Describe the science challenges expected to be solved in the 2020–2025 time frame using extant computing ecosystems.

Applying existing algorithms for metabolite docking on a genome (or multiple metagenome) scale is straightforward, due to the “embarrassingly parallel” nature of the calculations (i.e., each docking calculation for a protein-ligand pair can be performed on a single core). The scale of the problem can be estimated as ~10,000 virtual metabolites docked against structure or models of ~1,000 putative enzymes in a given organism, for a total of 10M individual docking calculations, each requiring a few minutes of computational time on a single core. Thus, this should be achievable.

Likewise, increasing computational power, using architectures similar to those already available, should make it possible to routinely deploy quantum mechanics in enzyme function prediction and enzyme design. Such applications might require QM calculations of transition state barriers for hundreds of possible ligands in one enzyme active site, or conversely, hundreds of possible active site designs for one chemical reaction. Computational costs vary widely based on levels of theory, etc., but for purposes of a rough estimate might be on the order of a few processor-days per enzyme-ligand complex. Thus, large-scale application becomes possible, potentially even approaching the genome scale.

Finally, computationally intensive simulations using methods like molecular dynamics or Rosetta are likely to improve the quality of comparative protein models, although the recent CASP meetings have made clear that improvements remain modest, thus far.

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

Treating entire proteins at a quantum mechanical level will likely remain intractable (with the exception of semi-empirical methods) due to the fundamental scaling of the methods (generally N^3 for commonly used methods, where N is the number of electrons) and difficulties associated with parallelizing the codes over very large numbers of processors. Likewise, treating macromolecular dynamics at the quantum mechanical level, for example, enzyme reactions fully incorporating thermal fluctuations of the protein, will remain impractical for most applications.

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 | Large GPU clusters will be transformative for molecular dynamics, for example. |

| Impede | Why? |
|--------------------------|--|
| 1. Workforce development | Across computational biophysics broadly, we need to train more people who are sophisticated about not only biological applications, but also algorithms and hardware architectures. |
| 2. Application codes | Updating (to take advantage of advances in hardware) and maintaining the complex codes required for docking, molecular dynamics, and quantum mechanics is highly challenging and difficult to support in academia. |

Exascale Virtual Cell

Leslie Loew, University of Connecticut School of Medicine,
and Phillip Colella, Lawrence Berkeley National Laboratory

Name, Institutional affiliation, and e-mail address: Leslie Loew, University of Connecticut School of Medicine, les@vlt.uchc.edu; Phillip Colella, Lawrence Berkeley National Laboratory, PColella@lbl.gov

Overview Description and Impact

Cells and the tissues that they form are composed of highly regulated dynamic chemical factories containing millions of different interacting molecular species within multiple flexible and geometrically intricate compartments. Transport of molecules through membranes separating these compartments is regulated by both chemical and electrical signals. The energy produced in biochemical reactions can also be transduced to generate mechanical force to drive alterations in cell shape, cell division, or cell migration. To understand how all these physical and chemical events are coordinated to produce the multitude of specialized cell functions is the long-term ambition of the Virtual Cell (VCell) Project.

Currently VCell (continuously funded by NIH since 1998) is able to address a large variety of specific biological problems focused on the function of cellular subsystems. However, to address and understand larger systems, and how they are integrated, is not possible with our current set of tools; this will require an exascale computing infrastructure. The scale and features of the VCell database will need to grow to enable the reuse of model components for simulation of increasingly large and complex systems. Expanding our code and tools to the exascale will allow us to tackle problems such as stem cell differentiation and organoid development, the diversity of cell-based immune response, the multiscale structure and function of neurons, the tumor micro-environment of cancer metastases, and more.

System Requirements

Code and Tools: The current VCell is a modular computational framework that is easily accessible to cell biologists and that permits construction of models, application of numerical solvers to perform simulations, and analysis of simulation results. VCell allows users to create and run simulations of biochemical networks, membrane transport, and electrophysiology. These are automatically translated using a combination of physics-based simulation frameworks: compartmental models numerically solved with a choice of ODE solvers; 3D spatial models solved as PDEs (incorporating realistic experimental geometries explicitly accounting for diffusion and flow); discrete compartmental simulations using stochastic or hybrid solvers (for reaction networks with species at low copy number); and full spatial simulations with stochastic reactions and Brownian dynamics (following trajectories of individual molecular species). Due to continuous enhancements in capabilities, and many unique features, VCell has achieved a fast-growing user base. As of November, 2015, more than 5,000 VCell users have run simulations on the system. They have collectively stored more than 70,000 models and 421,000 simulations in the VCell database system, and over 800 models were made public by their owners to be available to the worldwide VCell community.

Despite the success of VCell, the limitations in both computer hardware and mathematical algorithms keep us far from our ultimate ambition. Typical large-scale ODE problems in our database have $\sim 10^3$ variables, can be simulated in minutes, and do not pose major scaling issues. On the other hand, the largest PDE problem that has been solved with VCell (~ 60 variables, $\sim 10^5$ control volumes, 10^4 seconds time course) required two weeks to complete on our current hardware and cannot be practically solved at the desired scale even on high-end systems such as Titan. Furthermore, the equivalent mesoscale resolution stochastic simulations have an even higher computational cost. In addition, the current numerical methods are all optimized for fixed geometries and do not include formulations for cellular shape changes or mechanical forces; they also cannot account for the molecular-scale forces required to properly model polymers and high-turnover molecular clusters. To truly achieve an exascale “virtual cell,” we need not only more efficient numerical methods and higher performance hardware, but also new multiscale and multiphysics algorithms.

Models and Algorithms: Required advances in mathematical algorithms and software include new mesoscale physical descriptions for multi-molecular interactions in subcellular volumes; structured-grid adaptive mesh refinement algorithms with embedded-boundary and immersed-boundary descriptions for mechanics, deforming geometries, and electrophysiology; new computational infrastructure for large models and kinematics; and new algorithms for connecting complex dynamic microscopy imaging experiments to models. These capabilities would leverage a more general software framework, specifically the Chombo framework developed at LBNL. Specific use cases include solving reaction-transport problems in multi-compartment domains with moving boundaries; mechanical models of cytoskeletal dynamics; electrophysiology; the ability to simulate model ensembles matching probability distributions; automatic analysis of sloppy models; and enhanced rule-based modeling capabilities.

End-to-End Requirements: The ability to create models of increasing complexity will be enabled through enhancements to the VCell database, which now allows for model reuse and enhancement. The vision is that models can be modularized to contain fully annotated components that can be used as building blocks for ever more complex “supermodels.” Once this is realized, it will lead to an explosive growth in the community-driven use of VCell and in the size of the database, which in itself will require resources at the exascale. To enable effective model exploration, exchange, and reproducibility, we will need seamless bridging of VCell with external bioinformatics resources; enhanced support for standards such as BioPAX, SBML, and SED-ML; and novel domain-specific methods to visualize distributed large-scale complex datasets. Thus, the VCell database will be continuously expanding in scope and functionality. The central vision of the Exascale Virtual Cell project is that, coupled with multiscale/multiphysics modeling algorithms, the database will be the virtual cell.

Path to Exascale: Highly heterogeneous coupled multiphysics/multi-model (structured-grid/cut-cell/particle discretizations, deterministic/stochastic models) make static load balancing (even with introspection) untenable. For that reason, we will be using a hybrid of distributed-memory and threaded programming models, prototypes of which exist in Chombo, including for the embedded boundary capabilities. We will also need to interoperate with high-performance implementations of other elements of the exascale software stack, such as the PETSc solvers. Resilience through GVR or Containment Domains, using a combination of DRAM and NVRAM technologies, will also be required for local recovery and robust execution in the presence of platform degradation. Exascale platforms will also provide only a modest amount of high-bandwidth memory. Chombo is on track to reduce its computation working set to fit into 1/4 of the High-Bandwidth Memory on Cori with fully occupied cores. Used as a scratchpad, this reduced working set memory model can be translated to GPUs.

Related Research: There have been several attempts to create whole-cell models of prokaryotic cells. These models required HPC technologies and have moved the field forward significantly, even though they described the behavior of a very primitive cell. Indeed, these efforts prove that for eukaryotic cells exascale computing will be absolutely required (below); furthermore, because these efforts have been ad hoc, they do not provide a basis for enhancement and scale up. The proposed exascale VCell will be built within a database structure that will allow for the reuse of model components and thus benefit the entire systems biology modeling community, and the VCell database will also allow for ready access to external databases of experimental data (although limits on the availability of data for many cellular systems does represent an ongoing challenge). Furthermore, the mathematics, computer science, and scientific computing advances described above will be widely beneficial for computational modeling research, and will enable radical advances in computer-aided design and testing for new interventional strategies in agriculture and medicine. Examples of global impact include virtual validation of new pesticides and herbicides, in silico identification of drug targets, computational drug screening including toxicology, design of culture protocols and genetic engineering for the production of organoids to be used in regenerative medicine, simulation-based selection of personalized combination chemotherapy in cancer patients, etc.

10-Year Problem Target: Given the recent advances in understanding the governing principles of intracellular processes, coupled with high-throughput quantitative data on molecular components, we fully expect in a decade from now to see highly sophisticated attempts for detailed, predictive simulations of multiple eukaryotic cells acting in concert: development of an embryo, of an organoid, of cancer metastatic foci, etc. The computational cost for the recently achieved complete simulation of the life cycle of one of the most simple bacteria (*M. genitalium*) was ~0.1 TFlop·hr. The low-end limit for the problem target above can be estimated using a minimalistic eukaryotic organism (*C. elegans*) as an example, with known scaling factors for genome size, number of molecules, volume, cell number, and additional regulatory mechanisms/complexity; based on current technology, it is ~100 EFlop·hr for a single run. We expect that this number can be substantially reduced by the use of improved algorithms (e.g., high-order spatial discretizations, more aggressive use of adaptive mesh refinement, and linear and nonlinear solvers that enable more effective use of implicit discretization in time). However, even after such improvements, we expect that the simulations at this level of model complexity and fidelity will remain in the exascale class.

Other Considerations/Issues: Because biologists rarely have sufficient training in the mathematics and physics required to build quantitative models, modeling has been largely the purview of theoreticians who have the appropriate training but little experience in the laboratory. However theoreticians are often loath to dive into the many critical details required to fully describe a biological process. We have extensive experience in bridging these two scientific communities with the current VCell software environment, and we intend to keep the goal of serving both communities in the forefront as we build the exascale Virtual Cell.

BER White Paper: Atomistic Simulation Scaling to the Size of a Living Cell
Loukas Petridis, Oak Ridge National Laboratory

1. Please specify the current science drivers for your field of research.

Where we are today? Be sure to include broad impact, DOE interest, ties between experiment/theory, etc.

The most critical issues in 21st century biosciences involve definition of the principles linking genetic information, in the form of one-dimensional sequences of genomes, to how biological processes take place in time and space. Molecular dynamics (MD) simulation, involving stepwise integration of the equation of motion for a system of classical particles through an empirically derived potential energy function, has evolved into one of the most widely applied computational research tools in biology. MD has been employed successfully in understanding processes relevant to biofuel production, the functioning of molecular machines participating in cell signaling in microbes, and the structure of plant cell walls. MD simulations have been tightly integrated with various experimental techniques such as crystallography, nuclear magnetic resonance, small-angle scattering, and transmission electron microscopy. Of particular note is the complementarity of MD simulations with neutron scattering, two techniques that probe similar time and length scales.

2. Describe the science challenges expected to be solved in the 2020–2025 time frame using extant computing ecosystems.

What will probably be solved in the next 5–10 years? Why is this important to the field?

A critical scientific challenge concerns the complex interplay of the molecular systems within cells. More specifically, understanding how structure and dynamics give rise to function requires spatiotemporal characterization spanning decades of time and length scales. The overarching aim of this project is to employ simulation and neutron scattering techniques to obtain high-resolution spatial and temporal information on biological processes, and thus demonstrate the role that the interaction between the members of complexes plays in defining their function. The impact of this work will be on collecting intracellular data, collected over many orders of magnitude of length and time, and translating such data into a coherent understanding of coupled and interdependent biological processes.

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

What might not necessarily be solved in the next 5–10 years? Again, what is the importance?

Understanding how structure and dynamics give rise to cellular function requires spatiotemporal characterization spanning decades of time and length scales, from $\sim 10^{-12}$ to 10^{-4} s, and from $\sim 10^{-10}$ to 10^4 m. However, current biological MD applications are limited, due to hardware and software design, to $\sim 10^{-6}$ s and 10^{-7} m, three orders of magnitude smaller than those required to understand fully the physical processes leading to cellular function. The overarching challenge then is to employ HPC to extend the time and length scales accessible to the physical simulation of high-resolution molecular biological systems.

New multiscale algorithms, which couple more than one level of resolution across domains of time and space, need to be developed and ported so as to extend the timescales accessible to biological MD.

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

| Accelerate | Why? |
|-----------------------------|--|
| 1. Algorithms | A significant challenge is improving sampling of the configurational space in MD. Enhanced sampling methods need to be integrated into a “parallel ensemble” (multiple copies of loosely coupled systems) framework to further improve the configurational sampling. |
| 2. Application codes | Develop existing codes (GROMACS, NAMD, LAMMPP, etc.) to take advantage of heterogeneous architectures. The MD codes will thus take full advantage of future computing architectures. |
| 3. Hardware–network latency | Strong scaling in biological MD is primarily limited by the communication latency. Improving strong scaling will extend the timescales accessible to MD. |

| Impede | Why? |
|--------|------|
| 1. | |
| 2. | |
| 3. | |

5. (Optional) Characterize the data ecosystem aspects if the primary drivers for your field of research involve the transmission, analysis (including real-time analysis), or processing of data. For example:

- Classify the data as simulation, experimental, or something else.
Simulation.
- Characterize these data using the 3 V's (Velocity, Volume, and Variety).
Low variety data: trajectories consist of 3D coordinates of all atoms in the system.
How much data will be generated/stored with the next 5–10 years?
20 TB.
What will the data rates be?
About 200 gigabytes/hour.
- Describe current or planned solutions for any of these data challenges, e.g., in situ analysis.
Parallel analysis.
- Note any particular data security or privacy requirements.
N/A

6. References (please keep the reference length to no more than 10 critical references)

Schulz, R., Lindner, B., Petridis, L., and Smith, J., 2009, "Scaling of Multimillion-Atom Biological Molecular Dynamics Simulation on a Petascale Supercomputer," *J. Chem. Theory Comput.* 5(10): 2798–2808.

Lindner, B., Petridis, L., Schulz, R., and Smith, J.C., 2013, "Solvent-Driven Preferential Association of Lignin with Regions of Crystalline Cellulose in Molecular Dynamics Simulation," *Biomacromolecules* 14(10):3390–3398.

Vermaas, J., Petridis, L., Xi, Q., Schulz, R., Lindner, B., and Smith, J.C., 2015, "Mechanism of Lignin Inhibition of Enzymatic Biomass Deconstruction," *Biotechnol. Biofuel.*, accepted.

7. (Optional) Images

Model of plant cell wall; cellulose fibers are shown in dark green, lignin in brown, and hemicellulose in light green (Image credit: scistyle.com).

Roland Schulz, Intel Corporation

Current and Future Parallel Machines

The current largest systems are clusters of thousands of nodes with symmetric multiprocessors (SMPs) using specialized interconnects. Smaller HPC systems are typical clusters of SMPs with commodity interconnect. Even current workstations and notebooks are SMP nodes with many cores. The parallel processing available on a single node will keep increasing and will be the primary source of the expected increase in performance over the next 5–10 years.

Future Molecular Dynamics Simulations

Future larger machines will allow simulations of billions of atoms. They will also allow the creation of ensembles of thousands of independent simulations that can be used by enhanced sampling methods to compute statistical properties for timescales not directly accessible. The decrease in cost per simulation will make it increasingly economical to use simulations to replace or augment experiments in all areas of biophysical research. The increasingly powerful nodes will make small simulations (such as a single protein) run fastest on a single node. Without the bottleneck of communicating between nodes and with enough compute performance within the single node, small simulations will run faster than is currently possible on either a single node or a cluster.

Because the extra performance will come in the form of extra parallelism, obtaining the extra performance in MD will not be automatic. It will require sufficient funding in at least one MD community code, with extra parallelism and more optimized multi-level parallelism (Node, core, SIMD/SIMT) and data movement, compared to the best scaling code of today.

By itself, the extra compute performance provided by future hardware will not make longer simulations possible, especially those larger than can run efficiently on a single node. Making simulation of processes on longer timescale simulations possible will require improvements in strong scaling. Better scaling electrostatic solvers (multigrid or fast multipole method) will potentially improve performance for very large biophysical simulations, and thus enable very large simulations. However, they will probably have little impact on small- to medium-sized simulations such as those of single proteins or small assemblies with less than one million atoms. To achieve significantly improved strong scaling, the communication time for small messages needs to be significantly reduced. If current hardware trends continue, only modest improvements in small message latency and achievable timescale can be expected. Extra investment and prioritization of this hardware requirement would be needed. Special-purpose computers have demonstrated that an order of magnitude lower latency is possible with a tight integration between network and compute, and fine-grained event-driven computation. Field programmable gate arrays (FPGAs) allow tighter integration with the networking; general-purpose computer hardware and event-driven computation could be added as part of co-design. For small simulations running on a single node, synchronization and data movement latencies will be limiting factors for strong scaling. To reach simulation speeds of 1 μ s/day, low-latency computing cores and fine-grained synchronization will be required.

Like any HPC application, MD has large data and visualization requirements. But particle methods have modest requirements compared to volumetric methods. Thus, other applications will drive the requirements for future hardware, and MD will not be limited by IO or visualization requirements suitable for those.

From Molecular to Reach and Watershed Scales in Environmental Modeling

Jeremy Smith, Oak Ridge National Laboratory

While the cycling of contaminants and other materials is critical to DOE science, the requisite understanding is lacking, due largely to our inability to accurately model and predict biogeochemical transport and transformations in terrestrial surface and subsurface systems. Cycling on the reach and watershed scales depends on molecular (atomistic) scale biogeochemistry. Chemical changes are central to cycling. Although much recent progress has been made, quantitative characterization of chemical processes in complex environmental systems is needed. Each reaction can be characterized by a thermodynamic binding free energy (“log K ”) and a reaction rate, the latter depending on the activation free energy. While useful thermodynamic databases exist, substantial uncertainties remain in mechanisms, log K s, and kinetics of individual biotic and abiotic processes transforming carbon and contaminant materials. There remains substantial uncertainty in thermodynamic parameters required.

In addition to thermodynamic association constants, critical transformation rates either have not been quantified or have been described by simple, empirical parameters. Using Hg as an example, among the rates that need characterization are those for various oxidation and reduction, methylation and demethylation, precipitation and dissolution, and complexation reactions.

In most relevant meso-scale models the gas, solution, and solid phases are assumed to be well-mixed continua in each representative volume; equilibrium chemical reactions are represented by mass action equations; kinetic reactions are described by ordinary differential equations; and transport is simulated by the advection-diffusion equation. However, the deficiencies outlined above prevent us from integrating atomistic-scale data into these models to obtain a multi-scale model of cycling of use in environmental prediction. Consequently, even though geochemical codes have been under development for several decades and are widely used to simulate biogeochemical processes from lab to field scales, their application has been extremely limited. Obtaining the data required for accurate continuum-scale modeling and assembling them in a simulation framework are required for us to obtain a predictive understanding of cycling across spatial and temporal scales.

Although it has been increasingly recognized that linking atomistic modeling to these macroscopic application scales has the potential to improve overall model accuracy and predictability, atomistic modeling techniques rarely have been utilized to improve continuum-scale model predictions. To make this connection, atomistic computational chemistry toolkits need to accurately calculate interactions and to apply the results in meso-scale models for lab-to field-scale applications, providing first-principles input to reach- and watershed-scale modeling.

BER White Paper
D.J. Tantillo, UC-Davis

1. Please specify the current science drivers for your field of research.

The multi-scale modeling of biosynthetic reactions involving enzymes with nonpolar active sites is a challenging endeavor. Such enzymes play important roles in the biosynthesis of, for example, terpenes. Terpenes are hydrocarbons that have received considerable interest as fuels. Today we can model the inherent substrate chemistry (including dynamic effects), but modeling the chemistry on-enzymes is still a challenge and the area where high-end computing is essential. The main challenge arises in large part from the fact that both the substrate and surrounding active site lack polar groups that can participate in predictable hydrogen bonds, etc. Also, crystal structures of the enzymes involved often are not in productive conformations.

2. Describe the science challenges expected to be solved in the 2020–2025 time frame using extant computing ecosystems.

We are working on methods to address this problem, including specialized docking and homology modeling protocols that will facilitate subsequent dynamics simulations on relevant chemical reactions. Ideally, with these developments we will be able to predict changes to product distributions upon mutations. All of this will be facilitated by improved computing resources. In particular, computational resources are required that will allow for large quantum mechanical (QM) regions to be included in QM/molecular mechanics (MM) simulations on enzyme dynamics and reactivity. The substrates for these enzymes are large and weak, but important interactions (e.g., dispersion interactions, C–H••• π interactions) between them and surrounding enzyme residues are generally not well-treated by force fields.

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

We'd also like to predict oxidation patterns for downstream P450 reactions, thereby enabling more rational pathway engineering.

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

| Accelerate | Why? |
|---|--------------------------|
| 1. Hardware resources (at all scales) including I/O, memory, etc. | Need high throughput |
| 2. Models and algorithms | Integration of QM and MM |
| 3. Application codes (implementation, development, portability, etc.) | Faster QM |

| Impede | Why? |
|---|------|
| 1. Hardware resources (at all scales) including I/O, memory, etc. | |

Molecular Based Modeling of Biological Systems

Marat Valiev, Niri Govind, Edoardo Aprá, and Karol Kowalski Environmental Molecular Science Laboratory

1. Current science drivers for your field of research

The fate of biological systems is critically dependent upon the intricate details of numerous chemical transformations at the microscopic (angstrom) level. Many questions related to this process remain largely unknown, ranging from basic characterizations of chemical reactions in enzymatic active sites to more complex problems involving long-range allosteric transformation or electron transport. The immense dimensions of this problem make it an ideal candidate for massively parallel computing platforms, but the necessary software tools are still in their infancy. Our efforts are aimed at closing this gap through the development of computational, *in-silico* modeling approaches that can provide molecular-derived descriptions of biological processes.

Our developments take place within the framework of NWChem, an advanced computational chemistry code funded by BER. Efficient parallel implementations of a wide range of quantum-mechanical (QM) models and native QM molecular-mechanics (QM/MM) capabilities make NWChem perfectly suited for this task. The open source code and modular nature facilitate community-based development.

2. Science challenges expected to be solved in the 2020-2025 time frame using extant computing ecosystems

The key feature of a chemical process in a biological system is the presence of two disparate scales: electronic structure effects at the molecular level embedded into a long-range, nanometer-level DNA/protein structure. Computational modeling of such a system is best approached through multi-physics type methods, where different physical models are used to capture natural scale separation. QM/MM is one particular example of such an approach. The system is separated into two regions—the reactive region described at the QM level and the environment treated by means of classical molecular mechanics. This simple model captures a large class of important biological applications, including structural analysis, electronic structure properties, and free energies. Structural analysis may include accurate characterization of enzymatic active sites (reactant, transition, and product states), which are not directly accessible in the experiment. These types of calculations could also be used to estimate the structural impact of active site point mutations, providing initial clues to the potential impact on enzymatic activity. Electronic structure property calculations provide accurate spectroscopic signatures (e.g., NMR, UV-vis), connecting directly to experimental measurements. Free energy calculations could allow calculation of activation barriers and reaction energies.

While many of these capabilities have already been demonstrated, significant computational requirements and the complex nature of these simulations presented significant obstacles toward their widespread adoption in the general BER research community. The high cost of such calculations is typically related to the computational cost of QM calculations, especially those involving high levels of theory (e.g., CCSD(T)). Continuing growth in parallel computing resources and ongoing work in improving the efficiency of QM models should alleviate these issues and can be realistically expected in the coming years. A second limiting factor is the increased level of complexity of QM/MM calculations. These include system setup (e.g., force field parameters, topology), choice of QM region, proper parameters for the QM model, and appropriate computational resources (e.g., number nodes, simulation time). Many of these issues are currently being addressed within the framework of NWChem by abstracting input files into less complicated work flows, as well as developing interfaces to external MM codes, such as AMBER.

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

While QM/MM methods provide a good starting point for the analysis of reactive processes in biological systems, their applicability is ultimately restricted by the presence of a fixed rigid boundary between QM and MM regions. As a result, problems that involve large dynamical rearrangements of protein backbone or long-range electron transport will require different strategies. A brute force approach to this problem could consist of enlarging the QM region. The size of the latter may easily exceed thousands of atoms and will require substantial parallel computing resources for extended periods of time. A possible solution to this problem could be offered by lower-order approaches like DFT tight-binding (DFTB), but will require extensive benchmarking and parameter development. Another direction may involve using so-called QM/QM approaches, where multiple and potentially overlapping QM regions are used to accommodate dynamical changes in the system. The main challenge with these approaches is that the Hamiltonian of the system is no longer static but changes dynamically during the evolution of the system. In either case the success of any approach (defined as being available to the general user community by 2025) that will go beyond standard QM/MM techniques will be critically dependent upon sustained and focused multi-year development effort involving substantial commitment of computing resources.

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

Efficient QM methods capabilities: (1) Availability of new algorithms capable of utilizing the sparsity of first-principle formulations to push the envelope for system-size tractable by high-order coupled cluster (CC) methods. While these algorithms are available for ground-state formulations, there is no clear path forward for excited-state and linear-response formulations. (2) Development of novel tensor libraries that are needed to incorporate new programming models to deal with the increase in intra-node parallelism; data localization and reduction in intra-node communication; efficient utilization of runtime systems; development of topology-aware algorithms; and utilization of deeper memory hierarchies. (3) Hybrid DFT/CC description based on mathematically rigorous integration of various representations of quantum mechanics to build efficient embedding schemes based on the linear response coupled-cluster methods. (4) Development of lower order approaches like DFTB for extended length- and time-scale atomistic simulations on the ground and excited states. The availability of more complex quantum mechanical approaches in NWChem, such as DFT, MP2, and CC theory within the same program, positions us uniquely to test, validate, and develop new DFTB parameter sets that would be applicable to broader classes of systems and allow us to propagate accurate high-level corrections across length scales.

Asynchronous QM/MM algorithms: Conventional QM/MM approaches are based on strictly synchronous execution models between the QM and MM components. This leads to significant workload imbalance and idle processors, resulting in poor utilization of parallel computing resources. A key part of our project is the development of new asynchronous computational algorithms that can facilitate balanced partitioning of parallel computational resources between the two scales. This partitioning will account for the compute-intensive nature of the quantum-mechanical component (NWChem) and the latency-sensitive nature of the classical component.

Work flows and automation: Existing ecosystems at supercomputing centers are aimed at providing efficient parallel queuing systems but offer minimal support for ensuring proper simulation procedures and required computational resources. This creates an enormous entry barrier for a user whose area of specialization lies outside the code development. The fact is that many QM/MM simulations revolve around a handful of computational protocols and, in most cases, they can be reused with minor changes across many different systems. Currently, there is no easy mechanism which allows us to capture existing simulation procedures (e.g., in the form of workflows), distribute them to users, and provide a means of submission with an automated choice of proper computing resources.

Microbes to the Environment (corresponds to Section 3.1.3)**Reaction and Transport Parameters from the Electronic- to Nano-scale for the Multi-scale Modeling of Natural Organic Matter Stabilization/Decomposition Processes in Terrestrial, Aquatic, and Aerosol Systems**

Amity Andersen, Ryan Renslow, Marat Valiev, Niri Govind, and Eric Bylaska
Pacific Northwest National Laboratory, Richland, Washington

1. Current science and operational drivers

Natural organic matter (NOM) is found in all of the Earth's major carbon pools, and its stabilization and decomposition play a significant role in the Earth's carbon cycle. In the atmosphere, NOM is found in aerosol particles and plays a role in aerosol particle formation and chemistry.¹⁻³ In aquatic systems, dissolved organic matter (DOM) can be found, and its decomposition can lead to a net increase of CO₂ in the atmosphere. For the terrestrial carbon pool, soil is the primary organic carbon contributor (soil organic matter, SOM). In addition to soil's impact on climate change, prudent management of soil resources is essential to food, fiber, water, and ecosystem security.

Simulation and modeling protocols from quantum mechanics to classical atomistic and coarse-grained simulation describing the electronic to nanoscale regime of complex Earth systems are proposed to provide parameterization pertinent to reactive transport in micro- to global-scale modeling. Of specific interest here are the stabilization and decomposition processes of NOM, which directly impacts the Earth's carbon cycle and therefore the trajectory of climate change. Another aspect of NOM is its largely unknown supramolecular structure with intra- and intermolecular interactions in self-assemblies, including inorganic ionic species (e.g., metal cations Fe^{2+/3+}, Mn^{2+/3+/4+}, Cu⁺²⁺, Ca²⁺, Al³⁺, and anions such as phosphate) and interactions with mineral surfaces (e.g., clays, metal oxide/hydroxides). Structure characterization of NOM through computational spectroscopy (Mössbauer, NMR, EPR, and XAS) is also an important part of developing models for further computation of reactive transport parameters. *The aggregation of NOM in supramolecular self-assembled structures and on mineral surfaces is hypothesized to protect NOM, including biogenic organic matter (e.g., microbial exo-enzymes), from decomposition from biotic processes such as microbial metabolic consumption and abiotic processes.*¹⁻³ However, transition metal-containing mineral surfaces may facilitate abiotic decomposition of NOM through catalysis and make available decomposition products for microbial metabolic processes.

The proposed research is in line with BER's mission to advance "understanding of the roles of Earth's biogeochemical systems (the atmosphere, land, oceans, sea ice, and subsurface) in determining climate so we can predict climate decades or centuries into the future, information needed to plan for future energy and resource needs."⁴ The proposed research falls under EMSL's Terrestrial and Subsurface Ecosystem Science Theme in providing a fundamental (atomistic) understanding of the biogeochemical processes in soils that can provide qualitative and quantitative information for model development efforts at longer length and time scales.⁵ The proposed research also addresses the expanded focus areas of advanced spectroscopic capabilities to characterize NOM by providing computational tools to aid in interpretation of complicated multi-component spectra. NOM is also present in atmospheric aerosol particles; therefore, the proposed research also applies to the "Atmospheric Aerosol Systems" science theme.⁶

2. Broad computational and data challenges expected to be faced in the 2020–2025 time frame

With the myriad of chemical components, ways of organizing with other organics and with mineral surfaces, and reaction pathways for NOM, large databases of elementary biogeochemical reaction rate constants and transport parameters are required for multi-scale modeling involving multi-phase reactive flow. These parameters can be calculated from first principles using quantum mechanics (electronic scale) and classical molecular mechanics methods (atomistic to nanoscale). The former of these methods allows for the calculation of elementary rate processes involving chemical change (e.g., isomerization, polymerization, chemisorption, metal-ligand complexation, decomposition to CO₂ and other products, photochemistry) and some physical change (e.g., physisorption). The latter of these methods can be used to calculate properties related to transport, such as self-diffusion in the aqueous phase and adsorption to other organic molecules and to mineral surfaces.⁷ Classical molecular methods can also provide an ensemble of molecular complex configurations to be studied at finer electronic detail with quantum mechanical calculations. For organic systems with unknown supramolecular structure such as NOM, molecular structure characterization is an integral part of molecular model building for later calculations of reaction and transport parameters. Spectroscopic methods such as Mössbauer, nuclear magnetic resonance (NMR), electron paramagnetic resonance (EPR), and x-ray absorption spectroscopies (XAS) are employed for structure characterization of complex materials such as NOM. Quantum mechanical computational techniques can then be employed in an interpretation capacity to elucidate complex spectroscopic data.

For large supramolecular structures such as those found in NOM, combining quantum mechanical and molecular mechanics (QM/MM) methods will be desirable for accurate reaction rate constants, which depend not only on the local center of reaction but also on the degrees of freedom of the supramolecular system as a whole. In QM/MM simulations, a small region of the large molecular system (e.g., a catalytic metal center) is treated with quantum mechanical methods, whereas the rest of the large system is treated with the less computationally demanding molecular mechanics methods. Current systems of interest include NOM carboxylate side chains (e.g., from aspartate and glutamate peptide residues and fatty acids) complexed, directly or through water-bridges, to metal cations (especially Fe^{2+/3+}, Mn^{2+/3+/4+}, Ca²⁺, and Al³⁺) readily available in soil systems. These carboxylate-metal ligand complexes can create organic-organic molecular linkages and centers for catalytic reaction. Quantum mechanical and QM/MM calculations will also be used to simulate spectroscopic features^{8,9} of NOM systems and will be compared with available data. These calculations will be performed with the BER-sponsored quantum chemistry suite NWChem. AMBER (interfaced to NWChem for QM/MM calculations) and the BER-sponsored LAMMPS software packages will be employed for large-scale classical molecular mechanics and dynamics.

3. Broad computational and data challenges that cannot be solved in the 2020–2025 time frame using extant computing ecosystems

The natural environment contains a currently unknown, potentially daunting number of chemical species. To calculate a multitude of potential reactions and transport parameters for multi-scale modeling quickly, the CPU, memory, and I/O hardware technology must continue to become faster. With the continued evolution of hardware resources, the software applications must adapt to take advantage of improvements in hardware technologies. Moreover, robust, seamless, and efficient database population, storage, and access software must be developed.

| Accelerate or Impede | Why? |
|---|---|
| 1. Hardware resources (at all scales) including I/O, memory, etc. | Hardware systems need to continue to become faster in order to calculate large numbers of parameters in a reasonable time frame. |
| 2. Application codes (implementation, development, portability, etc.) | Application codes must continue to evolve to take advantage of improvements in hardware technology. |
| 3. Data work flow (including sharing, transmitting, archiving, etc.) | With the generation of large data sets, the work flow for analysis and utilization at the interface of multi-scale codes must become more efficient and seamless. |

References

- Schmidt, M. W. I.; Torn, M. S.; Abiven, S.; Dittmar, T.; Guggenberger, G.; Janssens, I. A.; Kleber, M.; Kogel-Knabner, I.; Lehmann, J.; Manning, D. A. C.; Nannipieri, P.; Rasse, D. P.; Weiner, S.; Trumbore, S. E. *Nature* **2011**, *478*, 49–56.
- Stockmann, U.; Adams, M. A.; Crawford, J. W.; Field, D. J.; Henakaarchchi, N.; Jenkins, M.; Minasny, B.; McBratney, A. B.; de Remy de Courcelles, V.; Singh, K.; Wheeler, I.; Abbott, L.; Angers, D. A.; Baldock, J.; Bird, M.; Brookes, P. C.; Chenu, C.; Jastrow, J. D.; Lal, R.; Lehmann, J.; O'Donnell, A. G.; Parton, W. J.; Whitehead, D.; Zimmerman, M. *Agriculture, Ecosystems, and Environment* **2013**, *164*, 80–99.
- von Lutzow, M.; Kogel-Knabner, I.; Ekschmitt, K.; Matzner, E.; Guggenberger, G.; Marschner, B.; Flessa, H. *Eur. J. Soil Sci.* **2006**, *57*, 426–445.
- <http://science.energy.gov/bet/>.
- <https://www.emsl.pnl.gov/emslweb/science/terrestrial>.
- <http://www.emsl.pnnl.gov/emslweb/science/atmospheric>.
- A. Andersen, P. Reardon, S. Chacon, N. Qafoku, N. Washton, M. Kleber, “Protein-mineral interactions: molecular dynamics simulations capture importance of variations in mineral surface composition and structure,” *to be submitted to Langmuir*.
- Fulton, J. L.; Govind, N.; Huthwelker, T.; Bylaska, E. J.; Vjunov, A.; Pin, S.; Smurthwaite, T. D., “Electronic and Chemical State of Aluminum from the Single- (K) and Double-Electron Excitation (KL_{II&III}, KL_I) X-ray Absorption Near-Edge Spectra of α -Alumina, Sodium Aluminate, Aqueous Al³⁺·(H₂O)₆, and Aqueous Al(OH)₄⁻,” *J. Phys. Chem. B*, **2015**, *119* (26), 8380–8388.
- Lipton, A. S.; Heck, R. W.; Staeheli, G. R.; Valiev, M.; De Jong, W. A.; and Ellis, P. D., *J. Am. Chem. Soc.*, **2008**, *130*, 6224–6230.

BER White Paper
On the State of Scalable Algorithms and Software for Metagenomics Research
Ananth Kalyanaraman
Associate Professor, School of EECS, Washington State University, Pullman, WA 99164
(email: ananth@eecs.wsu.edu)

1. Please specify the current science drivers for your field of research.

Metagenomics sampling and sequencing are at the forefront of environmental microbial community analysis. DNA sequencing, using next-generation technologies (short-read sequencing) and third-generation technologies (e.g., single molecule sequencing), is continuing to deliver raw data at the rate of Moore's law. Species taxonomical binning and functional characterization of communities continue to be the two most predominant use-cases of metagenomic sequencing.

2. Describe the science challenges expected to be solved in the 2020–2025 time frame using extant computing ecosystems.

However, the emphasis is gradually shifting to a number of other advanced analytics-driven scientific lines of inquiry. Notable directions include (1) comparative metagenomics and (2) longitudinal studies of a single community.

Comparative metagenomics enables the discovery of structural and/or functional features that distinguish any two community samples from similar habitats (e.g., soil metagenomes taken from similar agricultural lands) or from different habitats (e.g., human gut vs. cow rumen). Analytical methods to compare multiple (>2) communities are still required. While it is relatively easier to study the taxonomical distribution among multiple communities, functional annotation methods with a capacity to identify novel pathways unique to different communities, or shared among multiple communities, remain a significant challenge. Visualization methods are also generally limited to pie charts and dendrograms, while the high dimensionality of feature space, in comparison, calls for more sophisticated visual analytic methods.

Longitudinal studies are becoming a significant source of valuable information to monitor structural and functional changes to a given community over a course of time and/or changing environmental conditions. Such studies are necessary to observe actionable insights. However, such methods require deep sequencing across a wide timescale. Current methodologies to analyze such longitudinal data basically extend the application of single community algorithms. Temporal methods with alignment models to track changes need to be developed. Furthermore, graph representations and related analytical methods to model the relationships between samples of different time slices and to mine for interesting patterns also need to be developed. Scaling such approaches on deep sequencing data sets presents a significant challenge.

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

With the advent of single molecule sequencing, the possibility of sequencing individual genomes from within a metagenomic community is poised to become a reality in the next 5–10 years. Current technologies (e.g., PacBio SMRT) are already generating DNA reads with an average well over 10 Kbp, but with relatively large error rates (15% or higher) compared to short-read technologies. As these technologies continue to increase in throughput, prolong read lengths, reduce error rates, and decrease costs, sequencing full-length microbial genomes from metagenomic communities is a distinct possibility. However, algorithms and technologies that can operate at the scale of an entire

population (community or tissue) need to be developed around these new sequencing technologies in order to make the latter more effective for scientific discovery. Such technologies are likely to play a role in helping us transition from a dawning era in personalized medicine to community/population-scale healthcare.

A second transformation is likely to arise in mobile, real-time processing in the longer term. Current models, where data are generated through strenuous wetlab library construction protocols followed by specialized lab capabilities to sequence and characterize a metagenome, are likely to be superseded with mobile-enabled real-time processing capabilities using sophisticated sensors at the hand of every individual.

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

| Accelerate | Why? |
|--------------------------|--|
| 1. Application codes | Increased availability of open source, interoperable software libraries, and parallel codes for different platforms for community use. |
| 2. Models and algorithms | Novel graph/network representations, easy-to-program parallel programming models, and algorithmic paradigms are likely to accelerate code development and subsequent application. |
| 3. Hardware resources | Emerging on-chip architectures and vectorized architectures, coupled with new memory technologies and interconnection networks, are likely to improve capacity to compute at exascale. |

| Impede | Why? |
|---------------------------------|---|
| 1. Visualization | Visualization of large complex data sets, their network representations, alongside their high dimensional attributes, poses a significant stumbling block toward discovery. Novel data representations and visual analytical methods for scalable hypothesis generation are needed. |
| 2. Data movement | Graph models and other data-intensive computations leave a highly irregular footprint. Current architectures are not well suited for such computations. Data movement and preserving locality are likely to be significant challenges. |
| 3. Parallel algorithms to scale | Novel parallelization design schemes that can help expose and exploit parallelism on heterogeneous computing architectures at massive scales are needed to scale to exascale. Approximation and streaming methods are likely to prove resourceful. |
| 4. Data heterogeneity | Integrating data sets from various data sources, both experimentally obtained (at varying scales, resolutions, and conditions) and synthetically generated (e.g., publication corpus, social network), presents a significant challenge in analysis. |
| 5. Interdisciplinary training | Training the next-generation scientific workforce with inherently interdisciplinary skill sets will be key. |

Multiscale Simulation of Complex Biogeochemical Earth Systems

Timothy D. Scheibe, Pacific Northwest National Laboratory, Richland, Washington

1. Current science and operational drivers

Predictive simulation of complex earth systems (e.g., mass transport of water, other fluids, and constituent components through terrestrial and subsurface environments coupled to biogeochemical processes mediated by plants, microorganisms, fungi, and animals) is a scientific grand challenge that is essential to effective future management of increasingly scarce water, energy, and other natural resources. Predictive (as opposed to empirical, descriptive) simulation requires increased fidelity of simulations to properly capture effects of mechanistic processes at small physical and temporal scales relative to the scales at which impacts are observed and controlled. This is an extreme-scale computing challenge, as the scales of interest range from those of microorganisms experiencing local conditions within a single soil pore ($\sim 10^{-6}$ m) to those of global circulation models (grid resolution $\sim 10^5$ m).

Spanning more than 10 orders of magnitude of spatial scale (and a comparable range in time scales of interest) while maintaining the fidelity necessary to predict impacts of complex coupled and nonlinear processes at the smallest scales is a problem that is naturally suited to computing at the exascale and beyond. In the past, quantitative incorporation of small-scale information in application models has been limited to upscaling methods that involve assumptions that are often highly restrictive when applied to complex natural systems (e.g., Battiato et al. 2009). Only in the past decade, during which petascale computing resources have become available, has serious consideration begun to be given to fully multi-physics, multi-spatial model integration (see Keyes et al. 2013; Scheibe et al. 2015a; U.S. DOE 2015). Multi-scale, multi-physics model integration can be posed naturally in a many-task simulation framework (Scheibe et al. 2014), which exhibits high degrees of concurrency needed to effectively utilize current and future HPC architectures built around many-core paradigms. Parallelism is exposed at several levels: (1) individual at-scale models can be coded in parallel, typically using a domain decomposition approach; (2) small-scale models executed iteratively over subdomains of a large-scale model in a hybrid approach can be decoupled and run in parallel; and (3) many alternative scenarios can be executed in parallel in Monte Carlo fashion to support sensitivity analyses, data assimilation, and uncertainty quantification.

The current state of the art is such that pore-scale simulations of fluid flow and reactive transport can be performed over only relatively small model domains (on the order of cubic decimeters at the largest; Scheibe et al. 2015b). Such simulations currently require days to months to complete; in 10 years, one might envision performing simulations on this scale (which approaches that of the highest resolution of continuum models) in seconds to minutes and coupling many such simulations directly to larger-scale application models. Incorporation of advanced models of microbial and plant function (including genome-informed metabolic networks and regulatory processes as well as feedbacks to the physical environment) is also on the current cutting edge and will add additional computational complexity to future simulations. Current microbial simulations are limited to a small number of interacting species or functional guilds; increased computational power would facilitate more realistic simulation of complex microbial and plant communities based on metagenomic and other meta-omic data.

As pointed out in a recent workshop report (U.S. DOE 2015), development of these coupled model systems will require a highly flexible, modular, and interoperable set of high-performance codes constructed within a scientific community framework—a “software ecosystem.” Such a system would comprise many subcodes interconnected in run time, driven by an adaptive controller to define when and where to execute small-scale submodels within an overall simulation domain. Critical needs for such a modeling system include a well-developed high-performance work-flow management system, and improved methods for passing information between simulators beyond traditional message passing. For example, more widespread use of in-memory files rather than traditional file-based I/O would enhance the efficiency of information exchange among simulation modules operating at different spatial or temporal scales. Development of multi-spatial/multi-physics modeling languages is currently in its infancy (e.g., Falcone et al. 2010) and must be advanced significantly to facilitate the development of general toolkits for multi-scale simulation.

In conjunction with increased computational demands for high-fidelity simulations will come increased needs for high-resolution data to support physical realism. New methods are needed for generating, storing, and utilizing detailed descriptions of pore geometry and spatial distributions of minerals, biomass, and other reactants and reaction products. These will include advanced ultra-scale visualization techniques, able to visualize and animate complex 3D flow patterns and spatial material distributions, as well as efficient methods for data input/output and/or data persistence in models that will execute many times during the course of a single complex simulation.

Quantifying, identifying causes of, and adapting to the impacts of current and future climate change, addressing increased resource limitations (water, energy, and food being the most pressing), and ensuring the protection of critical ecosystems and environmental resources for future generations are all critical problems of national scope and importance that depend on complex natural earth systems. Their solution will be greatly facilitated in the future by a mechanistically based predictive capability that properly accounts for inherent system uncertainties and adequately represents the impacts of small-scale processes on large-scale phenomena. Such a capability will enable a rational basis for decision making and design of optimal control systems, but can only be developed in the context of continued major advances in high-performance computing technologies and software environments, toward exascale and beyond.

2. Current and future computational and data strategies

Models

Physically based models (mostly in the form of numerical solutions of coupled partial differential equations) of fluid flow, mass transport, and biogeochemical reactions (including plant dynamics and interactions with soil microbiomes).

New capabilities

Model coupling is a critical aspect. This can be in the form of interfaces and data mediators for physical process modules and/or multi-scale couplings of models of the same process but different fidelity.

Resolution

Many processes of interest may require resolution at the pore scale (tens of microns) for full process fidelity. Since this is not feasible over significant spatial domains, adaptive resolution (both in terms of numerical grid resolution and adaptive physical representations) will be needed. Adaptive model coupling and integrated data exchange across scales will be essential.

I/O

This is not likely to increase as rapidly as the computational demands, as many of the highest-fidelity results are used not as end products but to inform coarser models in an iterative manner. Limited output of high-fidelity results will be required, as well as increased input to drive high-fidelity models, but not everywhere in space and time.

Many-core and/or GPU readiness

Our codes mostly are not yet prepared for this. However, the task-based concept may be readily adaptable given adequate expertise. We are currently working on adapting some codes to utilize accelerators (e.g., performing independent reaction calculations on GPUs/accelerators and more strongly coupled flow calculations on primary processors).

Work flows

We currently utilize the SWIFT work-flow system in performing multi-scale model coupling (see details in Scheibe et al. 2014).

3. What top three computing ecosystem aspects will accelerate or impede progress to meeting the challenge described above in the next 5–10 years? Why?

| Accelerate | Why? |
|--------------------------|--|
| 1. Models and algorithms | Facilitate computational efficiency through adaptive model coupling, many-task parallelism |
| 2. Data work flow | Enable reproducibility, re-use, and generalization of complex modeling tasks |
| 3. Libraries/frameworks | Facilitate re-use and software productivity, good software engineering |

| Impede | Why? |
|--------------------------|---|
| 1. Application codes | Portability, good software engineering, sustainability of legacy codes; poor understanding of software engineering by domain scientists |
| 2. Workforce development | Limited expertise in HPC among domain scientists; challenge of interdisciplinary communications |

HPC services

Better workflow management environments; model coupling and data mediator algorithms and libraries; collaboration tools if they truly facilitate collaboration between computational scientists and domain scientists.

References

- Battiato, I., D. M. Tartakovsky, A. M. Tartakovsky, and T. Scheibe, "On breakdown of macroscopic models of mixing-controlled heterogeneous reactions in porous media," *Advances in Water Resources*, 32(11):1663–1673, 2009.
- Falcone, J.-L., B. Chopard, and A. Hoekstra, "MML: Towards a multiscale modeling language," *Procedia Computer Science* 1(1): 819–826, 2010.
- Keyes, D. E., L. C. McInnes, C. Woodward, W. Gropp, E. Myra, M. Pernice, J. Bell, J. Brown, A. Clo, J. Connors, E. Constantinescu, D. Estep, K. Evans, C. Farhat, A. Hakim, G. Hammond, G. Hansen, J. Hill, T. Isaac, X. Jiao, K. Jordan, D. Kaushik, E. Kaxiras, A. Koniges, K. Lee, A. Lott, Q. Lu, J. Magerlein, R. Maxwell, M. McCourt, M. Mehl, R. Pawlowski, A. P. Randles, D. Reynolds, B. Riviere, U. Rude, T. Scheibe, J. Shadid, B. Sheehan, M. Shephard, A. Siegel, B. Smith, X. Tang, C. Wilson, and B. Wohlmuth, "Multiphysics simulations: Challenges and Opportunities," *International Journal of High Performance Computing Applications*, 27(1):4–83, DOI:10.1177/1094342012468181, 2013.
- Scheibe, T. D., X. Yang, K. Schuchardt, K. Agarwal, J. Chase, B. Palmer, and A. Tartakovsky, "A many-task parallel approach for multiscale simulations of subsurface flow and reactive transport," in *Proceedings 7th ACM Workshop on MTAGS*, New Orleans, LA, USA, November, 2014.

- Scheibe, T. D., E. M. Murphy, X. Chen, A. K. Rice, K. C. Carroll, B. J. Palmer, A. M. Tartakovsky, I. Battiato, and B. D. Wood, "An analysis platform for multiscale hydrogeologic modeling with emphasis on hybrid multiscale methods," *Groundwater*, 53(1): 38–56, DOI: 10.1111/gwat.12179, 2015a.
- Scheibe, T. D., W. A. Perkins, M. C. Richmond, M. I. McKinley, P. D. J. Romero-Gomez, J. A. Serkowski, and J. M. Zachara, "Pore-scale and multiscale numerical simulation of flow and transport in a laboratory-scale column," *Water Resources Research*, 51(2): 1023–1035, DOI: 10.1002/2014WR015959, 2015b.
- U.S. DOE. 2015. *Building Virtual Ecosystems: Computational Challenges for Mechanistic Modeling of Terrestrial Environments: Workshop Report*, DOE/SC-0171. U.S. Department of Energy Office of Science. <http://doesbr.org/VirtualEcosystems/>.

Toward Exascale Simulations of Microbially Mediated Biogeochemical Cycling from Pore to Watershed Scale

Carl Steefel, David Trebotich, Eoin Brodie, Sergi Molins, and Eric King

Lawrence Berkeley National Laboratory, Berkeley, California

1. Describe a major science challenge expected to be solved in the 2020–2025 time frame that requires using extant computing ecosystems.

A major science challenge for the next 5–10 years is to understand the role of microbial communities on nutrient and metal cycling in terrestrial environments over a very large range of scales. Investigations of the distribution of organisms across aquifers and their responses to experimental stimulation and seasonal variations have led to a revised conceptualization of how the subsurface functions. Dramatic shifts in community composition are observed as the changing geochemical environment, in some cases perturbed by interactions with surface water and/or vegetation, selects from a microbial “seed bank.” These perturbations may operate at the finest pore scales in soils, or they may be important within the larger (km) scale watershed systems where surface and subsurface water interact with vegetation and the atmosphere. With the rapidly expanding accumulation of information about which functional traits occur in which organisms, it has become clear that the norm is not for one microbe to conduct all the steps of a biogeochemical cycle such as sulfur oxidation or nitrate reduction. The results of microbial community studies have provided direct insights into how biogeochemical models should be structured, but they also make clear the scientific challenge associated with understanding the mediation of complex biogeochemical reaction networks by multiple microbial communities with dynamically changing populations and functions.

In the biogeochemical reaction network, each microbial group within a functional guild is parameterized from metagenomic data with a unique combination of traits governing organism fitness under dynamic environmental conditions. The model simulates the thermodynamics of coupled electron donor and acceptor reactions to predict the energy available for cellular maintenance, respiration, biomass development, and enzyme production. In addition, the model allows for a faithful representation of the functional diversity of microbial populations, how microbial physiological traits affect fitness, how biogeochemical processes are affected by emerging microbial composition, and how biogeochemistry feeds back to alter microbial fitness and community assembly.

Ecosystem models are routinely employed to understand and predict biogeochemical dynamics, while the microbes responsible for reactions are either omitted or poorly represented. When included in models, microbes and their metabolism have been typically incorporated through, for example, constraint-based (Fang et al., 2012), kinetics-based (Loew and Schaff, 2001), or trait-based (Bouskill et al., 2012) approaches. However, constraint-based approaches linked to reaction transport models (RTMs) are computationally expensive, and it is a major exascale computational challenge to integrate these complex microbially mediated biogeochemical reaction networks into the reactive transport frameworks used to simulate heterogeneous environments, whether at the pore scale (soils) or larger watershed scale. The dense set of ODEs corresponding to the reaction network needs to be distributed across the complex heterogeneous domain by being linked to a set of PDEs for flow and transport at multiple scales. Problems with 2–10 billion degrees of freedom are likely. In addition, the simulators will need to incorporate data from large

genomic databases, as well as spatially distributed system properties either at the pore scale (e.g., x-ray synchrotron data images of the pore structure, see Figure 1) or at the watershed scale (e.g., high-resolution LiDAR data for topography and vegetation mapping). Thus, storage requirements for these multivariate data sets will be severe as well (Figure 2).

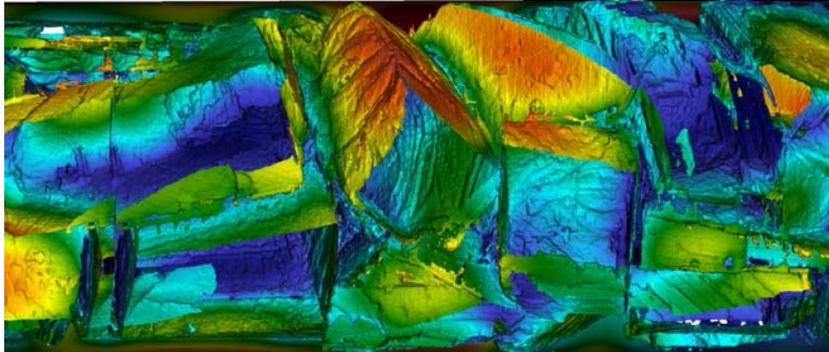


Figure 1: Digital image of the embedded boundary surfaces (pore and reactive solid phase) with superimposed pH values resulting from pore-scale flow and reaction (Molins et al., 2014).

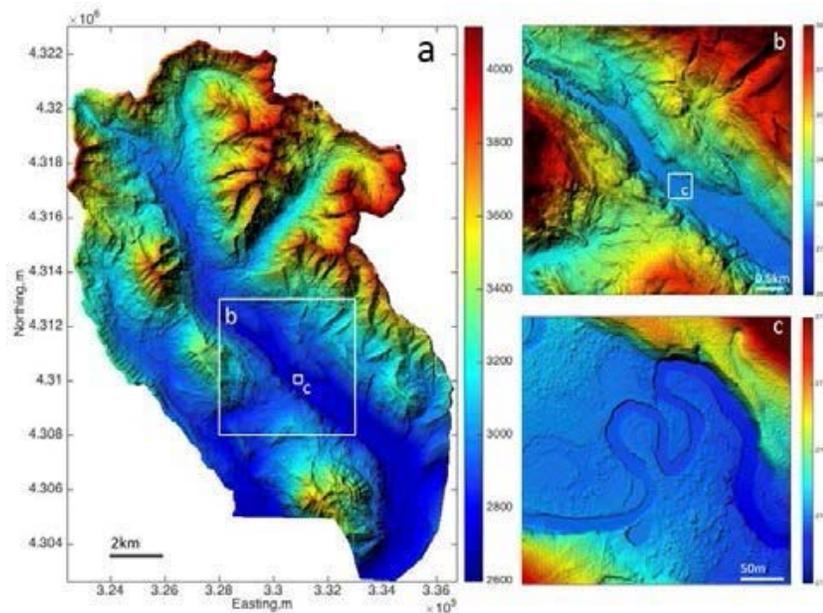


Figure 2: LiDAR-based digital elevation map of the East River Watershed, Colorado, illustrating the presence of nested heterogeneity.

Figure 3 depicts a BioCrunch simulation of microbial community and biogeochemical reaction network emergence developed for the Rifle, Colorado, floodplain, the first field site where these approaches were used. The challenge will become exascale when this approach is extended to complex heterogeneous pore-scale environments, as in soils interacting with vegetation and the atmosphere, or to the larger watershed scale with various terrestrial components (surface water, subsurface water, snow, and vegetation), with tracking of as many as 100 distinct microbial communities and 150 geochemical species within a multi-scale framework.

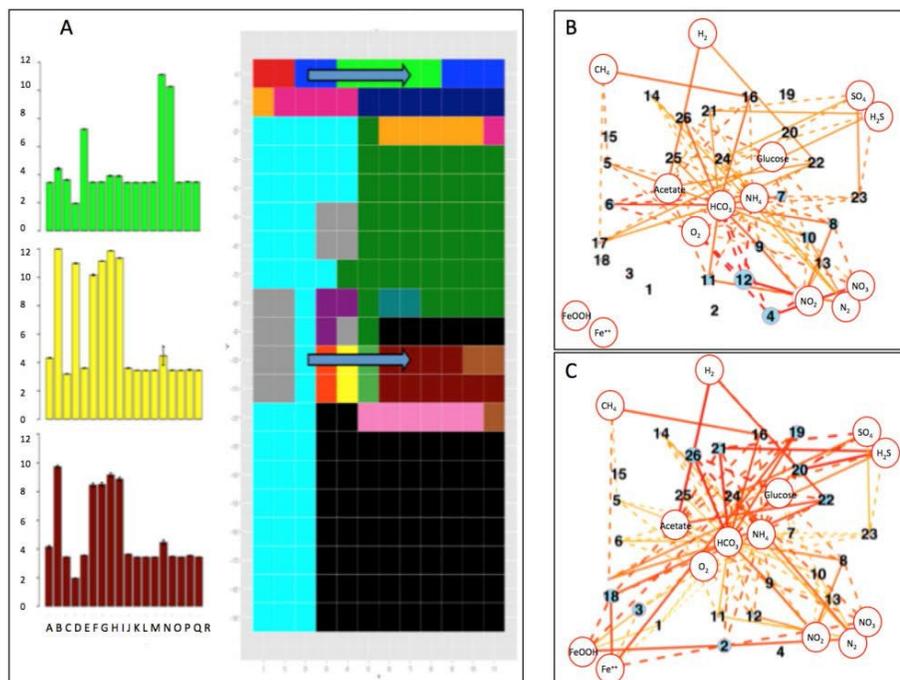


Figure 3: (A) Depiction of community emergence based on a 300-day simulation representative of conditions found in the Rifle, Colorado, floodplain. Each color represents the community composition, with three examples shown in bar plots along the left column. Y-axis values represent relative differences in log concentrations of biomass with each species identifier given on the x-axis. (B,C) Depiction of the reaction network at the upper (oxic, panel B) and lower (anoxic, panel C) arrows in panel A. Numbered vertices provide reaction labels (see below) with the relative log abundance of the microbe catalyzing the reaction indicated by the size of the blue circle. Solid lines indicate that the species is being created, while dashed lines indicate it being used by a numbered reaction. The size of the lines, along with the color (yellow<orange<red), indicates the relative log rate of species creation/destruction (King et al., in preparation).

2. Current and future computational and data strategies

Models

We propose to use the suite of application codes based on the production code ChomboCrunch (Molins et al., 2012, 2014; Trebotich et al., 2014). ChomboCrunch is a pore-to-continuum flow and reactive transport simulator based on the Chombo adaptive framework and the CrunchFlow multi-component reaction network. BioCrunch provides the coupling to microbial community dynamics through this interface, which links the dense set of ODEs to the spatially distributed set of PDEs for the pore-scale and watershed-scale problem (King et al., in preparation). At the pore scale, x-ray synchrotron data will be used to develop digital renditions of the heterogeneous soil pore structure, with flow based on the Navier-Stokes or

Stokes equation coupled to reaction network (Figure 1). Currently, Chombo assumes Fickian diffusion, but incorporation of the full Nernst-Planck equation is planned so as to handle multi-component diffusion and the electrochemical effects that are important in clay-rich media (Steeffel et al., 2015; Tournassat and Steefel, 2015). Currently under development is a Chombo-based watershed hydrology (Richards equations) and surface waters (shallow water, kinetic wave equation) simulator, in part based on the code parallel integrated surface-subsurface water code ParFlow (Kollet and Maxwell, 2006; Kollet et al., 2010). This last is particularly challenging due to the presence of nested heterogeneity (Figure 2).

New Capabilities

Highly heterogeneous media such as subsurface root-soil systems present the potential for load imbalance, making semistatic load balancing (even with introspection) untenable. A patch integration performance model is not accurate enough even now on NERSC Edison for flat MPI execution. OpenMP dynamic tasks might be sufficient, but this will require a PETSc matrix assembly to support threaded matrix creation for our AMG solvers. Matrix entries and non-zero structure change in an adaptive computation. Solver initialization is harder to scale up than the solver itself. Furthermore, an exascale energy budget means disk access is 100 times more time- and energy-intensive than DRAM access, which is 100 times more energy- and time-intensive than cache data access. Hence, data analysis must be pushed up into the simulation for exascale machines. Resilience through GVR or containment domains will also be required to do local recovery and robust execution in the presence of platform degradation. Stable resiliency technologies could lead to a significant I/O win by a discard of checkpoint/restart routines at the application level. We need to restart from a resilience API, which means resilience needs to expand its scope to include versioning out to long-term nonvolatile storage, as well as reconstitute a simulation in-place with resource degradation.

Resolution

At the pore scale, we will resolve pore structures down to 100-nm resolution based on x-ray synchrotron and FIB-SEM mapping. A more typical resolution will be a 1- to 2-micron discretization based on mapping at similar scale. The domains will extend typically up to scales 10,000 times larger, resulting in problems with up to 10 B DOF when the complex microbially mediated reaction networks (including microbial communities and aqueous- and solid-phase chemical components) are included.

At the watershed scale, we will model the mountainous East River Colorado Watershed, a 100-km² area, with local spatial resolution of 1 meter (and finer using AMR) for the terrestrial system. We propose full treatment of the coupled biogeochemical system. We will resolve a coupled surface-subsurface-soil hydrologic system along with dynamic vegetation, all conducted with high time resolution to capture hydrologic transients like storms and other extreme events. We propose full treatment of the coupled biogeochemical and microbial community, with consideration of up to 20 "microbial guilds" along with 20 chemical components. We will resolve a coupled surface-subsurface-soil hydrologic system along with dynamic vegetation, all done with high time resolution to capture hydrologic transients like storms and other extreme events. In addition, a larger domain (25 km by 25 km, corresponding to a GCM grid cell) with a reduced order treatment of microbial community dynamics and biogeochemistry and a locally 1- to 10-meter resolution using AMR will be included. An exascale (200 petaflop) machine will be needed as none of these targets are achievable with existing or intermediate, near-term architectures.

I/O

Per production run, the production code, Chombo-Crunch, for example, dumps approximately 1,000 plot files and checkpoint files, currently on the order of 1–8 TB each. Files are stored on HPSS systems using hsi, for a total of 1–2 PBs of offline storage while maintaining 100 TBs of online storage during production runs for visualization and post-processing. File sizes, and thus storage requirements, will obviously scale with problem size. In particular for microbially mediated watershed function, additional microbial guilds will add to the requirements; therefore the simulation will have to dump a reduced amount of variables on future systems or derive lightweight data structures to recover state variables.

Many-Core and/or GPU Readiness

The Chombo-based suite of codes currently scale up to 256K cores on 2 petaflop machines and are memory bandwidth limited. Future architectures will provide a modest amount of high bandwidth memory. Chombo-Crunch has reduced its computation working set to fit into 25% of the high bandwidth memory on NERSC Cori Phase 1 with fully occupied cores. Used as a scratchpad, this working memory programming model can be translated to GPUs. In the long term, Chombo-Crunch will utilize an embedded DSL AMRShift with an unstructured stencil language. Chombo-Crunch is currently working out next-generation adoption through the NERSC Exascale Application Program (NESAP).

Work flows

Production runs usually require 10,000–100,000 cores for 36 hours. Plot and checkpoint files are typically dumped every 100 time steps. Checkpoint files are deleted once it has been determined they are not needed. Data are archived on hpss. Visualization and data analytics are currently performed after the run using VisIt, which is launched locally and uses cores on the supercomputer. For future systems, data analysis will be pushed up into the simulation. We are beginning to make use of off-chip in transit data visualization and analysis using the Burst Buffer, or DataWarp, system on NERSC Cori.

References

- Bouskill, N. J., J. Tang, W. J. Riley, and E. L. Brodie (2012), Trait-based representation of biological nitrification: model development, testing, and predicted community composition, *Frontiers in Microbiology*, 3, 364, DOI: 10.3389/fmicb.2012.00364.
- Fang, Y., M. J. Wilkins, S. B. Yabusaki, M. S. Lipton, and P. E. Long (2012), Evaluation of a genome-scale *in silico* metabolic model of *Geobacter metallireducens* using Proteomic data from a field biostimulation experiment, *Applied and Environmental Microbiology*, 78(24), 8735–8742, DOI: 10.1128/AEM.01795-12.
- King, E., E. L. Brodie, and C. I. Steefel (in preparation), Genome-enabled reactive transport modeling reveals microbial metabolic interdependencies in space and time, *Nat Communications*.
- Kollet S.J., and R.M. Maxwell (2006), Integrated surface–groundwater flow modeling: a freesurface overland flow boundary condition in a parallel groundwater flow Model, *Advances in Water Resources* 29(7):945–58.
- Kollet S.J., R.M. Maxwell, C.S. Woodward, S.G. Smith, J. Vanderborght, H. Vereecken, and C. Simmer (2010), Proof-of-concept of regional scale hydrologic simulations at hydrologic resolution utilizing massively parallel computer resources, *Water Resources Research*, 46:W04201.

- Loew, L. M., and J. C. Schaff (2001), The Virtual Cell: a software environment for computational cell biology, *Trends Biotechnology*, 19(10), 401–406, DOI: 10.1016/S0167-7799(01)01740-1.
- Molins, S., D. Trebotich, C. I. Steefel, and C. Shen (2012), An investigation of the effect of pore scale flow on average geochemical reaction rates using direct numerical simulation. *Water Resources Research* 48, W03527, doi:10.1029/2011WR011404.
- Molins, S., Trebotich, D., Yang, L., Ajo-Franklin, J.B., Ligoeki, T.J., Shen, C., Steefel, C.I. (2014), Pore-scale controls on calcite dissolution rates from flow-through laboratory and numerical experiments. *Environmental Science and Technology*, DOI: 10.1021/es5013438.
- Steefel, C.I., Appelo, C.A.J., Arora, B., Jacques, D., Kalbacher, T., Kolditz, O., Lagneau, V., Lichtner, P.C., Mayer, K.U., Meeussen, J.C.L., Molins, S., Moulton, D., Shao, H., Šimůnek, J., Spycher, N., Yabusaki, S.B., and Yeh, G.T. (2015), Reactive transport codes for subsurface environmental simulation, *Computational Geosciences* 19: 445-478. DOI: 10.1007/s10596-014-9443-x.
- Tournassat, C., and C.I Steefel (2015), Ionic transport in nano-porous clays with consideration of electrostatic effects. *Reviews in Mineralogy and Geochemistry* 80:287–329.
- Trebotich, D.P., Adams, M.F., Molins, S., Steefel, C.I., and Shen, C. (2014), High resolution simulation of pore scale reactive transport processes associated with carbon sequestration, *Computing in Science and Engineering*, 10.1109/MCSE.2014.77.

Turning Uncertainty into Opportunity by Advancing Theory and Models

Will Wieder,^{a,b} Mark A. Bradford,^c Stuart Grandy,^d Jenifer Talbot^e

^a National Center for Atmospheric Research, Boulder, CO

^b University of Colorado, Boulder, CO

^c Yale University, New Haven, CT

^d University of New Hampshire, Durham, NH

^e Boston University, Boston, MA

Topic: From Microbes to the Environment.

Summary: Committed investment is needed to link multidisciplinary insights into soil processes with predictive models. Specific research aims should: (1) apply theoretical understanding of belowground processes in numerical models; and (2) investigate how data from new technologies can and build confidence in biogeochemical projections.

1. Key challenge: applying theory in models to advance understanding

High uncertainty and low confidence in terrestrial carbon (C) cycle projections reflect the incomplete understanding of how best to represent biologically driven C cycle processes at global scales.^{1–3} Advancing our understanding and representation of belowground processes in the face of environmental change is a key challenge facing earth system science. Addressing this challenge requires both **mechanistic insight** that can be applied in process-based models and **datasets that can evaluate models** across spatial and temporal scales.^{4,5} For example, contemporary thinking emphasizes microbial transformations as a major process driving stable soil organic matter (SOM) formation.^{6–9} Thus, microbial physiology (e.g., growth rates and growth efficiencies) and products (e.g., enzymes and necromass) are responsible for both the stabilization and loss of C from land surfaces through their effects on SOM.¹⁰

Under this framework, factors affecting the functional composition and activity of soil microbial communities may ultimately determine the balance between stabilization and decomposition of SOM. Although new sequencing technologies generate unprecedented insight into the genes, proteins, and metabolites present in soil systems, progress beyond these descriptive metrics remains measured. Recent advances in the functional categorization of microbial diversity^{11–13} and the application of such insights in numerical models¹⁴ afford new opportunities to accelerate linkages between soil ecological community structure and function, and their response to environmental change.

Observations provided by new technologies and experimental approaches are rapidly expanding the spatial and temporal scales over which measurements are made. Such advances reshape and refine the understanding of belowground processes¹⁵ and provide opportunities to meet fundamental scientific challenges. However, linking insight provided by these disparate data streams, which span orders of magnitude in spatial and temporal scale, to the global C cycle is needed to address societally and policy-relevant issues. The aim should be to improve confidence in belowground projections by identifying and representing key mechanisms such as microbial activity and physicochemical stabilization of soil organic matter under different environmental conditions. We suggest focusing in the near-term on improving confidence in model projections, as opposed to reducing uncertainty, because with better process-based representation, we can be confident that gains in certainty are actually meaningful. That is, many forms of uncertainty exist when modelling climate change and associated biosphere feedbacks, but they do not equally contribute to the *confidence* one has that projected changes will actually occur. We suggest that the primary, near-term focus should be on building confidence in “uncertain” projections by using the new technologies to guide how improved theoretical understanding of microbial processes can be best represented in models.

2. Scaling microbial understanding to regional and global biogeochemical function

Ecosystem theories, and consequently biogeochemical models, are based on the assumption that different belowground communities function similarly. This assumption of “scale invariance” posits that environmental conditions will change the rate of ecosystem processes, but the biotic response will be consistent across sites.¹⁶ Thus, regional- to global-scale biogeochemical responses should be predictable without understanding site-specific differences in belowground communities.

Emerging evidence challenges the validity of scale invariant assumptions, suggesting instead that the activities of belowground communities are shaped by particular local environmental conditions, such as climate.^{17–19} Under the assumption of “Scale dependence,” relationships generated by evolutionary tradeoffs in acquiring resources and withstanding environmental stress dictate the activities of belowground communities and their functional response to environmental change. Thus, understanding and projections of biogeochemical processes should require quantification and mechanistic evaluation of the controls on belowground activity at local scales.

Evaluating the assumptions of scale invariance versus scale dependence is critical to advancing the scientific understanding of microbial-mediated biogeochemical cycles and their application in models. Addressing this question over the next decade will require intellectual, experimental, and computational resources aimed at:

- A. Using microbial-omics to define microbial functional traits that are broadly characteristic of different ecological strategies and to identify the environmental or edaphic factors responsible for determining their relative abundances.
- B. Identifying how interactions between plant chemistry, microbial products, and soil mineralogy influence the stabilization of SOM and its response to environmental change.
- C. Scaling (from A and B) to predict ecosystem-level soil C stabilization by incorporating microbial functional traits into models that assess regional and global biogeochemical responses on a changing planet.

To address these challenges we need greater investment, integration, and communication among empiricists and modelers. These collaborations should transcend disciplinary boundaries to apply theory, generate hypotheses, collect data, and refine understanding through an iterative exchange of ideas and information. The work should *not* be motivated by the need to parameterize models. Instead, the primary aim of such work should be to provide a deep understanding of the factors regulating soil biogeochemical processes across scales. Developing such knowledge is critical to advance basic and applied research in microbial, agricultural, ecosystem, and earth system science. Building these interdisciplinary connections will expand scientific opportunities and address societal needs to more accurately project belowground responses and feedbacks to environmental perturbations.

- 1 Arora, V. K., *et al.* Carbon-Concentration and Carbon-Climate Feedbacks in CMIP5 Earth System Models. *J Climate* 26, 5289–5314, doi:10.1175/jcli-d-12-00494.1 (2013).
- 2 Jones, C., *et al.* Twenty-First-Century Compatible CO₂ Emissions and Airborne Fraction Simulated by CMIP5 Earth System Models under Four Representative Concentration Pathways. *J Climate* 26, 4398–4413, doi:10.1175/jcli-d-12-00554.1 (2013).
- 3 Wieder, W. R., Cleveland, C. C., Smith, W. K., & Todd-Brown, K. Future productivity and carbon storage limited by terrestrial nutrient availability. *Nat. Geosci.* 8, 441–444, doi:10.1038/ngeo2413 (2015).
- 4 Luo, Y., *et al.* Toward more realistic projections of soil carbon dynamics by Earth system models. *Global Biogeochemical Cycles* 30, 40–56, doi:10.1002/2015GB005239 (2016).

- 5 Wieder, W. R., *et al.* Explicitly representing soil microbial processes in Earth system models. *Global Biogeochemical Cycles* 29, 1782–1800, doi:10.1002/2015gb005188 (2015).
- 6 Grandy, A. S., & Neff, J. C. Molecular C dynamics downstream: The biochemical decomposition sequence and its impact on soil organic matter structure and function. *Science of The Total Environment* 404, 297–307, doi:10.1016/j.scitotenv.2007.11.013 (2008).
- 7 Schmidt, M. W., *et al.* Persistence of soil organic matter as an ecosystem property. *Nature* 478, 49–56, doi:10.1038/nature10386 (2011).
- 8 Cotrufo, M. F., Wallenstein, M. D., Boot, C. M., Denef, K., & Paul, E. The Microbial Efficiency- Matrix Stabilization (MEMS) framework integrates plant litter decomposition with soil organic matter stabilization: do labile plant inputs form stable soil organic matter? *Global Change Biology* 19, 988–995, doi:10.1111/gcb.12113 (2013).
- 9 Lehmann, J., & Kleber, M. The contentious nature of soil organic matter. *Nature* 528, 60–68, doi:10.1038/nature16069 (2015).
- 10 Kallenbach, C. M., Grandy, A. S., Frey, S. D., & Diefendorf, A. F. Microbial physiology and necromass regulate agricultural soil carbon accumulation. *Soil Biology and Biochemistry* 91, 279–290, doi:10.1016/j.soilbio.2015.09.005 (2015).
- 11 Fierer, N., *et al.* Cross-biome metagenomic analyses of soil microbial communities and their functional attributes. *Proceedings of the National Academy of Sciences* 109, 21390–21395, doi:10.1073/pnas.1215210110 (2012).
- 12 Leff, J. W., *et al.* Consistent responses of soil microbial communities to elevated nutrient inputs in grasslands across the globe. *Proceedings of the National Academy of Sciences*, doi:10.1073/pnas.1508382112 (2015).
- 13 Tedersoo, L., *et al.* Global diversity and geography of soil fungi. *Science* 346, doi:10.1126/science.1256688 (2014).
- 14 Wieder, W. R., Grandy, A. S., Kallenbach, C. M., Taylor, P. G., & Bonan, G. B. Representing life in the Earth system with soil microbial functional traits in the MIMICS model. *Geoscientific Model Development* 8, 1789–1808, doi:10.5194/gmd-8-1789-2015 (2015).
- 15 Hinckley, E.-L. S., Wieder, W., Fierer, N., & Paul, E. Digging into the World beneath Our Feet: Bridging across Scales in the Age of Global Change. *Eos, Transactions American Geophysical Union* 95, 96–97, doi:10.1002/2014eo110004 (2014).
- 16 Bradford, M. A., & Fierer, N. in *Soil Ecology and Ecosystem Services* (eds. D. H. Wall *et al.*) 189–200 (Oxford University Press, 2012).
- 17 Averill, C., Waring, B. G., & Hawkes, C. V. Historical precipitation predictably alters the shape and magnitude of microbial functional response to soil moisture. *Global Change Biology*, n/a- n/a, doi:10.1111/gcb.13219 (2016).
- 18 Strickland, M. S., Keiser, A. D., & Bradford, M. A. Climate history shapes contemporary leaf litter decomposition. *Biogeochemistry* 122, 165–174, doi:10.1007/s10533-014-0065-0 (2015).
- 19 Crowther, T. W., *et al.* Biotic interactions mediate soil microbial feedbacks to climate change. *Proceedings of the National Academy of Sciences* 112, 7033–7038, doi:10.1073/pnas.1502956112 (2015).

Biological Big Data Challenges (corresponds to Section 3.1.4)**Genomes to predictive biology: cognitive integration of intra- and inter-kingdom functional genomics datasets for dynamic, real-time, big data contextualization and interpretation for gene function discovery**

Ian Blaby,¹ iblaby@bnl.gov; Crysten Blaby-Haas,¹ cblaby@bnl.gov; Shinjae Yoo,² sjyoo@bnl.gov

¹Biology Department, Brookhaven National Laboratory, 50 Bell Avenue, Building 463, Upton, NY 11973, USA ²Computational Science Initiative, Brookhaven National Laboratory, 50 Bell Avenue, Building 463B, Upton, NY 11973, USA

1. Please specify the current science drivers for your field of research.

Functional genomics uses genomic and post-genomic data to understand how genomes are translated into biological processes, living organisms, and communities. This understanding provides the foundation for the systems and predictive biology that forms a cornerstone in DOE-sponsored biological research. A key revelation made with early genome sequences, and subsequently echoed with each new genome, is that we only know the function of a small percentage of genes. At the same time, a genome sequence enables genome-wide investigations that can collect experimental data for large numbers of genes and gene products. Functional genomics lies at the interface of these two observations and strives to employ and integrate genome-wide analyses to uncover gene function.

2. Describe the science challenges expected to be solved in the 2020–2025 time frame using extant computing ecosystems.

It is likely that the number of genomes and genome-based datasets will continue to grow at an exponential rate, while accurate gene annotations will progress linearly. Extant computing ecosystems will provide individual scientists with greater access to tools for analysis of individual experiments. Researchers will be able to analyze data in a modular manner (i.e., one data type at a time). Functional annotation pipelines are expected to become more proficient and discriminatory. We expect to see more organism-specific and data-type-specific network analysis (for example, organism-specific gene co-expression databases). These analyses will remain largely static and provide a snapshot of functional data at the time of analysis. We expect that metabolic modeling will be applied to an ever-growing number of organisms, but the value of those models will plateau.

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

A functional genomics dataset provides a single snapshot of a system's behavior. To provide a truly holistic and systems-level perspective, datasets will need to be interpreted in the context of all other pre-existing experiments. This virtual gene-to-function-to-system platform will require the ability to make logical links between experimental data via a co-inference network. This processing must occur in real time so that (i) the pre-existing network is calibrated in the context of new data and (ii) every researcher can participate with their own data as it becomes available rather than interacting with a static, predetermined system. Confidence in an observation made in one experiment increases when a compatible observation is made in an entirely different type of experiment or different organism. A simple example follows: a specific gene has higher expression under nitrogen-limited growth compared to nitrogen-replete in an alga than 500 other genes; a homolog of this gene is 1 out of 50 genes that when knocked out has a growth defect in the presence of protein as a nitrogen source in a land plant; in 1982, a paper was published on the proteolytic activity of a homologous protein in a bacterium. Therefore, the computing ecosystem would need to recognize the significance between these observations, including machine-learning, to extract experimental data and

observations from the literary archive. Generation of co-inference networks would produce the hypothesis that this gene encodes a protein that degrades proteins to salvage nitrogen in plants and algae.

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

| Accelerate | Why? |
|---|--|
| 1. Incremental learning algorithms | Incremental learning algorithms can handle constant updates of large-volume data, so that we can build up-to-date predictive biological models without re-computing predictive models from scratch |
| 2. Multimodal analysis | Multimodal analysis is required to handle diverse data types and sources, enabling better data integration beyond homology |
| 3. Extrapolation of data from the literary experimental archive | To accelerate knowledge discovery using experimentally validated (gold-standard) functional characterization |

| Impede | Why? |
|--|--|
| 1. Absence of an open computing platform | Non-universal access to a platform to develop and deploy the described system |
| 2. Data sharing | Literature, experimental data, and reprocessed data sharing are key to successful implementation and dissemination to the research community |
| 3. Workforce development | To develop the sophisticated system that can achieve big data research goals requires diverse expertise, including domain science, statistics, system architecture, databases, knowledge integration, etc. |

5. Characterize the data ecosystem aspects if the primary drivers for your field of research involve the transmission, analysis (including real-time analysis), or processing of data.

A tool with the capabilities that we are describing will require the need to handle highly variable datasets from experimental data, including transcriptome, proteome, phylogenetic evidence, and published journal articles. The vast number of previously acquired datasets means that we face a high **volume** of data and our model is based on such diverse datasets (and newly uploaded datasets) being updated in real time. A single expression profile may require several Mb (since we do not envision working directly with raw experimental datasets but instead will rely on processed data directly indicating mRNA/protein abundance, mutant growth rate, etc.), and an experiment may comprise 50 conditions/samples. Although each file is of manageable size, to take advantage of all available expression datasets in all organisms we should anticipate requiring Tb(s) storage (SRA presently contains >66k records, each on the order of tens of Mb). However, while transcriptomic, such as RNA-Seq, data has benefitted from cost reductions more than other experiment types, this value does not encompass proteomes, metabolomes, etc., or critically, the $\sim 2.1 \times 10^7$ records in PubMed. Another concern is the **veracity** of data given the fact that we contextualize all available evidence and the reliability of each source and each piece of evidence is different. Specifically, it is highly unlikely that experiments performed by different laboratories are conducted under identical conditions. Consequently, conclusions drawn may be due to these differing conditions and not due to genetic perturbations. To overcome our concerns regarding volume, variety, and veracity, we may consider **streaming multimodal analysis**.

Mass Spectrometry Imaging and LC-MS: Data Storage, Access, and Analysis

Ben Bowen

Lawrence Berkeley National Laboratory, Berkeley, California

Bioremediation, biomanufacturing, precision agriculture, biomolecular nanotechnology, bioenergy, and other concepts have emerged as industrial goals with horizons on a 20-year time scale. Applying robust engineering models toward these goals will require greater understanding of the complex biological processes involved. To achieve these goals, mass-spectrometry-based metabolite assessment will likely be a necessary capability. Currently, samples can be routinely analyzed for metabolite characterization, and the surface is only being scratched in terms of the comprehensive measurement that is feasible and needed. Although identified as areas of need, standardizations in sample preparation, data acquisition, quality control, data analysis, data management, and knowledge storage are missing.

These objectives will be realized through integration of advanced tools for biological analysis with computational tools for management of petabytes of scientific observations, real-time integration of new results with historical findings, identification of best practices for sample preparation/analysis, and dereplication of results to identify new insight from repetition.

Existing computing ecosystems that lower the barrier to integrated data analysis are empowering everyday researchers to do extraordinary computational tasks. For example, the metagenomic binning process, where short-read nucleotide sequences are assembled into a long, continuous sequence, is an urgent need that will likely be met in the 2020–2025 time frame. Meeting this need will involve both the improvement of assembly algorithms and, to a greater extent, new sequencing technologies that measure the sequence of long reads directly. Thus, in the 2020–2025 time frame, scientists will be measuring whole genomes from the majority of microbes living in environments.

What lies beyond is understanding the mechanisms associated with environmental selection of microbes through existence of required metabolic capabilities. Metabolic characteristics of microbial communities, their interspecies interactions, and their interactions with the environment are the furnace of global cycles. A mechanistic understanding of these processes will give humanity the potential to mitigate climate change, stabilize food supplies, and produce high-value hydrocarbons without relying on fossil fuels. Already, mass spectrometry integrated with chromatography (LC-MS) and mass spectrometry imaging (MSI) are primary tools used for these measurements.

The features of computing ecosystems that will enable or impede a more effective application of LCMS and MSI in the future are highlighted below.

| Accelerate | Why? |
|--|---|
| 1. Models and ontologies | Data integration will require ontologies that push models and algorithms toward physical law and facilitate meaningful re-use of knowledge. |
| 2. Data work flow and management | Each mass spectrometry lab is tasked with being its own IT department. Centralizing the task of primary data analysis for standardized work flows will lower the responsibility for data management. |
| 3. Workforce development and easy-to-use, yet powerful, computing ecosystems | The phrase, “biologists can’t analyze their own data,, is over-used and outdated. We are rapidly transitioning to the point where everyone analyzes their own data. Many computational biologists are performing unnecessary, simple tasks that a small amount of training and easy-to-use computing ecosystems will eliminate. |

| Impede | Why? |
|---|---|
| 1. Other: unrealistic expectations of large-scale analysis and data integration | The promise of large-scale measurement techniques is often over-hyped. By increasing the transparency and educating scientists about the limitations of specific techniques, experiments and analysis will be designed with realistic expectations in mind. |
| 2. Other: reliance on correlation, p-values, and machine learning, instead of outputs that suggest testable physical mechanisms and magnitude of change | Often probabilities and correlation analysis supersede physical law and fold change. This has the potential for discovery of new, unexpected discoveries, but often is not integrated with fundamental principles. |

References

1. Bowen, B.P., and Northen, T.R. (2010). Dealing with the unknown: metabolomics and metabolite atlases. *Journal of the American Society for Mass Spectrometry*, 21(9), 1471–1476.
2. Baran, R., Brodie, E.L., Mayberry-Lewis, J., Hummel, E., Nunes Da Rocha, U., Chakraborty, R., Bowen, B.P., Karaoz, U., Cadillo-Quiroz, H., Garcia-Pichel, F., and Northen, T.R. (2015). Exometabolite niche partitioning among sympatric soil bacteria. *Nature Communications*, 2015 6(6): 8289. DOI:10.1038/ncomms9289.
3. Fischer, C.R., Ruebel, O., and Bowen, B.P. (2016). An accessible, scalable ecosystem for enabling and sharing diverse mass spectrometry imaging analyses. *Archives of Biochemistry and Biophysics*, 589, 18–26.
4. Jenkins, H., Hardy, N., Beckmann, M., Draper, J., Smith, A.R., Taylor, J., et al. (2004). A proposed framework for the description of plant metabolomics experiments and their results. *Nature Biotechnology*, 22(12), 1601–1606. <http://doi.org/10.1038/nbt1041>.
5. Kell, D.B., Brown, M., Davey, H.M., Dunn, W.B., Spasic, I., and Oliver, S.G. (2005). Metabolic footprinting and systems biology: the medium is the message. *Nature Reviews Microbiology*, 3(7), 557–565. <http://doi.org/10.1038/nrmicro1177>.
6. Rübél, O., Greiner, A., Cholia, S., Louie, K., Bethel, E.W., Northen, T.R., and Bowen, B.P. (2013). OpenMSI: A High-Performance Web-Based Platform for Mass Spectrometry Imaging. *Analytical Chemistry*, 85(21), 10354–10361. <http://doi.org/10.1021/ac402540a>.
7. Silva, L.P., and Northen, T.R. (2015). Exometabolomics and MSI: deconstructing how cells interact to transform their small molecule environment. *Current Opinion in Biotechnology*, 34, 209–216. <http://doi.org/10.1016/j.copbio.2015.03.015>.
8. Tautenhahn, R., Patti, G.J., Rinehart, D., and Siuzdak, G. (2012). XCMS Online: a web-based platform to process untargeted metabolomic data. *Analytical Chemistry*, 84(11), 5035–5039. <http://doi.org/10.1021/ac300698c>.
9. Yao, Y., Bowen, B.P., Baron, D., and Poznanski, D. (2015a). SciDB for High-Performance Array-Structured Science Data at NERSC. *Computing in Science & Engineering*, 17(3), 44–52.
10. Yao, Y., Sun, T., Wang, T., Ruebel, O., Northen, T., and Bowen, B.P. (2015b). Analysis of Metabolomics Datasets with High-Performance Computing and Metabolite Atlases. *Metabolites*, 5(3), 431–442.

Big Data for Knowledge Discovery in Biology and Multi-scale Simulation of Biophysical Processes from Molecules to Biological Interfaces

James (Ben) Brown (Lawrence Berkeley National Laboratory)

1. Please specify the current science drivers for your field of research.

Field

Understanding ecosystem functions in the context of anthropogenesis.

Approach

Use perturbations to dissect interactions in microbial communities and complex ecosystems. We leverage principles from exposure biology and phylogenetics to model organismal and community environmental interactions across the entire tree of life.

Rationale

Genetic analysis, in which we modulate the activity of genes one at a time, has been incisive for understanding gene function in individual species. Similarly, chemical perturbations to ecologies or microbial communities that target particular members or functions are revealing the organizational principles of ecosystem architectures.

Impact

Success will yield predictive models of ecologies and environmental systems, as well as new engineering principles for the design and construction of ecosystems with specified (desired) services.

Major science drivers

Omics technologies—particularly transcriptomics, metabolomics, and fluxomics and phenomics—multi-scale imaging, high-performance computing, algorithms based on stochastic multiplex networks, advances in feature selection in machine learning frameworks, advances in the detection of higher order interactions using machine learning.

Areas of rapid advancement and potential

We have developed a theoretical framework for the modeling of organismal functions across the entire tree of life using tractable collections of chemical perturbations across a small numbers of species (manuscript in preparation). Our current implementation of this framework demands computing decompositions of multi-tensors based on generalized PARAFAC2 models. Extant objective functions are nonconvex, and stability is a problem—new procedures for fitting latent variable models on multiplex networks are needed. In addition, ensemble approaches could be formulated that would short-cut part of the currently optimization-based challenges, but these would bring with them new computational demands.

2. Describe the science challenges expected to be solved in the 2020–2025 time frame using extant computing ecosystems.

We will derive quantitative and interpretable models for organismal and community responses to environmental perturbations. For a small cohort of chemical classes, we will derive biochemical and physiological models with substantial predictive power across a broad swath of the phylogeny. These models will have profound impacts on the design of pesticides, prebiotics and probiotics, other soil amendments, antimicrobials, and plant growth modifiers.

3. 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.” |

5. Characterize the data ecosystem aspects if the primary drivers for your field of research involve the transmission, analysis (including real-time analysis), or processing of data.

Velocity

We work primarily on real data generated by omics and imaging technologies. With the notable exception of nanopore sequencing data, the velocity is relatively slow, in that real-time analysis ala LHC is not required. Nanopore sequencing is likely to generate up to 1 TB/sec by 2020 in parallel implementations, and new algorithms are needed to cope with these data rates—suitable algorithms are in development in my group and elsewhere. After in situ processing, nanopore data volumes will be within a log of current sequencers and hence tractable. Mass spectrometry data, however, may remain less reducible, unless deep metabolome annotation projects improve data interpretability and therefore compressibility. Multi-scale imaging initiatives are poised to generate enormous data sets with substantial velocity deriving from single instruments.

Volume

Current data volumes range from hundreds of gigs to the low petascale (~1 PB). The time horizon for data sets >100 PB is likely in 2019–2023. Sequence data are already growing in volume faster than Moore’s law and hence will eventually pose foundational challenges to data storage; methods are in development at NCBI, EBI, and elsewhere. Metabolomic imaging has the potential to provide peta-scale images of single biological samples. By 2025, biology will likely have overtaken astronomy in both data volume and data velocity, and while much of the data will be highly reducible by “straightforward” in situ analysis, other objects, particularly images and spectra, will have lasting value in their raw states, creating challenges for data volume.

Variety

We have enormously diverse data-type and intriguing coregistration problems. Data heterogeneity is a foundational challenge in my field, and one of the principal foci of my group. Currently, it is challenging to conduct more than a few assays on a single sample, so replicated samples are used to coregister data, although this is intrinsically inexact. Weak or “fuzzy” data coregistration is a massive challenge in our field, and “massively parallel” molecular imaging has continued to prove elusive. Improved methods of assay coregistration will continue to improve the ease of integration, but jointly modeling broad varieties of data on (sometimes) radically different scales will continue to require foundational advances in mathematical statistics, machine learning, and computer science.

Veracity (the fourth V)

Beyond velocity, volume, and variety (the 3 Vs), Colbourne has proposed (and now IBM has adopted) the importance of data “veracity”—the degree by which data are reliable and relevant. The veracity of biological data is frequently challenging to confirm or quantify, and hence we often rely on measures of stability or reproducibility. New statistical procedures to assess and ensure stability are needed.

Combinatorial Applications for Systems Biology

Dan Jacobson, Oak Ridge National Laboratory

1. Please specify the current science drivers for your field of research.

The U.S. Department of Energy (DOE) and the U.S. Department of Agriculture have independently launched ambitious research agendas to accelerate the development of domestic, renewable alternatives to liquid fossil fuels and bio-based products. Fuels converted from cellulosic biomass offer one alternative to conventional energy sources, which in turn foster economic growth and energy security of the United States. Furthermore, DOE has interests in terrestrial ecosystems and has launched projects such as the Next Generation Ecosystem Experiment in order to better understand mechanisms and processes controlling primary production and carbon cycling, biogeochemistry, and the impacts of disturbance on terrestrial ecosystems. DOE has also expressed interest in precision medicine and the White House BRAIN Initiative (Brain Research through Advancing Innovative Neurotechnologies). As such, the DOE has strong interests in the fundamental understanding of complex biology and its interactions with the environment. Biological organisms are derived from complex genetic systems that are composed of pleiotropic functional networks of interacting molecules and macromolecules. The subsequent phenotypes are the result of orchestrated, hierarchal, varied collections of expressed genomic variants regulated by and related to biotic and abiotic signals. However, the measured effects of these genomic variants can be viewed as the result of historic selective pressure and current environmental as well as epigenetic interactions. Thus, their co-occurrence can be seen as genome-wide associations in a number of different manners. We are currently using data derived from the re-sequenced genomes from thousands of genotypes and tens of thousands of genomes in combination with transcriptomics, metabolomics, microbiomics and phenomics data. Genome-Wide Association Study (GWAS) networks, integrated with Single Nucleotide Polymorphism (SNP) correlations and co-expression networks, are already proving to be a powerful approach to gain a better understanding of cellular functions.

2. Describe the science challenges expected to be solved in the 2020–2025 time frame using extant computing ecosystems.

At present, we are working on petascale computational approaches that use combinatorics (rather than the simple pairwise comparisons often utilized) to explore the associations between genome structure and variance with physical and molecular phenotypes. In doing so, we aim to determine the pleiotropic and epistatic relationships underlying cellular functions that will thus allow us to gain further insights into the molecular basis of complex, multigenic interactions responsible for biological complexity and the emergent properties of complex systems. The tools that allow us to better understand complex biological systems at this level of sophistication will have applications to bioenergy, carbon cycles, ecosystem studies, climate, agronomics, neuroscience, precision medicine, and even material science and advanced manufacturing. As such, there is the potential for considerable scientific, economic, and societal impact.

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

Petascale systems will not allow us to explore as deep into the combinatorial space of biological function as we would like to in order to gain deeper and deeper understandings of the complex interactions responsible for the emergence of complicated structure and behavior of biological organisms and ecosystems.

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

| Accelerate | Why? |
|---|--|
| 1. Application codes | We are developing codes that are optimized for Titan and GPUs, which is greatly speeding up our applications and giving us access to expanded computing power. |
| 2. Algorithms | We are developing new algorithms that will make maximal use of HPC resources in order to make data-driven discoveries that have simply never been possible before. |
| 3. Visualization and analysis resources | We currently have reasonable visualization approaches, and it is possible that the utilization of resources such as Everest will enhance our current capabilities. |

| Impede | Why? |
|-----------------------|---|
| 1. Hardware resources | Limitations in compute power, memory, and I/O slow our work down. More resources will yield faster progress. |
| Workforce development | We have excellent collaborations with OLCF that provide us with high-level expertise in the creation of code optimized for GPUs. However, we will need to train/acquire personnel within our group in order to meet our growing need for GPU code development. |
| 3. Data workflow | At present, the migration of very large datasets across multiple sites and across a complex ecosystem of compute platforms/architectures is complicated and increases the time-to-answer for projects. We are working on solutions to these issues, but more effort is needed here. |

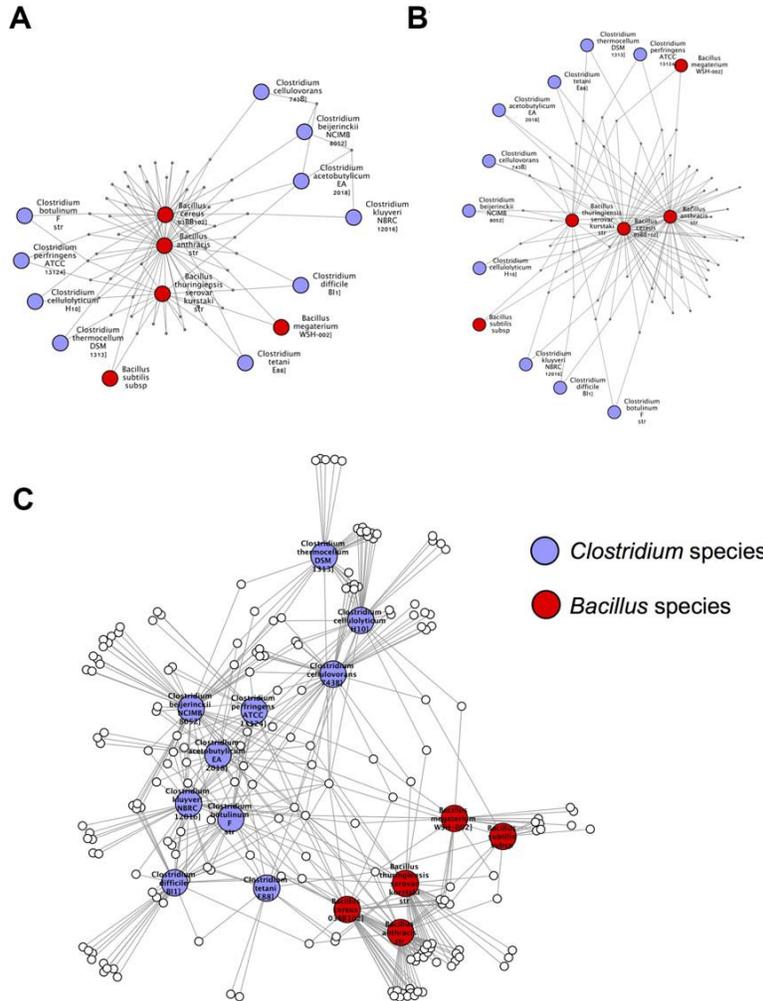
5. (Optional) Characterize the data ecosystem aspects if the primary drivers for your field of research involve the transmission, analysis (including real-time analysis), or processing of data.

Currently, all of our data come from experimental sources. In the future, we aim to couple this with modeling and simulation via land use and integrated earth system models. Data generation is accelerating exponentially and we need to store and have rapid access to all of our experimental data. We will combine, integrate, mine, and re-mine datasets many times with many different methods as part of an ongoing process to derive new knowledge from new and extant datasets. Our data is definitely multimodal, coming from many different experimental and field-based data sources/types and collection efforts. Our data can also have spatial and temporal dimensions. We are standing up larger and larger storage platforms as well as investigating archival solutions.

6. References

1. D. A. Weighill and D. A. Jacobson, “3-way Networks: Application of Hypergraphs for Modelling Increased Complexity in Comparative Genomics,” *PLoS Comput. Biol.*, vol. 11, no. 3, 2015.
2. D. Joubert, W. Nance, J. Weighill, and D. Jacobson, “GPU-enabled Parallel Similarity Metrics Computations for Genomics Applications,” in preparation, 2016.
3. D. Jacobson and G. Emerton, “GSA-PCA: gene set generation by principal component analysis of the Laplacian matrix of a metabolic network,” *BMC Bioinformatics*, vol. 13, no. 1, p. 197, Aug. 2012.
4. S. Climer, *Limit Crossing: A Guide for Solving Combinatorial Problems by Exploiting Bounds*. VDM Verlag, 2008.
5. S. Climer and W. Zhang, “A linear search strategy using bounds,” in *14th International Conference on Automated Planning and Scheduling - ICAPS'04*, 2004, pp. 132–141.
6. M. E. Setati, D. Jacobson, U.-C. Andong, and F. Bauer, “The vineyard yeast microbiome, a mixed model microbial map,” *PLoS One*, vol. 7, no. 12, p. e52609, Jan. 2012.
7. K. Zheng, X. Wang, D. A. Weighill, H.-B. Guo, M. Xie, Y. Yang, J. Yang, S. Wang, D. A. Jacobson, H. Guo, W. Muchero, G. A. Tuskan, and J.-G. Chen, “Characterization of DWARF14 Genes in *Populus*,” *Sci. Rep.*, vol. 6, p. 21593, 2016.
8. I. Porth, J. Klapte, O. Skyba, J. Hannemann, A. D. Mckown, R. D. Guy, S. P. Difazio, W. Muchero, P. Ranjan, G. A. Tuskan, M. C. Friedmann, J. Ehling, Q. C. B. Cronk, Y. A. El-Kassaby, C. J. Douglas, and S. D. Mansfield, “Genome-wide association mapping for wood characteristics in *Populus* identifies an array of candidate single nucleotide polymorphisms,” *New Phytol.*, vol. 200, no. 3, pp. 710–726, 2013.

7. (Optional) Images



$$C_3(X, Y, Z) = \frac{\frac{3}{2} \sum_i (\min(X_i, Y_i) + \min(X_i, Z_i) + \min(Y_i, Z_i) - \min(X_i, Y_i, Z_i))}{\sum_i (X_i + Y_i + Z_i)}$$

Subnetworks containing the *Clostridium* and *Bacillus* species selected from (a) 3-way best edge Sørensen Network; (b) 3-way best edge Czekanowski Network; and (c) Gene family enrichment network.

Modified from: D. A. Weighill and D. A. Jacobson, "3-way Networks: Application of Hypergraphs for Modelling Increased Complexity in Comparative Genomics," *PLoS Comput. Biol.*, vol. 11, no. 3, 2015.

Extreme-scale Partial Correlation Estimation for High-Dimensional Data Analysis

Sang-Yun Oh (syoh@pstat.ucsb.edu), Ariful Azad, Aydin Buluc (abuluc@lbl.gov), Penporn Koanantakool, Dmitriy Morozov, Leonid Oliker, Katherine Yelick (KAYelick@lbl.gov)

Overview: Characterizing relationships in high dimensional data is often an important research goal in many scientific disciplines, including environmental sciences and biology. In biology, for example, gene co-expression analysis and eQTL mapping are two popular methods for identifying so-called “omic” relationships. Gene co-expression study analyzes microarray or RNA-seq data for inferring gene regulatory networks (GRN) [1], and eQTL (expression quantitative trait loci) mapping mines gene expression and single nucleotide polymorphism (SNP) datasets jointly in order to determine pairwise gene co-expressions as well as gene-expression-to-genotype relationships [2, 3, 4, 5].

Estimating partial correlation is an important part of aforementioned methods. Compared to Pearson correlation, partial correlation models direct pairwise relationships between two variables after removing any indirect relationships through the rest. Therefore, partial correlation is more conducive to qualitative interpretation and is also more suitable for graphical representation in node-edge graphs, allowing for further analysis using graph algorithms. In fact, one of the most prominent open-source software programs for eQTL analysis, GeneNetwork, has support for partial correlations [10], albeit this is slow.

Despite the many advantageous properties of partial correlation, their application has been limited due to the computational complexity involved in computing their estimators. Typically, the computational complexity of partial correlation inference procedures on a p -dimensional dataset is $O(p^3)$ using $O(p^2)$ memory. Current state-of-the-art algorithms in the literature for partial correlation estimation can handle a number of features on the order of 200 K dimensional data in about 5 hours using 32 CPU cores with 128 GB of shared memory [6].

Considering that the numbers of genes and SNPs are in tens of thousands and millions, respectively, eQTL analysis using partial correlations in a general setting is currently intractable using a shared-memory machine. In fact, tradeoffs such as including only a small number of loci and/or assuming independence of all SNP-SNP pairs are unavoidable for computational tractability [4, 5]. To this end, a high-performance partial correlation estimation framework suitable for distributed parallel environment, dubbed HP-CONCORD, is being developed. HP-CONCORD computes an estimator according to the CONCORD framework [7, 8] on a distributed-memory parallel environment. The CONCORD estimator does not rely on a Gaussian assumption and, hence, is more robust to outliers.

Analyzing a dataset of size $p = 110K$, an initial proof-of-concept implementation of HP- CONCORD can estimate partial correlation structure (6 billion parameters) in less than 30 minutes on 512 cores. However, problem sizes are limited due to memory layout. To increase scalability in time and memory, a communication-avoiding sparse-dense linear algebra algorithm has been developed [9], with excellent scaling in both theory and experiments.

Impact: Understanding biological systems at the genome-wide level is perhaps one of the most important scientific challenges of our time. HP-CONCORD has the potential to be useful for many computational genomics analyses that require partial correlation computation, either as a final estimator or as an intermediate computation.

First, increased scalability to a distributed memory environment would allow more biomarkers to be included. As a result, tradeoffs necessary to attain computational tractability can be reduced. For example, partial correlation computation may be used to model linkage disequilibrium [5], and increased scalability would require fewer assumptions to be placed on dependency structure of SNPs.

Second, many biological problems will not fit in a single shared memory. Combining fast high-performance networks in future exascale systems with the distributed HP-CONCORD implementation will allow partial correlation computation for unprecedented datasets.

eQTL analysis: A potential use-case for HP-CONCORD may be for biologists to perform eQTL analysis for understanding plants that are candidates for biofuels at a genome-wide level. This genome-level understanding may lead to better genetic engineering strategies. Analysis at genome-wide level may reveal new insights by including biomarkers, which may have been removed by imposing assumptions for computational tractability.

Multimodal analysis: Note that eQTL mapping considers gene expression (abundance) and SNP calls (genetic variation) jointly (although raw reads come from the same technology). In this sense, eQTL analysis can be considered to be an example of multimodal analysis, which is part of a larger trend in data-driven scientific discovery. Paleoclimate reconstruction in environmental science analyzes temperature and proxies (tree rings, ice cores, etc.) jointly using partial correlations. Neuroscientists frequently use multimodal data to align and correct for individual variations. Functional connectivity between brain regions is often inferred with partial correlations. These areas are just a few examples where joint analysis of multimodal data can be potentially interesting. In multimodal analysis applications, within-mode and inter-mode dependency structures would be modeled jointly, resulting in significant increases in dimensionality.

Challenges: While currently ongoing HP-CONCORD research would increase the scalability of partial correlation estimation, more effort is needed to deploy full-fledged genomic analysis algorithms and their software implementations, which are challenging for independent reasons. Furthermore, an analysis method is often a part of a larger workflow, and interfacing to existing software tools is not always straightforward and requires domain specific expertise.

Computational challenges aside, a fundamental difficulty of high dimensional data analysis is gathering, storing, and sharing large enough coherent datasets and devising models for increasing statistical power. Noting that the number of free parameters is also $O(p^2)$, larger sample size is always desirable in order to achieve more accurate and stable estimates.

Regardless of these challenges, partial correlation is one of the fundamental measures of relationships in statistics. Reliable inference of partial correlations in a high-dimensional setting would prove useful in many real data analysis settings.

References

- [1] Shengjun Hong, Xiangning Chen, Li Jin, and Momiao Xiong. Canonical correlation analysis for RNA-seq co-expression networks. *Nucleic Acids Research*, 41(8):e95, 2013.
- [2] Jianxin Yin and Hongzhe Li. A sparse conditional Gaussian graphical model for analysis of genetical genomics data. *Ann. Appl. Stat.*, 5(4):2630–2650, 12, 2011.
- [3] Calvin McCarter and Seyoung Kim. Large-Scale Optimization Algorithms for Sparse Conditional Gaussian Graphical Models. ArXiv e-prints, September 2015 (to appear, AISTATS 2016).
- [4] Jen-hwa Chu, Scott T. Weiss, Vincent J. Carey, and Benjamin A. Raby. A graphical model approach for inferring large-scale networks integrating gene expression and genetic polymorphism. *BMC Systems Biology*, 3(1):1–9, 2009.
- [5] Raphael Mourad, Christine Sinoquet, and Philippe Leray. Probabilistic graphical models for genetic association studies. *Briefings in Bioinformatics*, 13(1):20–33, 2012.
- [6] Cho-Jui Hsieh, Matyas A. Sustik, Inderjit S. Dhillon, Pradeep K. Ravikumar, and Russell Poldrack. Big & QUIC: Sparse inverse covariance estimation for a million variables. In C.J.C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 3165–3173. Curran Associates, Inc., 2013.
- [7] Sang-Yun Oh, Onkar Dalal, Kshitij Khare, and Bala Rajaratnam. Optimization methods for sparse pseudo-likelihood graphical model selection. *Advances in Neural Information Processing Systems (NIPS) 27*, pages 667–675, 2014.
- [8] Kshitij Khare, Sang-Yun Oh, and Bala Rajaratnam. A convex pseudolikelihood framework for high dimensional partial correlation estimation with convergence guarantees. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 77(4):803– 825, 2015.
- [9] Penporn Koanantakool, Ariful Azad, Aydin Buluc, Dmitriy Morozov, Sang-Yun Oh, Leonid Olikier, and Katherine Yelick. Communication-avoiding parallel sparse-dense matrix-matrix multiplication. *International Parallel and Distributed Processing Symposium (IPDPS)*, IEEE Computer Society, 2016.
- [10] GeneNetwork, Frequently Asked Questions, <http://www.genenetwork.org/faq.html#Q-25>.

SINGLE CELL GENOMICS – CHALLENGES AND FUTURE PERSPECTIVES (BER WHITE PAPER)

Allon Wagner and Nir Yosef, UC-Berkeley

Please specify the current science drivers for your field of research.

Over the past decade, computational genomics has become a major catalyst of innovation in the biomedical sciences. Technological breakthroughs made over the past 3 years now allow, for the first time, massively parallel collection of genomics data (RNA, proteins, DNA, methylation, chromatin accessibility, histone marks) from single cells. The ability to study individual cells at that resolution opens research avenues that were thus far out of reach, but at the same time requires surmounting unique computational challenges.

Describe the science challenges expected to be solved in the 2020–2025 time frame using extant computing ecosystems.

Traditional *bulk genomics* averages molecular quantities across thousands of cells. In contrast, single-cell genomics allows one to observe the full spectrum of cellular phenotypes, which in turn allows distinct and hitherto unknown subtypes of cells to be discerned. Such discoveries have considerable implications for our understanding of human physiology in health and disease. For example, our lab has recently used single-cell RNA-Seq to characterize the heterogeneity of cells in the immune system, which led to the discovery of novel regulators of immune functions. Single-cell technology also has the potential to alleviate the need for large amounts of microbial DNA in metagenomics studies. Multiple studies that explore tissue heterogeneity and its implications for human health are currently underway, and will no doubt enrich our knowledge of the function and dysfunction of multiple organs in the following decade even if they rely only on extant computing infrastructure.

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

While single-cell genomics can contribute to the study of any biosystem in isolation, a great challenge is seeing the forest for the trees – namely, building a comprehensive atlas of all the human cell types, and studying their function as well as their intricate relationships across multiple experimental systems and varying conditions. This will require (a) scaling up the computing power available to process and store single-cell data, and (b) developing an infrastructure that will allow hundreds of labs around the country (and around the globe) to communicate their research findings and share their raw data. Assembling the human cell atlas, which will revolutionize our understanding of human physiology, will require considerable improvements in the current computing infrastructure before it can be attained.

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

We chart our perspective of these computational challenges in three domains: (a) statistical models and algorithms for single-cell genomics, (b) community ontologies and data-sharing practices, and (c) mega-scale analysis and visualization.

Challenge A: statistical models and algorithms for single-cell genomics [refers to option b: Models and algorithms]

The development of new generations of genomics technologies always spurred research on pertinent statistical issues. For example, next-generation RNA-Seq measures discrete reads in contrast to the continuous fluorescence of microarrays, which led to much interest in proper statistical models of count data, and sound methods to detect differentially expressed genes while controlling for unwanted variation. By the same token, the research community soon realized that single-cell RNA-Seq's most prominent characteristic was the considerable number of undetected genes (dropouts) it produced. This led to a surge of work on zero-inflated models and other statistical methods to account for the presence of dropouts carrying routine analyses, such as differential expression, gene set enrichment analysis, and dimensionality reduction. Furthermore, technical constraints in conducting single-cell RNA-Seq experiments result in a much

lower signal-to-noise ratio than in bulk, which led to research into computational methods for assessing and handling technical confounders and batch effects in single-cell experiments. Importantly, the community has not settled on best practices yet. We expect those to be established in the near future. This will lead to an acute need for the development of robust and high-quality software libraries that are accessible to the wide scientific community, in contrast to existing software packages that only experts in the domain can utilize.

Challenge B: community ontologies and data-sharing practices [refers to option *d*: Data workflow]

The greatest reward of single-cell genomics will arguably be the curation of a comprehensive atlas of human cell types. Much of the present work in single-cell RNA-Seq revolves around the delineation and validation of novel cell subtypes in various tissues. As these efforts grow in scale and are conducted by multiple labs around the globe, questions of a data workflow come to the forefront. A major effort to set ontologies and naming conventions is necessary to allow labs to share their discoveries, and to test whether they reproducibly observe the same cell types that other labs working in similar experimental systems observed. These issues are highly complex due to several factors. First and most important, although one may hope that distinct cellular subtypes will have conspicuous transcriptomic markers (such as one or two elevated marker genes), it is plausible that many important cell types can be characterized only by complex transcriptional signatures. Second, experimental protocols (both wet procedures and computational processing) have drastic effects on the observed transcriptomic features. Third, cellular identity, as reflected in the executed transcriptional program, is a superposition of multiple factors, such as spatial position, cell cycle progression, and stochastic fluctuation in RNA abundance. Importantly, some of these factors can be considered, depending on the research question at hand, either a technical confounder or an immanent facet of the biological phenotype (e.g., cell cycle progression in proliferating tumor cells).

Challenge C: mega-scale analysis and visualization [refers to option *e*: Visualization and analysis resources]

Whereas mass spectrometry studies now routinely measure small sets of proteins across millions of single cells, single-cell RNA-Seq has thus far been limited to several thousands of cells. However, in 2015, emerging droplet technologies allowed for the first time RNA abundance to be quantified genome-wide from up to ~50k cells, and their scale is projected to increase by orders of magnitude in the near future. These developments will enable much better understanding of single-cell behavior owing to increased statistical power and the better representation of rarely occurring cell types. However, they require careful attention to the scalability of computational building blocks that are now taken for granted (e.g., alignment algorithms used to quantify the gene expression, and the zero-inflated statistical inference discussed above). We expect that the scalability issue will be tackled by a combination of algorithmic advances (such as the fast approximation algorithm for tSNE, or the recent concept of pseudo-alignment) and infrastructure deployment (CPU, storage). The increased scale also requires rethinking of the exploratory data analysis tools at hand, such as devising visualizations to fit single-cell resolutions at these scales.

To give a sense of the problem domain, we consider two bottlenecks that determine the number of cells available in one study: (a) experimental capacity, and (b) sequencing costs. Experimental technologies have been making great strides, and within three years the state-of-the-art has increased by three orders of magnitude from dozens of cells per study (Figure 1). The decrease in sequencing costs over the past decade has been outpacing Moore’s law (Figure 2; source: National Human Genome Research Institute, <https://www.genome.gov/sequencingcosts>; the sharp step around 2008 corresponds to the emergence next-generation sequencing). Consequently, we expect that within the next decade we will be seeing single-cell RNA-Seq studies measuring on the order of hundreds of thousands or even millions of cells each. This magnitude is currently achieved only by mass-spectroscopy studies, which produce relatively little data on each individual cell compared with single-cell RNA-Seq (dozens and thousands of numerical features per cell, respectively). A rough estimate of the required resources is as follows:

Storage: 0.5–1 GB per cell for long-term storage of processed files. Alignment results (bam) files are the most data intensive, and one may choose not to store them and instead re-compute them when necessary (a tradeoff with CPU follows). Short-term storage used while processing each cells is one order of magnitude larger.

CPU: Between 30 and 180 minutes of CPU processing on one modern core per cell. Required resources depend on (a) sequencing depth, a parameter that the experimenter determines to balance cost versus information coverage; and (b) software used. Our estimate is based on standard alignment and quantification packages (Tuxedo suite and RSEM). Note, however, that if the recently proposed pseudo-alignment algorithm (implemented in Kallisto) becomes mainstream in single-cell experiments, then CPU time per cell will decrease by ~ 1 order of magnitude. We do not expect it to decrease further because pseudo-alignment opens the gate to novel data-handling practices through bootstrapping, which will certainly be adopted and require additional CPU processing per cell, thus offsetting the significant CPU gain usually obtained by pseudo-alignment.

This page is intentionally left blank.

C.2 White Papers Addressing Environmental Research

Atmospheric Simulation and Data Assimilation within the Earth System (corresponds to Section 3.2.1)

BER/CESD Exascale Facilities Workshop White Paper

SIMULATION OF CONVECTIVE SYSTEMS

William I. Gustafson Jr.

Atmospheric Sciences & Global Change Division, Pacific Northwest National Laboratory

This brief white paper will be a source document for a meeting scheduled in March, 2016, and the subsequent report for ASCR, BER, and the Office of Science. We ask for your help to determine the requirements for a computing ecosystem (which includes data, software, libraries/tools, as well as computing requirements) and the potential impact in advancing the scientific mission of Climate and Environmental Sciences.

Your program may have some of these materials requested below from previous exercises. In all of these sections, feel free to pull text from those efforts and provide references to those materials. If there are a few key papers, conferences, or reports you feel strongly bring greater detail to your white paper, please reference them.

1. Describe a major science challenge expected to be solved in the 2020-2025 timeframe that requires using extant computing ecosystems.

What will probably be solved in the next 5-10 years? Why is this important to the field? Please give a high-level description of your research and the roles played by high-end computing, storage, and networking.

Understanding the full convective lifecycle well enough to parameterize it in atmospheric models with grid spacings ranging from several to hundreds of kilometers is critical for advancing climate models. This is particularly important as computers begin to permit climate modeling with grid spacings in the 10's of kilometers, where convection is partially resolved. To-date, aspects of convection are understood, but with significant gaps. Of particular importance is gaining a better understanding of convective initiation, and upscale growth, which leads in some cases to large organized systems. The dynamics that sustain these organized convective systems is not fully understood, which makes parameterization difficult. And, cloud parameterizations for models coarser than about 5 km are unable to capture many of the system characteristics [Clark *et al.*, 2007]. This leads to incorrect timing of precipitation in models, e.g., with too much precipitation falling during the day when it should fall at night over portions of the central United States.

Current computers are on the verge of being able to reasonably simulate small convective systems using large-eddy simulations (LES's) with grid spacings close to 100 m. However, computational constraints require domains that are still too small to encompass the full convective lifecycle that can extend beyond a day, and therefore cannot simulate systems as they propagate across the country. Current methodologies to get around this are most commonly to use coarser grid spacings, on the order of 1–3 km. This permits reproducing many of the storm characteristics, but misses critical details needed for

fully understanding the processes, such as what controls the storm initiation and how merging and splitting of updraft cores drive the overall storm organization. Alternatively, smaller domains with higher resolution can be used for targeted periods and locations within the storm lifecycle. Ultimately, researchers need the ability to do routine, continental scale simulations at sub-100 m grid spacing to enable study of entire cloud populations and how they interact with host weather systems. This will permit a much clearer understanding of the convective lifecycle and, hopefully, will concurrently lead to better parameterizations for climate models.

The ability to do routine simulations of large convective systems will be a large part of the breakthrough needed to address this problem. Currently, only “hero” simulations can simulate realistic cloud populations over long periods. Global simulations with 800 m grid spacing have been used for studying cloud populations [Miyamoto *et al.*, 2013] and grid spacing closer to true LES resolution have been used for the area of small countries [Schalkwijk *et al.*, 2015]. But there is hope beyond these initial attempts. As LES codes with the complexity necessary to perform these sorts of simulations are better optimized for the next generation compute architectures, the ability to do these sorts of runs will become more routine for typical cloud researchers.

2. Current and Future Computational and Data Strategies

With respect to the challenge described above, please provide a short, high-level description of your computational and data strategies used to solve the challenge (or its precursor) today and how those might change in the next decade.

Models: *What types of models and/or modeling systems will be required?*

Models needed for this type of research are both global and limited area LES models. Most of the simulations would be better handled using limited area models, which decreases cost, and thus enables increased number of simulations for doing ensembles and testing scenarios. However, some simulations will require global domains to see the overall impact of the convection on the global-scale dynamics, where the clouds can interact with the synoptic-to-global wave patterns. Many LES models exist, with commonly used ones being the Weather Research and Forecast (WRF) model, which is one of the more versatile models at the trade-off of increased cost, the System for Atmospheric Modeling (SAM), which is used as part of the Community Atmosphere Model (CAM) superparameterization methodology, and the Dutch Atmospheric LES (DALES) model. These are a few examples, with many others in existence that are used to various degrees.

The difficulty with most of these models is that they were designed and coded prior to the common existence of accelerators in high-end computing. Typical usage on DOE supercomputers is to use the heavyweight cores and to neglect the lightweight cores, such as graphical processing units (GPUs) and Many Integrated Cores (MICs). This works today, but will become increasingly problematic as more and more Flops derive from the lightweight cores. Investment is needed to adapt the current models, which would benefit from the range of support software already available for these models. For example, WRF is capable of ingesting a wide range of boundary conditions due to its wide use in the research and forecasting communities. Where adapting current models becomes too difficult, one must consider developing new models that fully utilize state-of-the-art computing architectures, as well as those that are anticipated in the exascale era.

New Capabilities: *What new capabilities will have to be developed for these models? As compared to today's models, how will these new capabilities increase the computational cost and memory requirements?*

Computationally, the LES models must make efficient use of available resources, as noted above. However, there are scientific limitations as well to fully realize the vision of using LES for the full lifecycle of deep convection. Current understanding of ice within clouds is insufficient to accurately reproduce the full microphysics of mixed-phase clouds, even in models with the most detailed handling of microphysics. Investment is needed to understand this critical limitation and to then find ways to represent the relevant processes in the model. This would benefit both the LES models as well as the climate models, which both suffer from this problem.

Resolution: *What model resolution do you anticipate will be required? How does this compare to the resolutions currently in use? Can you describe the resolution requirements in more detail – e.g. global uniform resolution, regional resolution, regionally refined resolution or adaptive resolution?*

Current models used for simulating deep convection typically use grid spacings around 1 km. *Lebo and Morrison [2015]* identified a threshold around 250 m grid spacing where the cloud characteristics noticeably change, implying a grid spacing smaller than this will be required. Understanding how nighttime planetary boundary layers, which are decoupled from the surface, interact with convection will likely require much finer grid spacing due to the stable conditions and strong gradients associated with low-level jets. Grid spacing on the order of 10's of meters will be required for these sorts of situations.

I/O *As compared to today's models, how much additional I/O will be required? Will the increase be due to adding additional variables, or higher frequency or other considerations? How much of this data will need to be permanently archived?*

Increased I/O compared to today's simulations will be due primarily to the increased model resolution combined with larger domain extent. Depending on particular research needs, there might also be the requirement of higher output frequency, on the order of minutes, to better capture transient cloud phenomena. The increase in output going from 1 km to 100 m grid spacing is two orders of magnitude. Extending the domain from a 10 km box to a 1000 km box adds an additional four orders of magnitude. So, combining these together provides a roughly six order of magnitude increase in anticipated data generation.

Many-Core and/or GPU Readiness: *Future systems will contain "lightweight" cores and/or hardware accelerators (e.g., GPUs) with deepening memory hierarchies. Are your codes ready for this? If yes, please explain your strategy for exploiting these technologies. If not, what are your plans for dealing with such systems and what do you need to help you successfully transition to them?*

Atmospheric codes are far from "ready" for these changing compute architectures. Everybody knows this is a problem, and everybody hopes somebody will develop codes that will be better. However, no agency has committed enough money to develop (and sustain) replacement models. Plus, the architecture has been changing sufficiently fast that it has often been viewed as not worth one's time to attempt to modify codes because the next generation of lightweight cores requires rewriting the code again. So, many wait for the computing architecture to be stable enough to make efforts worth it to

redesign the general-use codes. Most atmospheric scientists take the approach of trying to ignore the problem, but at some point this will no longer be possible.

Workflows: *Please briefly describe your current workflows and requirements for 2020 and 2025.*

Based on past experience, the compute cycles used grows substantially over time. One million core hours was considered extravagant for most research five years ago. Today, this is a drop in the bucket for many types of atmospheric research. Projects now commonly request multiples of millions. It is not uncommon for single simulations to use 100,000 core-hours. Based on this growth, common requests in ten years could easily be for increments of 10’s of millions of hours, with some annual requests approaching 100’s of millions of hours. While not always a one-to-one correspondence, the estimated six orders of magnitude in model output cited above is also a rough estimate of the increased need for computational time, with the caveat that one also needs to account for the increased number of timesteps required when using 100 m versus 1 km grid spacing. This could potentially add an additional order of magnitude to the overall allocation needed bringing the total to seven orders of magnitude more computing time needed compared to present requests for atmospheric modeling. The limiting factor will ultimately be the usability of hours if that many are granted. If users cannot easily do simulations with fast turnaround time, they will most likely stick to smaller simulations so they can make progress in their work. Over allocation is a perpetual problem limiting use of large models since they wait in the queue too long to be useful.

OPTIONAL

3. What top three computing ecosystem aspects will accelerate or impede progress to meeting the challenge described above in the next 5-10 years? Why? Possible topics include the following.

- Application codes (implementation, development, portability, etc.)*
- Models and algorithms*
- Hardware resources (at all scales) including I/O, memory, etc.*
- Data workflow (including sharing, transmitting, archiving, etc.)*
- Visualization and analysis resources*
- Internal/external libraries/frameworks*
- Workforce development*
- Other*

| Accelerate | Why? |
|---|--|
| 1. Algorithms are needed to get around the limitation imposed by small timesteps in LES models. | The anticipation is that future compute cores will use slower clock speeds than today, which compounds the difficulty of working with small model timesteps. As model resolution increases, timesteps will decrease, leading to longer time-to-solution. |
| 2. Parallel I/O libraries | Increased core counts in models is leading to increased difficulty in optimally handling I/O to disk. High-performing, generalized parallel I/O |

| | |
|--|---|
| | libraries would greatly help in this area. This is especially true when specialized I/O capabilities exist on machines, such as the Burst Buffer on Cori. |
| 3. Compilers that automate handling of memory hierarchies. | Memory hierarchies vary greatly between computers, e.g., GPU vs. MIC vs. traditional configurations. Typical scientists lack the knowledge to optimize their codes as they move between environments. So, compilers that can do this for them would increase efficiency greatly. It would also lead to codes that more seamlessly work on multiple architectures. |

| Impede | Why? |
|--|---|
| 1. Large datasets | Datasets on the order of many terabytes are becoming increasingly common. Most analysis tools do not make parallel processing of this data easy, which can noticeably slow research. Data growth to 10's and 100's of TB will only aggravate this difficulty. |
| 2. Over-allocated resources | Expectations are to use increasingly large amounts of compute resources. However, it is very common for DOE resources to be over allocated, with queue wait times approaching many days or longer. This encourages smaller jobs and results in more limited research. |
| 3. Lack of computer science expertise for cutting-edge computers | The quick transition of computing to many-core techniques will catch scientists off guard, with most of them not knowledgeable enough to deal with the problem. Increased interaction between computing scientists and domain scientists is needed with resources committed to bringing codes up-to-date. |

References

- Clark, A. J., W. A. Gallus, and T.-C. Chen (2007), Comparison of the diurnal precipitation cycle in convection-resolving and non-convection-resolving mesoscale models, *Mon. Wea. Rev.*, *135*, 3456–3473, doi:10.1175/MWR3467.1.
- Lebo, Z. J., and H. Morrison (2015), Effects of horizontal and vertical grid spacing on mixing in simulated squall lines and implications for convective strength and structure, *Mon. Wea. Rev.*, *143*, 4355–4375, doi:10.1175/MWR-D-15-0154.1.

- Miyamoto, Y., Y. Kajikawa, R. Yoshida, T. Yamaura, H. Yashiro, and H. Tomita (2013), Deep moist atmospheric convection in a subkilometer global simulation, *Geophys. Res. Lett.*, *40*, 4922-4926, doi:10.1002/grl.50944.
- Schalkwijk, J., H. J. J. Jonker, A. P. Siebesma, and E. Van Meijgaard (2015), Weather forecasting using GPU-based large-eddy simulations, *Bull. Amer. Meteor. Soc.*, *96*, 715–724, doi:10.1175/bams-d-14-00114.1.

Title: Integrating Exascale Streaming Analysis and HPC with Particle-Resolved DNS to Address Cloud Challenges

Yangang Liu (Environmental and Climate Sciences Department, Brookhaven National Laboratory)
Shinjae Yoo (Computer Science Initiative, Brookhaven National Laboratory)
Nicholas D'Imperio (Computer Science Initiative, Brookhaven National Laboratory)
Xiaolin Li (Applied Math and Statistics Department, Stony Brook University)
Zheng Gao (Applied Math and Statistics Department, Stony Brook University)

1. Please specify the current science drivers for your field of research

Clouds continue to pose daunting challenges to climate modeling and observations because of their turbulent multiphysics-multiscale nature. Particular knowledge gaps exist for vital processes that occur at scales smaller than typical grid sizes of large eddy simulation (LES) models (e.g., 100 m) or cloud-resolving models (CRM, e.g., 1 km), including, but not limited to, microphysics, turbulent entrainment-mixing between clouds and environmental air, and turbulence as well as their mutual interactions. These processes and the associated subgrid variability/structure are either not represented at all or are represented rudimentarily in major types of models such as climate models, CRM models, and LES models, hindering further progress on climate modeling. Addressing the challenges at the most fundamental level calls for a cross-cutting particle-resolved direct numerical simulation (DNS) model that realistically simulates ambient clouds and has direct online data mining and streaming analysis to guide physical understanding and model development and integration with observations in a timely manner. This unique model can serve as a benchmarking tool for developing/improving understanding and parameterization of cloud-related processes and multiscale variability/structures. The online advanced streaming analysis capability will also guide development of next-generation observations and model-observation integration.

2. Describe the science challenges expected to be solved in the 2020-2025 timeframe using extant computing ecosystems.

The particle-resolved DNS model will resolve the smallest turbulent eddies (~mm size), track evolution of individual cloud droplets (e.g., $10^7 \sim 10^9 \text{ m}^3$), and cover a domain about the size of the LES grid (10 m), mimicking what happens in real turbulent clouds in the most realistic way possible. This DNS, together with advanced data mining and streaming analysis, poses new challenges to exascale computational environment. Here we propose to form a multidisciplinary team from areas of cloud and climate sciences, computer science, and mathematics to address the challenges facing domain science, streaming analysis, and HPC computation.

3. Describe the science challenges that cannot be solved in the 2020-2025 timeframe using extant computing ecosystems

Clouds are inseparable from aerosols seeding cloud droplets and precipitation resulting from cloud droplet interactions; full investigation of the aerosol-cloud-precipitation system involves processes occurring at scales even larger than the current typical LES resolution (e.g., 10 m). Further extending the model domain size to cover broader scales (e.g., 100 or 1000 m) and considering aerosol and precipitation particles in addition to cloud droplets will likely pose further challenges to computing ecosystems. Another potential challenge is the integration with the increasing amount of high-resolution measurements, which will likely become available.

Significantly large model domain, more particles and streaming of ever-increasing resolution measurements pose additional challenges on modeling, data processing, and streaming analysis.

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

| Accelerate | Why? |
|-------------------------------|--|
| 1. External libraries | Our model is set up by coupling the particle dynamics and fluid fields. To accelerate the computation of particles, multiple levels of parallelism using many-core and heterogeneous architectures are required. The combination of using OpenMP, MPI, CUDA, OpenACC, and so forth will further reduce our computational time with current computational |
| 2. Streaming analysis library | Exascale DNS simulation will generate high velocities and high volumes of data that cannot be efficiently stored or analyzed. Streaming analysis will enable us to understand, predict, and guide exascale DNS simulation. |
| 3. Data workflow | Taking advantage of modern parallel I/O packages such as MPIIO will improve scalability and simplify file management. A collective I/O combined with noncontiguous accesses will yield high performance. |

| Impede | Why? |
|-----------------------|--|
| 1. Hardware resources | The speed of CPUs and memory are still the main bottleneck of DNS, whose time complexity increases as a cubic function of the Reynold number. |
| 2. External library | We currently adopt PETSc as our linear solver. However, we can only achieve good scaling up to 256 cores for the current code on our Linux cluster. Extending to larger computational domain requires better scaling on more cores and/or GPU-based environment. |
| 3. Data workflow | Data size increases quickly as the domain size is extended and more particles are considered. Building an efficient database for data query and analysis may become a challenge in the future. |

5. (Optional) Characterize the data ecosystem aspects if the primary drivers for your field of research involve the transmission, analysis (including real-time analysis), or processing of data.

The DNS poses high-velocity and high-volume simulation data stream challenges. Preliminary DNS simulation on the 1-m domain size can easily generate 400 TB of data, and the data size will likely be over 400 PB with 10-m domain simulations. The data problem will be more severe if we integrate high-resolution observational data of comparable size. To overcome such issues of high velocity, volume, and variety, we believe in-situ analysis such as streaming algorithms is mandatory to cope with the future data challenges. With streaming analytics, we can understand and predict the simulation behavior and verify it with observations, which also give us an opportunity to steer the simulation on the fly to be more meaningful. Furthermore, such streaming analytics can help us to save computational time and reduce data rates by adjusting the resolution of the simulation instead of achieving the same resolution by brute force.

Extreme Precipitation Events in Present-Day and Future Climates

David Randall

Colorado State University

1. Extreme precipitation events as a major science challenge

Understanding and predicting the climate sensitivity of extreme precipitation events is a major scientific problem that is also very important to society. There is some evidence (e.g., Lenderink, and Attema, 2015) that local precipitation maxima may increase *twice as fast as Clausius-Clapeyron scaling* as the climate warms. Many extreme precipitation events are enhanced by the effects of local topography, which can also control the fate of the precipitation after it hits the ground.

We have developed a second-generation “super-parameterization” of cloud and precipitation processes, called the Q3D MMF, which can directly simulate atmospheric deep convection, including realistic topographic effects (Jung and Arakawa, 2015), on the entire global domain, using today’s supercomputers. The Q3D MMF is ready now for implementation in global atmospheric models. It has the potential to simulate the first realistic global climatology of extreme precipitation events, and to predict how that will change in the future.

The Q3D MMF is scale-aware in the sense that it converges to a global cloud-resolving model (GCRM) when the horizontal grid spacing of the dynamical core is on the order of 4 km or less. We envision that eventually (I will not make a forecast) it will be possible to run the model as a GCRM, and this will enable even more realistic simulations.

2. Current and Future Computational Strategy

To obtain optimal computational performance in terms of wall-clock time, the Q3D MMF should be implemented in such a way that the physical processes run on a much larger set of nodes than the dynamical core. Ideally, each node should have a thousand or more cores. It will also be useful to run the physical processes concurrently with the dynamical core. It may even be advantageous to split the physical-process calculations into multiple concurrently executing sets, e.g., radiation as one set and microphysics as another.

In the GCRM limit, the model will need to run on millions of cores in order to deliver acceptable simulation speeds on the order of several simulated years per wall-clock day.

Super-parameterization and regional grid refinement can be viewed as alternative strategies to fill the gap until climate simulations with GCRMs become possible. We are betting on super-parameterization.

As atmospheric scientists, our strategy for dealing with future computing architectures is to team with computational-science professionals who can guide us through the evolving landscape.

References

- Jung, J.-H. and A. Arakawa, 2015: Simulation of subgrid orographic precipitation with an embedded 2-D cloud-resolving model. *J. Adv. Model. Earth Syst.*, **8**, 1-10. DOI: 10.1002/2015MS000539.
- Lenderink, G., and J. Attema, 2015: A simple scaling approach to produce climate scenarios of local precipitation extremes for the Netherlands. *Environmental Research Letters*, 10(8), 085001.

Atmospheric Radiation Measurement Climate Research Facility

Exascale 2016 Workshop

Chitra Sivaraman
Laura Riihimaki
Bill Gustafson
Pacific Northwest National Laboratory

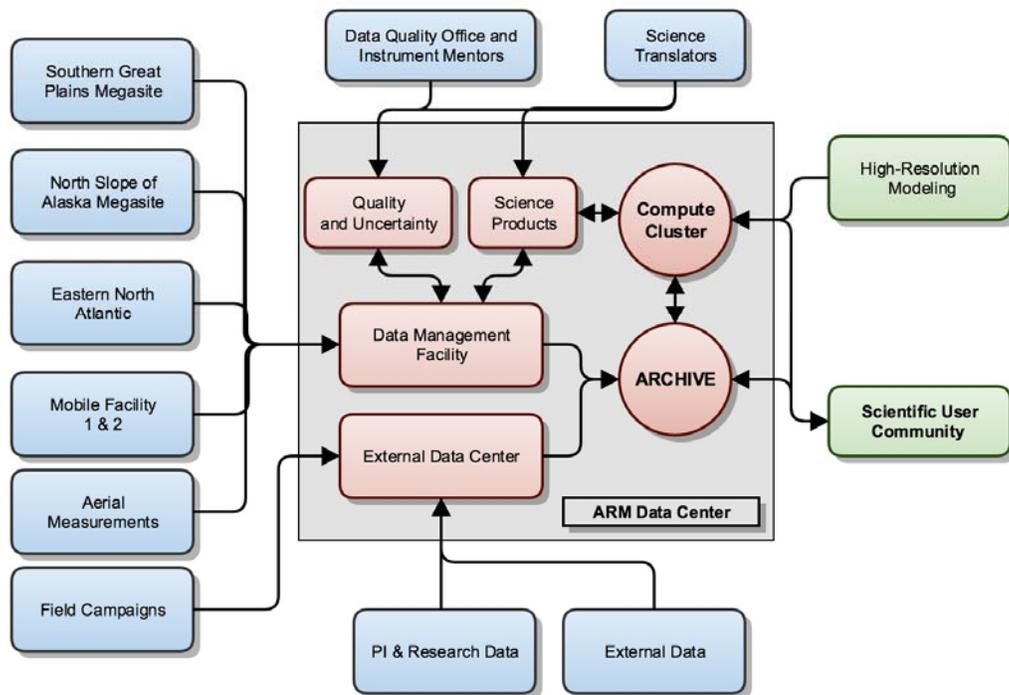
Giri Prakash
Oakridge National Laboratory

Introduction to ARM:

The Atmospheric Radiation Measurement (ARM) Climate Research Facility, a DOE scientific user facility, provides the climate research community with strategically located *in situ* and remote sensing observatories designed to improve the understanding and representation, in climate and earth system models, of clouds and aerosols as well as their interactions and coupling with the Earth's surface. ARM operates a network of surface stations including fixed sites with long-term 24-7 measurements, mobile facilities deployed for several months or years (see Figure 1.1), and an Aerial Facility with a G-1 Aircraft that makes atmospheric *in situ* measurements. Mobile and Aerial Facility campaigns are chosen through a competitive proposal process where proposed campaigns are selected based on scientific priority and practical feasibility. Ground site data comes from over 350 instruments, with another 25-30 instruments used on board the aircraft during field campaigns. All measurements and derived Value Added Products (VAPs) are stored in the ARM Archive and are available for download from ARM's webpage (www.arm.gov). ARM also interacts closely with scientists funded by the Office of Biological and Environmental Research (BER) and other science users to identify facility priorities for new instrumentation and VAP development.



Figure 1.1: Map of fixed ARM sites, mobile facility deployments, and aerial facility campaign.



160304 Jimmy Voyles and Giri Palanisamy

Figure 1.2: Diagram showing sources and stages of data processing from measurement to user download.

Figure 1.2 shows the movement of data from initial measurements to the ARM data archive where it can be downloaded by the user community. Currently about 20 TB of data is archived each month from the over 350 instruments and there are about 12 million files from 1600 different data streams in the ARM archive. A number of new instruments or upgraded instruments were added to the facility during the Recovery Act which provide new information on some key gaps in our measurements of cloud microphysical processes and aerosol composition. However, many of these instruments (scanning radars, high spectral resolution radiometers, and Aerosol Observing System measurements) produce very large and complex data streams. Thus many of these data streams remain unmined resources, and are quickly becoming the largest fraction of our data by volume. Many of these data streams are now being used only for case studies for a few days per year because they must be manually interpreted by experts for data quality and meaning.

Over the next few years, we expect that the annual observational data rate will be 10 times what is currently being archived annually. So by this estimate, and without further data product development, perhaps 90% of data volume will soon fall into the category of only accessible for small case studies by experts. Additionally, even some well- defined algorithms to retrieve useful information from these data streams are testing the limits of our current method of processing data. For example, optimal estimation retrieval methods to retrieve atmospheric humidity and cloud liquid water paths from infrared and microwave radiances call radiative transfer codes iteratively. To handle the processing, we often subsample and only retrieve 1-10% of the time steps, but if we parallelize this code and ran it in a high performance computing environment we could run all time steps which will give better cloud statistics.

2.Current and Future Computational and Data Strategies:

Observations can only get us so far in fully understanding and quantifying atmospheric processes because some key quantities can't be measured, or can't be measured at an optimal spatial scale. ARM is now beginning a new approach where high resolution Large Eddy Scale (LES) models forced by observational data will be run routinely in order to fill in some of those gaps. In order to accomplish this, additional instrumentation has been added to the Southern Great Plains site in Oklahoma, and a team has been selected to do a 2-year pilot project to develop the infrastructure needed for routine modeling. The pilot project will focus on modeling cases of shallow cumulus clouds. Several scientific motivations for this focus are the impact of shallow cumulus on global climate model temperature bias in that region, the importance of land surface heterogeneity on cloud properties, and the ability to develop statistics for parameterizations that incorporate higher order terms of the covariance between multiple parameters. Once developed for this application, it is anticipated that the modeling framework will be applied to other more complex scientific problems, possibly including deep convection at the SGP as well as phenomena at other ARM sites.

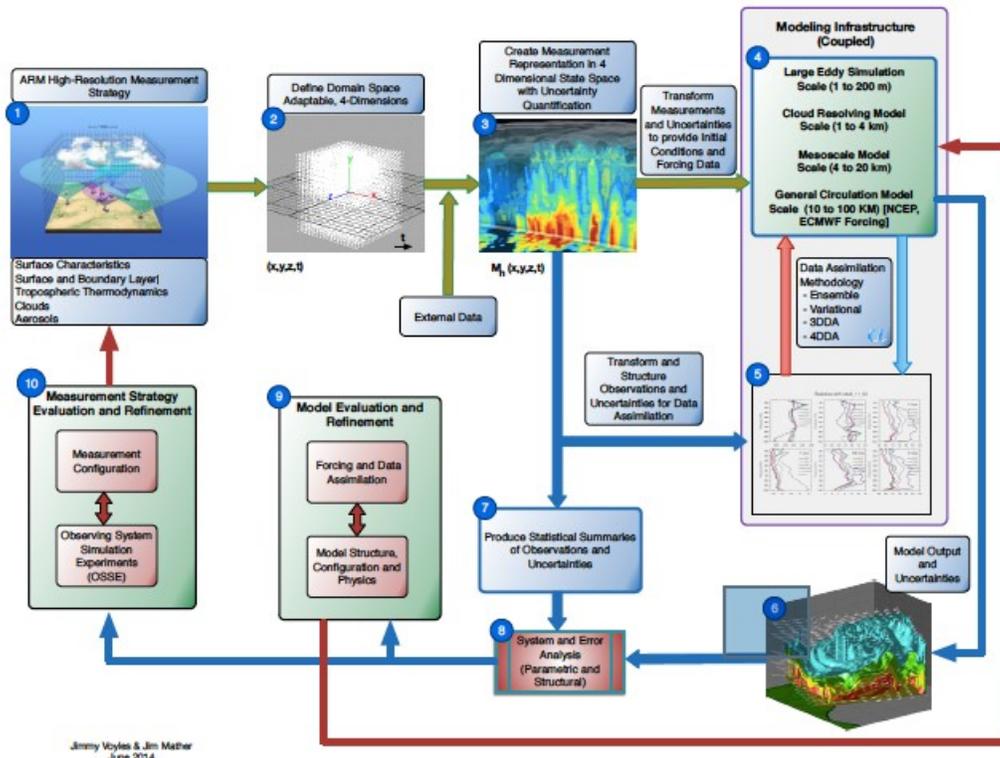


Figure 1.3: Diagram showing steps in new vision for more integration between measurement and model data.

This next generation paradigm for ARM that more closely links observations and modeling is illustrated in Figure 1.3. Box 1 shows the higher density of observational measurements at the site, along with the new transformations needed to put the data into a format for forcing the model (boxes 2-3). Observations from instruments will be assimilated (box 5) to provide forcing for the higher resolution runs (box 4). The LES output will then be further compared to observations in an interactive fashion (boxes 6-8), to refine the model (box 9) and the measurement strategy (box 10). This integrated approach will be used to build a “4-D data cube” that can give a more complete picture of the atmospheric state over time. Incorporating atmospheric modeling into ARM’s observational strategy will increase the computing needs. The anticipated data rate of the LES models is over 1 PB per year. LES modeling will also significantly increase the

computational processing requirements for the program. One of the goals of the pilot project is to examine the costs and benefits of how much model output to store and what model settings (e.g. resolution) to use to run the model. Current estimates of the LES domain to be used are a 25 km square with 100 m grid spacing. An ensemble of roughly a half dozen simulations would be run for each shallow cloud event, with different forcing conditions used for each ensemble member. Turn-around time is approximately 1.6 times the simulated time when using 512 cores on the NERSC Edison computer with the Weather Research and Forecasting (WRF) model and a bulk microphysics representation of the clouds. Additionally, plans include running one simulation per case with a spectral bin representation of the clouds, which is approximately a magnitude more expensive for a given simulation. Based on reasonable assumptions, the cost of doing this would be approximately 17–20 million core hours per year. Using a more generalized LES configuration with nested boundary conditions to more accurately capture spatiotemporally varying boundary conditions would raise the cost to 65–80 million core hours per year.

One of the challenges, however, is to accommodate the scientific needs of a variety of users in these decisions. Initially, beta users will be identified that fit into ARM and other BER program's scientific priorities for the LES model output. But in order to support a wider array of users, creative solutions will need to be found. For example, we will need to store the required forcing data sets and model settings in a way that users can rerun the model for specific output they need. In addition, we will need model metrics, statistical summaries, and indices that allow users to find the subset of data of interest to them. This either requires predefining data so that a smaller subset of data can give needed information, or alternative ways to query large datasets without downloading it. A further challenge is developing the needed tools to integrate observational and model data through improved retrievals and instrument simulators. This will likely initially only be done for the subset of instruments that are expected to have the largest impact on our scientific question.

Current Operational Needs:

ARM's current biggest operational needs in terms of computational capabilities are to run radar processes and value-added products (VAPs). A VAP is a higher-order data product that may include algorithms that are applied to observations to fulfill unmet measurement needs of the program. VAPs provide derived quantities that cannot be measured directly or routinely by the program.

Two current VAPs that are compute-intense processes are the Broadband Heating

Rate Profile (BBHRP) VAP and the Atmospheric Emitted Radiance Interferometer (AERI) optimal estimation (AERIOE) VAP. BBHRP provides a structure for a comprehensive assessment of model atmospheric radiative transfer for all conditions. The BBHRP VAP includes measurements such as surface albedo and profiles of atmospheric state (temperature, humidity), gas concentrations, aerosol properties, and cloud properties from various input datasets. Approximately 1600 cores are necessary to process one year of data. Although the BBHRP VAP is itself not compute-intensive, any reprocessing effort on any of the eleven inputs can quickly escalate this to a computationally intense process. The AERIOE VAP retrieves profiles of temperature and water vapor mixing ratio, together with cloud properties for a single-layer cloud from AERI-observed infrared radiance spectrum. It would require 200K core hours to process one year of data. ARM currently operates 5 AERIs across several facilities and a processing or reprocessing effort for a 10-yr record would need 10 million core hours.

Similarly, another high-end computational need is radar processing. The Ka-band ARM zenith radar (KAZR) determines the first three Doppler moments (reflectivity, vertical velocity, and spectral width) at a range resolution of approximately 30 meters from near-ground to nearly 20 km in altitude. One year of KAZR spectra processing within 4 hours will require about 2336 CPU's. ARM operates KAZRs across the facility. Another example of a VAP that uses GPU to process data is the microsarscl VAP that extracts approximately 30 parameters quantifying key properties of Doppler spectra collected by any of ARM's zenith pointing radars - performing insect clutter detection using a neural network based algorithm operating on the Doppler spectra. The VAP currently processes 160,000 Doppler spectra per second (limited by I/O) and can process one year of data, or roughly one terabyte, in an hour. To optimize available I/O, input data are loaded from disk and prepared for GPU processing in one-day chunks by four asynchronous I/O threads running on the CPUs.

The ARM Data Center is being designed to include compute clusters that will provide a unified cyber-infrastructure for high-volume processing, analysis, synthesis, and visualization of large ARM data products that will meet the needs of the modeling and radar community. The compute cluster's hardware architecture will provide terascale (local) and petascale (deep storage) capabilities with scalable parallel processing servers. The hardware architecture will be complimented with the software tools and framework needed to support both the infrastructure and the science community.

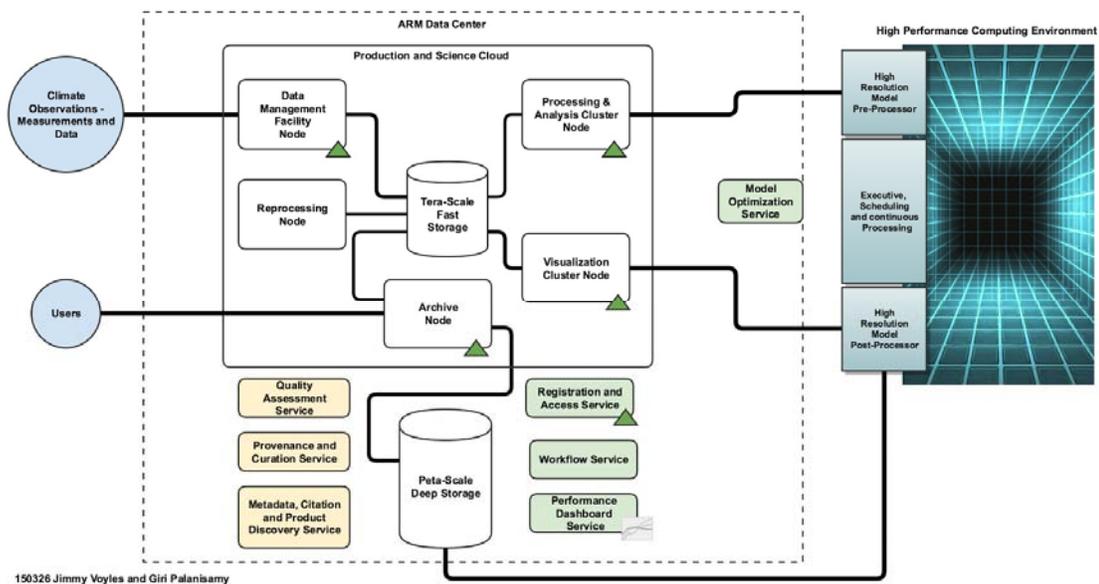


Figure 3 shows the data storage architecture of the compute cluster that will closely be integrated with high-performance parallel computing to efficiently support ARM High-Resolution modeling.

ARM is currently identifying the effort needed to migrate the ARM data process to a mid-level cluster architecture. This will then feed into the ARM next generation computing architecture. ARM is establishing a flexible adaptive architecture to meet future data management needs and challenges. This new architecture will allow individual cyber-infrastructure components to function independently but communicate with other components to support the overall data system. These individual components can be scaled independently without affecting the overall workflow. As an example, databases can be modernized to take advantage of new technologies while maintaining the communication layer between tools and applications using standards-based web-services and Application Programming Interface (API) calls.

Gaps and Challenges:

Here are a few examples of challenges that ARM faces.

- Processing and managing large data streams:
 - Reliable ways to compress data: ARM needs a reliable way to compress large data streams for storage without compromising on I/O until algorithms are developed to reduce their size without significantly reducing the needed information. As the science of

interpreting the data progresses, we may have a need to process or reprocess large data streams for new content, but the challenge of storing this data over time is significant.

- Reading and writing large data volumes in processing: When ARM process large data, like KAZR spectra, the I/O requirements are significant and sometimes prohibitive. Trivial parallelization cannot be considered for VAP processes that cannot be run individually without impacting others.
- Parallelizing data processing when sequential data is needed. Many algorithms use data from times before or after a measurement to give context to a measurement. This makes it more difficult to process algorithms in parallel.
- Monitoring and improving data quality:
 - Ways to visualize data remotely in real time to monitor the health of an instrument. This is particularly a challenge at remote deployments with low bandwidth (1.5 Mbps at Oliktok and Brazil, 512 kbps at Antarctica). Part of the difficulty is a lack of resources to create visualizations of data that will be useful for spotting problems. It is often a research effort to understand what problems can go wrong with an instrument and how to identify them. Ways that could speed up the discovery process so that simple summary plots could be decided on more quickly would help.
 - How to achieve usable data quality (calibration, field conditions, etc.) with limited human resources to manually inspect data and develop algorithms to automate this process. Even simple things like recurring shading of instruments from a tree, or an instrument that is covered in direct sunlight are now often screened manually for lack of a quick system to automate these processes. There is a need for very complex automation of data quality as well. For example, it can take 20-30 hours of a trained scientist's time to get the best usable data out of 3-hours of X-band scanning radar data (about 10 GB).
- Providing ways for diverse users to access and interpret data for different scientific needs:
 - As data volume grows, in order for users to be able to do statistical analysis of new data streams and not be confined to a few cases, we need an interface to search and subset data without having to download it. Our plan for the short term is to create pre-defined, static indices like cloud classifications, LES model metrics, etc., that allow users to find cases of interest and then

either download them or analyze them on an ARM computing cluster. But even better would be something that would allow users to interact with the data to define their own metrics on the fly.

- Computing resources to process large data sets or ways to transfer large data to institutional computing resources. Increasingly, with larger radar datasets and especially new high resolution modeling output, users will need ways to transfer large amounts of data to the computing resources where they can do analysis.

BER/CESD Exascale Facilities for Next-Generation Climate Models**(White Paper)****Shaocheng Xie (LLNL) and Wuyin Lin (BNL)**

Describe a major science challenge expected to be solved in the 2020-2025 timeframe that requires using extant computing ecosystems.

Convective- and meso-scale cloud systems are expected to be explicitly resolved in next-generation climate models with expected orders of increase in supercomputing powers in the next 5–10 years. This would significantly reduce most long-standing climate biases in both the mean climate state (e.g., double ITCZ biases) and variability (e.g., low skills in simulating diurnal precipitation and MJO signals). The interactions among convection, clouds, and aerosols are expected to be largely improved in convection-permitting climate models, which would lead to a better estimate of climate forcing and climate sensitivity in future climate change projections.

BER/CESD exascale facilities can also empower ultra-high resolution process modeling (e.g., Large-Eddy Simulation (LES) and Direct Numerical Simulation (DNS) models) across scales to aid in the improvement of 1) parameterization for remaining sub-grid processes (e.g. boundary layer turbulence and microphysics) to better couple with the resolved convection or 2) scale-aware deep convective parameterizations that can skillfully adapt to high-resolution models that might still operate at the grey-zone for deep convection over which the explicit and parameterized convection may co-exist.

Massive amounts of data will come at various stages of developments and routine simulations at increasingly convection-permitting resolutions. Large on-disk storages are necessary during the simulations and also to facilitate onsite post-processing and analysis. Even for current moderately high-resolution (e.g., $\frac{1}{4}$ degree) simulations, mandatory output can easily take up multiple terabytes of disk space. Fast networking is critical to archive the massive simulation outputs and transfer data to other distribution centers for sharing.

2. Current and Future Computational and Data Strategies

Current global climate models rely heavily on parameterizations of various physical processes. A large portion of the convective spectrum will become explicitly resolved with future global convection-permitting climate models. Data reduction procedure and big-data analytics for post-simulations will be increasingly important for future ultra-high-resolution global climate simulations.

Models:

It would require global convection-permitting models for climate simulations, and LES models for process studies that help further improve representation of remaining sub-grid processes. Investment on direct numerical simulation (DNS) modeling may also be crucial toward improving the representation of cloud microphysical processes, which will become the next weak point, to couple with explicit deep convection.

New Capabilities:

With grid resolution approaching cloud and convection scale, a non-hydrostatic dynamical core becomes a necessity. Shorter time stepping is required as a result, which will proportionally increase the computational cost. Moreover, until deep convection can be fully resolved, which requires a resolution less than 100 m, explicit and sub-grid convection may coexist in climate models of the next decade. Scale-aware parameterization for this deep convection grey-zone will be critically needed.

Resolution:

The resolution will be 1–5 km for next-generation climate models. This is at least one order higher compared to typical resolutions used currently. The ultimate goal is to be able to explicitly represent convective- and meso-scale cloud systems. Static regional refinement is needed as a cost effective complement to accelerate model developments.

I/O:

Expect to have two orders of increase in I/O. The increase is mostly due to increased grid resolution. Additional chemical species and higher order moments for process representation at high resolution also contribute to the increase in I/O. Temporal frequency is expected to be unchanged for global simulations. Data that needs to be permanently archived will be 10 times less compared to current generation models with necessary data reduction procedures adopted.

Many-Core and/or GPU Readiness:

The codes are built to run hybrid massively parallel processing and shared memory multiple-processor architecture. Some elements of the codes are ready to use accelerators, but much of the codes need to be transitioned to use GPUs. The codes have a well-established structure for profiling. Dedicated software engineering personnel are needed to analyze the computationally-intensive elements and implement the transitions.

Workflows:

The system currently uses the Common Infrastructure for Modeling the Earth (CIME) to govern the workflow from configuration, build, and job submission to the creation of checkpoints, job resubmission, and archiving. The future workflow should include data reduction and comprehensive post-processing, including data analytics and visualization. Primary computational systems and analysis clusters are usually independent machines. Federated authentication is necessary to include post-processing in the single workflow.

OPTIONAL

3. What *top three* computing ecosystem aspects will accelerate or impede progress to meeting the challenge described above in the next 5-10 years? Why? Possible topics include the following.

| Accelerate | Why? |
|---|--|
| 1. Application codes | Smooth transition with robust portability to using GPUs will ensure quick adoption of future systems for target high-resolution climate model development and testing |
| 2. Models and algorithms | Utilization of regional refinement capability and a short-range hindcast framework will accelerate the development of new features targeting globally uniform high-resolution modeling. UQ perturbation tuning is critical for finding an optimal combination of parameters. |
| 3. Visualization and analysis resources | Efficient analysis of massive simulation output will shorten the development-validation cycle. |

| Impede | Why? |
|--------------------------|--|
| 1. Application codes | Smooth transition to using GPUs is key but takes time to implement and validate |
| 2. Hardware resources | Memory capacity limits full utilization of the many-core architecture, even for current generation of high-core-density machines. |
| 3. Models and algorithms | To fully realize the potential of high-resolution, it requires to closely with parameterizations for the remaining subgrid processes, which are still challenging. |

4. Software Applications, Libraries, and Tools: *If you think your needs for HPC software (applications / libraries / tools / compilers / languages / etc) will change between now and 2020, 2025, please describe them here? Be sure to consider workflows, analytics, and I/O software.*

- i) Easy-to-use debugging tool under HPC environment
- ii) Effective machine-learning tools for big data analytics that can be used without deep knowledge of machine-learning algorithms

Terrestrial and Subsurface Research (corresponds to Section 3.2.2)

Enhanced Software Productivity to Deliver BER Science at Exascale

Lois Curfman McInnes, Argonne National Laboratory, Argonne, Illinois;
J. David Moulton, Los Alamos National Laboratory, Los Alamos, New Mexico;
David Bernholdt, Oak Ridge National Laboratory, Oak Ridge, Tennessee;
Michael Heroux, Sandia National Laboratories, Albuquerque; New Mexico
and Hans Johansen, Lawrence Berkeley National Laboratory, Berkeley, California

1. Science challenges in BER terrestrial simulations

The exascale era holds the potential to support the significant advances in multi-scale/multi-physics modeling and simulation that are necessary to enable the robust predictive understanding of complex systems sought by BER. These advances must include significant improvements in multi-scale modeling, bridging mechanistic and parameterized models across disparate spatial and temporal scales with rigorous one-way and two-way approaches. Similarly, in these simulations the number of coupled processes, as well as their complexity, must be able to grow to support integration with rapidly increasing amounts of experimental and observational data. However, these advances in modeling and simulation must be made despite the disruptive changes to both hardware architectures and programming models. In this setting, it is increasingly clear that interdisciplinary teams are needed to successfully overcome these challenges [Moulton Steefel et al. 2015].

A variety of important projects in BER highlight these challenges. For example, work to develop a predictive understanding of climate impacts on the Upper Colorado River System involves new complex biogeochemistry models that are informed by, and model, genomic information critical to carbon and nutrient cycling. These genome-informed biogeochemistry models are in turn targeted at multi-scale reactive transport models that can be coupled to high-resolution integrated hydrology (surface/subsurface flow) models of subcatchments or the entire watershed. A key feature in these systems is stream meanders in the floodplain, where hyporheic exchange may play a disproportionate role in the carbon and nutrient cycle. Similarly, understanding and predicting hydrologic response to changing climate and forest cover requires explicit consideration of groundwater–land surface–atmosphere feedbacks. New work in this area is exploring the potential of hyper-resolution, integrated hydrology models at regional and continental scales, including their coupling to both land process models and atmospheric models. An important role for these high-fidelity mechanistic models is to advance the scientific understanding of critical processes and their couplings to enable the development of models that are efficient and predictive at a particular scale of interest. This type of model reduction work is being pursued for microtopography-resolving models of thawing permafrost in the Arctic tundra. Here process-driven, simplified, intermediate-scale models are being developed that are capable of capturing the influence of tightly coupled fine-scale thermal hydrology on regional scales that are relevant to climate.

As motivated by these modeling and simulation needs, the IDEAS scientific software productivity project (<http://www.ideas-productivity.org>) is beginning to address these challenges in a BER/ASCR partnership among domain scientists, applied mathematicians, and computer scientists. Our approach is to adapt software productivity methodologies and software engineering principles for the unique needs of extreme-scale computational science, in order to help manage complexity in an era of

disruptive architectural changes. We are working to increase software development productivity—a key aspect of overall scientific productivity—thereby facilitating multi-institutional collaboration via reusable software and helping to ensure that simulations are robust, efficient, scalable, and portable across emerging architectures [Johansen McInnes et al. 2014].

2. Current and future computational and data strategies

As explained in detail in complementary white papers [Hammond and Johnson 2016; Maxwell, Johansen, and Moulton 2016; Scheibe 2016; Steefel et al., 2016], a common need for next-generation subsurface simulations is high fidelity, including multiple levels of parallelism, which in turn requires scalable algebraic solvers. Also required is better model/data integration (e.g., to model the integrated hydrology and carbon cycle). Because simulation capabilities exist for individual subsurface processes, but not all in the same application, another important need is refactoring of existing code and attention to interface issues to enable interoperability of components developed by different groups. Likewise, multi-scale and multi-physics coupling, for a spectrum of weak through strong coupling of physical phenomena, is increasingly important and requires attention to issues in software interfaces, modeling, algorithms, and coupling technologies [Keyes McInnes Woodward et al. 2013].

Application readiness for emerging extreme-scale architectures, including many-core and GPU functionalities, is essential and demands attention to low-level data structures throughout applications and libraries. These cross-cutting needs in subsurface simulations all require the use of good software design and refactoring of current codes in order to address the layers of complexity in next-generation modeling, algorithms, analysis, and performance.

3. Critical aspects of the computing ecosystem

Rapid, efficient production of high-quality, sustainable exascale applications is best accomplished using a rich ecosystem of state-of-the-art reusable libraries, tools, lightweight frameworks, and defined software methodologies, developed by a community of scientists who are striving to identify, adapt, and adopt best practices in software engineering. A key aspect of work in the IDEAS scientific software productivity project is development of an extreme-scale software development kit (xSDK), the vision of which is to provide infrastructure for and interoperability of a collection of related and complementary software elements—developed by diverse, independent teams throughout the high-performance-computing community—that provide the building blocks, tools, models, processes, and related artifacts for rapid and efficient development of high-quality applications.

Collaborative community work to develop a rich ecosystem for exascale BER simulations, possibly considering the prototype xSDK as a foundation, has the potential to dramatically accelerate progress in meeting next-generation simulation challenges. The use of best practices in scientific software engineering and productivity (e.g., object-oriented design, thorough testing, repository usage, attention to performance portability, etc.) will help scientific application codes to manage complexity and achieve good performance, whether the application software is intended for use in a single context or in modest reuse across applications in the same domain (e.g., as domain components). When pursued in a community context with emphasis on layers of software interoperability, such work enables effective reuse of internal/external libraries and frameworks, as well as the incorporation of new models and algorithms over time. For example, draft xSDK package compliance standards, including standard configure and CMake options, help to address challenges in interoperability and sustainability of software developed by diverse groups at different institutions. This work provides a foundation for addressing deeper levels of interoperability that involve exchanging, controlling, and interpreting data between packages.

4. Software applications, libraries, and tools

As the subsurface simulation community advances to address next-generation challenges, the need for reusable, extensible, and interoperable applications and libraries only increases. For example, the BER use cases introduced above already rely on solver technologies in the numerical libraries hypre, PETSc, SuperLU, and Trilinos for scalable solver technologies. These libraries provide easy access to sophisticated mathematical algorithms and high-performance data structures, so that application users do not have to write this complex code and can instead focus on their areas of primary interest. As multi-physics simulations increasingly integrate capabilities from multiple subsurface applications, the combined use of independently developed libraries becomes imperative, and this is being addressed in the xSDK. Over the long term, tools are needed to enable applications and libraries to work at higher level abstractions, in order to reduce the labor-intensive process of modifying and extending existing code in order to port to new platforms, support new data types, and improve performance. Moreover, such capabilities must fit within the broader ecosystem, requiring work flow to address the entire cycle of computational science [Johansen McInnes et al. 2014].

5. HPC services

The development and use of scientific software is a team sport, in which the team is typically geographically distributed and multiple HPC facilities may be used. The ability to work effectively across geographical and institutional boundaries, as well as across the diverse background and experience of the team members, is critical to success. As the sites that all team members can access, HPC facilities can play a central role in facilitating the success of such teams. For example, the IDEAS project is currently working with the ASCR computing facilities to expand their training activities to encompass best practices not only in high-performance and parallel computing but also in software development, tailored to the large, complex software systems typical in BER and other scientific areas. IDEAS is likewise working to engage the larger computational science and engineering community in gathering and sharing software development experiences for the benefit of others. The facilities, as a focal point for many such experiences, could also contribute strongly to encouraging such dialogue. Strong testing regimes are widely recognized as a key “best practice” for producing quality, reliable software. In this vein, the IDEAS project is also working with the facilities to explore how automated testing, including continuous integration-style testing, can be made easier both within individual facilities and across facilities and institutions where software teams work. This poses challenges very similar to other scientific work flows and will likely benefit from federated authentication, gateways, and other types of services.

References

H. Johansen, L. C. McInnes, D. Bernholdt, J. Carver, M. Heroux, R. Hornung, P. Jones, B. Lucas, A. Siegel, and T. Ndousse-Fetter. Software Productivity for Extreme-Scale Science, 2014. Report on DOE Workshop, Jan. 13–14, 2014.

[D. Moulton, C. Steefel, T. Scheibe, S. Painter, J. Lynch, S. Long, M. Heroux, L. C. McInnes, E. Brodie, and D. Watson.](#) Building Virtual Ecosystems: Computational Challenges for Mechanistic Modeling of Terrestrial Environments: Workshop Report, DOE/SC-0171. U.S. Department of Energy Office of Science, 2015. Report on DOE Workshop, March 26–27, 2014, Germantown, MD.

G. Hammond and J. Johnson, Exascale Computing for Subsurface Science: A Case Study Based on the PFLOTRAN Simulator, white paper submitted to BER/CESD Exascale Facilities Workshop, March 2016.

R. Maxwell, H. Johansen, and J. D. Moulton, An Exascale Community Challenge: Predictive Integrated Hydrology over the Continental U.S., white paper submitted to BER/CESD Exascale Facilities Workshop, March 2016.

T. Scheibe. Multiscale Simulation of Complex Biogeochemical Earth Systems, white paper submitted to BER/CESD Exascale Facilities Workshop, March 2016.

[C. Steefel et al.](#), Toward Exascale Simulations of Microbially-Mediated Biogeochemical Cycling from Pore to Watershed Scale, white paper submitted to BER/CESD Exascale Facilities Workshop, March 2016.

D. E. Keyes, L. C. McInnes, C. Woodward, W. Gropp, E. Myra, M. Pernice, J. Bell, J. Brown, A. Clo, J. Connors, E. Constantinescu, D. Estep, K. Evans, C. Farhat, A. Hakim, G. Hammond, G. Hansen, J. Hill, T. Isaac, X. Jiao, K. Jordan, D. Kaushik, E. Kaxiras, A. Koniges, K. Lee, A. Lott, Q. Lu, J. Magerlein, R. Maxwell, M. McCourt, M. Mehl, R. Pawlowski, A. P. Randles, D. Reynolds, B. Rivière, U. Rude, T. Scheibe, J. Shadid, B. Sheehan, M. Shephard, A. Siegel, B. Smith, X. Tang, C. Wilson, and B. Wohlmuth. Multiphysics simulations: Challenges and opportunities. *International Journal of High Performance Computing Applications*, 27(1):4–83, Feb 2013. Special issue.

Exascale Computing for Subsurface Science: A Case Study Based on the PFLOTRAN Simulator

Glenn Hammond, Sandia National Laboratories, Albuquerque, New Mexico and

Jeff Johnson, Lawrence Berkeley National Laboratory, Berkeley, California

1. Science challenges

Modeling at high-fidelity earth-system processes:

- Fundamental subsurface biogeochemical processes
- Impact of climate change on terrestrial ecosystem processes
- Fate and transport of contaminants at DOE facilities, legacy waste sites, and nuclear waste repositories
- Terrestrial storage of supercritical CO₂.

The following are areas in which major advances are expected in the next 5–10 years:

- Full 3D representation of geologic formations at 10–100x the level of detail
- Tighter coupling of multi-scale and multi-physics processes
- Maturation of data assimilation methodologies (that leverage subsurface simulators as engines to run ensemble calculations for coupled process models).

2. Current and future computational and data strategies

The current approach to modeling subsurface processes within subsurface simulators is to couple these processes in either a fully coupled or sequentially coupled approach, or a hybrid of the two. For instance, the code PFLOTRAN (Hammond et al., 2014) solves for multi-phase flow and reactive multi-component biogeochemical transport as separate nonlinear systems of partial differential equations (PDEs) that are coupled sequentially. However, multi-phase flow fully couples the PDEs for two fluid phases and energy, solving the resulting nonlinear system of equations using a Newton-Krylov solver. Reactive transport also solves for its PDEs simultaneously. It is not anticipated that this implicit solution approach will change significantly in the future as it is required to stably accommodate the disparate time scales involved.

3. Aspects of computing that accelerate or impede progress

Accelerators

1. *Scalable Newton-Krylov solvers.* Iterative solvers are the only option for solving large systems of linear equations and are an elementary building block for any simulation approach that integrates disparate time scales. These solvers are currently our largest bottleneck for scalability because they require global communication (Hammond et al., 2012). “Multi-level” methods such as multi-grid can significantly improve scalability, but they still need some form of global communication, and their application to system PDEs is an open area of research. Further, these solvers need preconditioners, and nonlinear preconditioning is currently very problem-specific (“dark art/sorcery”), and needs either more systematic analysis or to be separated from the critical path.
2. *Workforce development.* Teams comprising researchers and software developers that trust each other and work well together would greatly accelerate the development of next-generation models. The cohesion of the project team would require long-term funding for stability, which can be counterbalanced by accountability for tangible results. Without cohesive teams of interested and appropriately selected personnel with time for focused work, enthusiasm wanes and progress is

slow. Without long-term funding, staff becomes overcommitted and multi-institutional efforts encourage the pursuit of pet projects. Without accountability, mediocrity prevails.

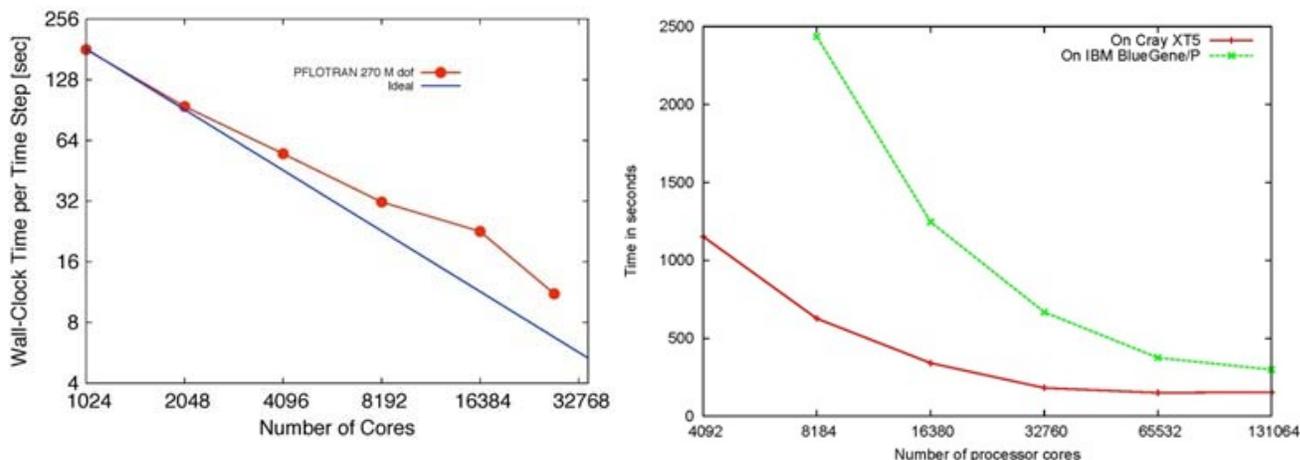
3. *Improved solution methodologies.* As posed, many multi-phase models have discontinuities that hinder the performance of scalable solvers. Focusing efforts on alternatives with nicer continuity properties would help solvers do a better job.

Impediments

1. *Current trends in hardware and programming model complexity.* Even in current-generation software, the complexity of models makes high-performance-code bases difficult to maintain. In an exascale landscape where researchers are forced to consider details such as hierarchical memory, the population of qualified HPC developers is becoming increasingly rarified and specialized. Interdisciplinary projects suffer the most from this issue.
2. *Reliance upon global communication.* At larger process counts, the cost of global communication becomes the dominant factor in performance. Communication-minimizing algorithms must be developed and deployed in applications.
3. *Lack of parallel input (the “I” in I/O).* As models become larger and more complex with large data sets being assigned to the model domain, the reading of these data sets (of increasing size) will become prohibitive.

4. Software applications, libraries, and tools

Although PFLOTRAN makes use of PETSc’s scalable solvers, the HDF5 parallel I/O library, and MPI for message passing, the largest barriers to exascale subsurface simulators are still the lack of scalable solver algorithms and I/O. This was demonstrated through a SciDAC II groundwater project in which PFLOTRAN was run on the Jaguar XT5 supercomputer at ORNL utilizing up to 262,144 cores. The performance of the code was assessed on Jaguar XT4 on up to 27,580 cores and Jaguar XT5/ANL’s IBM BGP on up to 131,064 cores (see figures below). The major bottleneck in the latter was attributed to the global reduction within the iterative Krylov solver based on profiling results (Hammond et al., 2012).



It is likely that OpenMP can be used to accelerate certain portions of the code in an expedient fashion. Through ORNL’s Center for Accelerated Application Readiness in 2011, the 0D chemical reactions within PFLOTRAN were refactored by Cray and nVidia developers to use accelerators. The results of the one-year effort can be summarized as follows: (1) for an exaggerated biogeochemical system (e.g., composed of ~30 chemical species), the maximum speedup experienced was ~14x. For a more realistic biogeochemical system (e.g., 10–15 species), the maximum speedup was < 4.

5. Summary

Based on our experience applying HPC to BER research problems, the major barriers to exascale performance are the same algorithmic barriers that curtail petascale performance, with the additional constraints introduced by accelerators and hierarchical memory. To realize scalable performance at the exascale, we believe that greater emphasis should be placed on software development, both in the study and development of scalable algorithms and in the practices of staffing, training, and motivating teams of engineers and scientists.

References

- Hammond, G.E., P.C. Lichtner, C. Lu, and R.T. Mills (2012) Chapter: PFLOTRAN: Reactive flow and transport code for use on laptops to leadership-class supercomputers, *Ebook: Groundwater Reactive Transport Models*, Editors: Zhang F., G.T. Yeh, and J.C. Parker, Bentham Science Publishers, p141–159, eISBN: 978-1-60805-306-3, doi:10.2174/978160805306311201010141.
- Hammond, G.E., P.C. Lichtner, and R.T. Mills (2014) Evaluating the Performance of Parallel Subsurface Simulators: An Illustrative Example with PFLOTRAN, *Water Resources Research*, 50, doi:10.1002/2012WR013483.

An Exascale Simulator for Scale-Aware Watershed Function Modeling

Carl Steefel, William Collins, David Trebotich, and Hans Johansen,
Lawrence Berkeley National Laboratory, Berkeley, California,
and Tim Scheibe, L. Ruby Leung, Xingyuan Chen, and Alex Tartakovsky,
Pacific Northwest National Laboratory, Richland, Washington

1. Describe a major science challenge expected to be solved in the 2020–2025 time frame that requires using extant computing ecosystems.

What will probably be solved in the next 5-10 years?

Recent research has demonstrated the use of coupled mechanistic watershed hydrologic models to project how climate-driven changes will influence watershed functioning (e.g., Engdahl and Maxwell, 2015; Troch et al., 2013; Riley and Shen, 2014). While such frontier approaches are at a very early stage of development (e.g., NRC–NAS, 2012), they offer significant potential for improving our predictive understanding of watershed function and dynamics associated with punctuated events as well as seasonal to decadal time scales. The modeling of watershed function, particularly when water flow is coupled to biogeochemical and vegetation dynamics in a complex heterogeneous environment subject to punctuated events, is an exascale computational problem. The large variations in topography, vegetation, and subsurface-rock type in mountainous watersheds suggest the use of a scale-aware modeling framework that is able to incorporate the entire watershed (typically on the order of 10–100 km²) while retaining the ability to capture finer scale gradients and fluxes within subsystems. This work is an important part of the LBNL *Watershed Function* Science Focus Area (SFA) that focuses on the Colorado River Basin. It is also an important component of PNNL's SFA that focuses on the Hanford Reach of the Columbia River, and it leverages work in the DOE–ASCR/BER IDEAS project. In a past SciDAC project, a multi-scale framework for coupling pore and continuum models for high-fidelity reactive transport simulations in subsurface has been developed. In parallel, functioning prototypes for simulating the atmospheric hydrological inputs into watersheds using global, dynamically adaptive, and cloud-system-resolving models are under active development in DOE's *Multiscale* SciDAC project. Such simulations also constitute an inherently exascale computation challenge.

Why is this important to the field?

Climate change and human activities—whose numbers and lifestyles drive an ever-increasing demand for clean water, food, and energy—are significantly reshaping interactions between vegetation, soils, subsurface, and fluvial compartments of watersheds throughout the world. These interactions occur within watersheds, which are a critically important functional unit of the Earth's surface that govern terrestrial states, stocks, and flows of water. The Colorado River watershed, for example, is the main water supply for 33 million people in Arizona, California, Colorado, New Mexico, Nevada, Utah, Wyoming, and Mexico. Watershed hydrological processes in turn mediate a wide range of biogeochemical interactions. The availability, cycling, and quality of water are intimately linked to most if not all biophysical processes that support life on Earth (NRC–NAS, 2012). In addition to the importance of watersheds to the sustainability of various Earth's systems, watersheds also provide significant feedbacks to climate (e.g., DOE, 2012). In spite of the importance of watersheds to society, there are huge uncertainties associated with how watersheds function to cycle water, nutrients, carbon, and other elements along environmental gradients. There are commensurate uncertainties stemming from large changes anticipated in the amount, phase, and seasonal cycle of precipitation inputs into watersheds as global warming accelerates during the twenty-first century.

Please give a high-level description of your research and the roles played by high-end computing, storage, and networking.

The scale-aware approach for the terrestrial component of the exascale watershed modeling problem will be based on the use of adaptive mesh refinement (AMR) as implemented in the Chombo framework and applied to terrestrial water-biogeochemical simulators like Parflow-CLM. Chombo provides a set of software libraries for solving applied partial differential equations using high-resolution numerical approaches to finite volume methods, including adaptive mesh refinement (AMR) as well as embedded boundary methods to treat complex geometries like hillslopes and river beds found in watersheds.

Scale-aware approaches should not only address numerical resolution, but also accommodate adaptation of mechanistic representations where feasible and necessary. The PNNL SFA is developing and testing scale-aware approaches for modeling surface-groundwater interactions over the Hanford Reach of the Columbia River, based on embedding high-fidelity representations of heterogeneous material properties and hydrobiogeochemical processes in selected grid elements of a lower fidelity model. Combining AMR with adaptive process models within selected grid elements not only provides increased process fidelity (and predictive power) but also opens up opportunities for increased computational concurrency for exascale computing (through task-parallel execution of subgrid models (e.g., Scheibe et al., 2015)). The high- and low-fidelity models are used in a multi-fidelity approach for data assimilation (DA) and uncertainty quantification (UQ). The multi-fidelity approach offers accurate and low-cost surrogate models, making it possible to conduct DA, UQ, sensitivity analysis, and parameter estimation in complex systems with a large number of degrees of freedom.

Similarly, the atmosphere is a new global, nonhydrostatic climate code with adaptive space-time resolution. By using both space- and time-adaptive mesh refinement, the solver allocates computational effort only where greater accuracy is needed to resolve critical phenomena. These dynamics are combined with the same physics used in superparameterized versions of DOE's Accelerated Climate Model for Energy (ACME). In superparameterization, conventional statistical representations of clouds and convection are replaced with explicit dynamical representations for greater fidelity to observed atmospheric processes. Superparameterization has been shown to improve key features of climate simulations, including the daily cycle of convective orographic precipitation and the statistics of extreme rainfall. Because the physics columns are effectively independent, they can be calculated in parallel with good scaling.

This simulator would be directly pertinent to one of the key scientific objectives for DOE's ACME project: to quantify how the hydrological cycle and water resources interact with the climate system on local to global scales. Focused simulations, together with atmospheric and land-surface process models enhanced with AMR, would be freely contributed to and further developed through active engagement with the domain and computational scientists in ACME.

2. Current and future computational and data strategies

Models

For the terrestrial component of the simulator, we propose to use the application codes Chombo-Crunch (LBNL) and PFLOTRAN (PNNL). Chombo-Crunch is a pore-to-continuum flow and reactive transport simulator based on the Chombo adaptive framework. It has been validated by pore-scale reactive transport experiments and has performed first-ever computations of resolved flow in shale. For the atmospheric component, we propose to use a variant of Chombo optimized to solve the compressible Euler equations on 3D thin spherical shells, where the radial direction is treated implicitly

to eliminate time step constraints from vertical acoustic waves. The atmosphere utilizes the same moist physics parameterizations as the System for Atmospheric Modeling (SAM). SAM is a large-eddy simulation code that solves an anelastic system of momentum equations, employs multi-moment bulk microphysics, and uses the DOE-developed Rapid Radiative Transfer Model (RRTM) radiation package.

Chombo algorithms are based on adaptive, embedded boundary methods that permit high resolution of landscape to obtain accurate fluxes into and out of the river systems. The block structured adaptive framework allows for multi-scale models to be used at differing resolutions in the watershed modeling framework and in the overlying atmospheric multiscale model. BioCrunch (based on the CrunchFlow software and under development in the LBNL SFA) will provide the coupling to microbial community dynamics through this interface.

Block-structured AMR techniques and the Chombo framework have been supported through the ASCR Applied Mathematics Program, the SciDAC Applied PDE Center at LBNL, and the SciDAC FASTMath Institute. Chombo has been developed and thoroughly tested through its application to a wide range of physics problems, including tokamak and astrophysical magnetohydrodynamics (MHD), cosmological and space plasma computational fluid dynamics (CFD), hydrological modeling, microscale fluids, and supernova explosions. Members of the white paper team have been instrumental in applying Chombo to the challenge of modeling the dynamics of the Antarctic land-ice sheet in a warmer climate.

New capabilities

These models are designed to maximize parallel scaling, by using implicit-explicit coupling of physics, and avoiding global implicit solvers, which can create extensive communication overhead. Block-structured refinement improves effective resolution for a given computational resource (such as refining over steep orography or to capture extreme events), but can create challenges for dynamic load balancing. Isolating watersheds.

Resolution

We will model the East River Colorado Watershed, a 100-km² area, with local spatial resolution of 1 meter (and finer using AMR) for the terrestrial system. We propose full treatment of the coupled biogeochemical system. We will resolve a coupled surface-subsurface-soil hydrologic system along with dynamic vegetation, all conducted with high time resolution to capture hydrologic transients like storms and other extreme events. A second objective will be to model the 60-km-long Columbia River reach adjacent to the Hanford Reservation in Washington to capture hyporheic exchange with groundwater, including the effect of daily to monthly river stage on flow and carbon cycling. A multi-fidelity approach will be used to build surrogate models for sensitivity analysis, parameter estimation, and uncertainty quantification.

Atmospheric hydrological inputs into these watersheds will be simulated using a combination of static and adaptive refinement superimposed on our global AMR model operating at resolutions commensurate with the transition from hydrostatic to nonhydrostatic dynamics at O (10 km). Static refinement to ≤ 1 km will be introduced over the target watersheds and the orographic terrain surrounding these watersheds to resolve the snow packs and rain “shadows” on the windward and leeward sides of the adjacent mountain ranges. Adaptive refinement to comparable resolutions will be introduced to track incoming synoptic-scale storm systems, for example, atmospheric rivers incident

from the Pacific Ocean, and to follow the formation and evolution of mesoscale convective complexes (MCCs) during the summer monsoon season.

I/O

Chombo utilizes parallel data I/O (HDF5), while VisIt provides visualization and data analytics. Files are stored on HPSS systems, currently totaling 1–2 PBs of offline storage while maintaining 100 TBs online. File sizes will obviously scale with problem size at the exascale. Data analysis will be pushed up into the simulation for exascale machines (e.g., asynchronous off node communication to a Burst Buffer). Further research into optimizing Chombo's parallel HDF5 I/O for in situ analytics and indexing to improve performance of tracking space-time features (like Atmospheric Rivers) will be necessary for assessing the climate impact on watershed hydrology and terrestrial ecosystems.

Many-core and/or GPU readiness

Chombo-based application codes currently run on 2 petaflop machines, up to 256K processor cores on DOE supercomputers. We expect that a Chombo-based watershed simulation capability will scale similarly because the root driver is a native variable coefficient elliptic solver in Chombo for Richards equation (subsurface) and a hyperbolic solver for the surface water equations. The atmospheric model is designed and proven to be both highly accurate and highly scalable to $>10^5$ processing elements. We are adapting and optimizing our codes for deployment onto NERSC's Cori Phase 1 and 2 through the NESAP process, and anticipate being ready to run, with demonstrated strong scaling, on O (100 Ks) of processing elements when Cori Phase 2 is operational in late 2016, taking advantage of SIMD and multi-threaded performance. Given the architectural similarities between Cori and the planned Aurora machine at ALCF, we should be in good position to effectively utilize Aurora once it is operational in 2018.

While Chombo has traditionally been parallelized exclusively using MPI, we are now actively incorporating OpenMP into it to create a hybrid code capable of exploiting the increased efficiencies from many-core parallelism and shared memory on a single die. We have started to implement cache- blocking and memory/stream management techniques that will be required because of reduced dependency on a compiler's ability to optimize memory management.

At present, we have not initiated an effort to port to GPU accelerators. Challenges with GPU programming models (and coprocessor performance in coupled simulations) continue to discourage us from writing GPU-specific codes. As less disruptive coding standards (such as OpenACC and OpenMP) continue to evolve, we will explore what parts of the simulation can be offloaded to GPU and how to implement better latency/communication hiding for those platforms.

Work flows

Current work flows are primarily large-scale file management for initialization and analysis. For initialization, these are typical "spinup" processes from coarser resolution inputs, where resolution is added dynamically. Post-simulation analysis consists of HDF5 file management, indexing, and reduction to generate time series and ensemble statistics across resolutions. This shares many of the file I/O challenges of climate applications at scale, but with the additional challenge of dynamic data structures from AMR. In the future, the ability to index and reorganize work flow analysis computations will greatly reduce the amount of file I/O at exascale.

3. What top three computing ecosystem aspects will accelerate or impede progress to meeting the challenge described above in the next 5–10 years? Why?

| Accelerate | Why? |
|--|--|
| 1. Programming models for optimal many-core SIMD performance | Managing memory, SIMD, and threading for any given platform requires platform-specific performance expertise, and time-consuming trial-and-error optimization. |
| 2. Asynchronous, communication-hiding libraries and hardware support | Latency, communication optimization for dynamic load balancing, and I/O performance |

| Impede | Why? |
|--|--|
| 1. Emphasis on flat MPI, static load balancing | Does not reflect the reality of unbalanced systems, models, and workloads. |
| 2. Memory management | Explicit memory management is problem dependent; performance can vary greatly. |

References

Scheibe, T. D., K. Schuchardt, K. Agarwal, J. Chase, X. Yang, B. J. Palmer, A. M. Tartakovsky, T. Elsethagen, and G. Redden, “Hybrid multiscale simulation of a mixing-controlled reaction,” *Advances in Water Resources*, 83: 228–239, doi:10.1016/j.advwatres.2015.06.006, 2015.

Oceans and Cryospheric Research (corresponds to Section 3.2.3)

BER/CESD Exascale Facilities Workshop White Paper/Case Study: Robustly and Probabilistically Assessing Ice Sheet Stability Due to Anthropogenic Activity

Jeremy Fyke, Los Alamos National Laboratory

The Challenge: Ice sheets are (by and large) “slow-responding” components of the climate system, in the sense that ice-sheet-scale response times to imposed perturbations range from 10^2 to 10^5 years (Price et al., 2011; Robinson et al., 2012). Critically, this timeframe mirrors the expected length of the Earth system response to anthropogenic carbon emissions (IPCC 2013). For example, the majority of the CO_2 concentration perturbation due to anthropogenic emissions will persist for order 10^3 years after the cessation of emissions due to the timescales of oceanic carbon sequestration and carbonate dissolution; atmospheric/oceanic temperature perturbations will last even longer, for $>10^4$ years (Archer 2005). The challenge in the next 5–10 years will be to provide model-based estimates of ice sheet stability based on simulations of the coupled ice-sheet/carbon/climate system over 10^3 – 10^4 years so that (for example) robust thresholds on cumulative carbon emissions beyond which ice sheet loss is “guaranteed” can be identified. This challenge has obvious societal ramifications: long-term societal, technical, and policy trends that result in exceedance of cumulative carbon emissions consistent with, for example, Greenland Ice Sheet loss (Fyke and Matthews, 2015) will commit 10^1 – 10^2 future generations to inexorably rising ocean levels with millennial persistence (Winkelmann et al., 2015; Clark et al., 2016).

Assessments of cumulative carbon-based ice sheet stability thresholds out of necessity require large ensembles of millennial-scale full glacial-to-deglacial simulations to identify cumulative emissions consistent with ice sheet loss. Furthermore, large ensembles are required to identify uncertainty in these thresholds related to climate sensitivity and polar amplification (Fyke et al., 2014a) and internal variability (Fyke et al., 2014b,c). These assessments will clearly be very computationally demanding and could thus benefit greatly from exascale computing capabilities.

Current research is approaching this challenge via use of single-core, low-resolution Earth system models of intermediate complexity (EMIC) to simulate coupled ice-sheet/carbon/climate dynamics (e.g., Fyke et al., 2011). Pertinent initial results of these simulations are distilled policy-relevant statements such as “it is *likely* that an irreversible Greenland Ice Sheet threshold will be passed due to anthropogenic emissions, and that threshold passing is most likely to occur early in the 22st century.” Such statements beg confirmation and expansion using more robust Earth system models run on exascale architectures.

Current and Future Computational and Data Strategies: To arrive at the above types of statements, 10^2 EMIC simulations are typically performed, which generate terabytes of output data. The natural progression over the next 5–10 years in the context of simulations exploring ice sheet stability thresholds will be to maintain the scientifically required simulation length/ensemble sizes, and apply exascale computing resources to a movement up the model hierarchy from EMIC-grade to CMIP5-grade model grid resolutions and explicit resolution of physical mechanisms. This transfer will result in greatly increased demands on computing power and storage systems.

New Capabilities/Models: Given the process coupling necessary for model-based assessments of ice sheet stability thresholds, simulations must be performed using fully coupled ice-sheet/carbon/climate models forced with carbon emissions. Full coupling is essential in order to resolve pertinent Earth system feedbacks. Such process-rich coupled model frameworks are just emerging from development.

Resolution: Given the required processes, timescales, ensemble sizes and spin-up requirements (e.g., Fyke et al., 2014d) simulations quantifying ice sheet stability thresholds will always be performed at spatial resolutions significantly lower than that of cutting-edge high-resolution models (e.g., ACME).

As noted above, increases in computational resources will be directed toward closing this gap, while still maintaining sufficient simulation length and ensemble size. Regional grid refinement capabilities around (for example) the Antarctic coast could greatly aid in the compromise between resolution and run length/ensemble size in (for example) Antarctic experiments. However, the computational advantage of regional refinement in the context of operational simulations can be limited by the (current) lack of spatially varying time steps, adaptive grid refinement, and scale-aware parameterizations.

I/O: I/O requirements for ensembles of millennial-scale coupled ice-sheet/carbon/climate simulations are large, but these requirements are likely to scale in a predictable manner that reflects simulated processes, run length, and ensemble size. One avenue to reduce data volume is to reduce output frequency during the majority of long simulations; another is generating in situ analyses. However, both approaches are potentially hindered by the need to make a priori guesses as to what the important climate diagnostics will be and when significant coupled ice-sheet/carbon/climate processes (e.g., ice sheet dome separation and related freshwater surges, Gregoire et al., 2012) will emerge.

Many-Core and/or GPU Readiness: Since much of the computational expense of ice sheet stability quantification studies is taken up by large, embarrassingly parallel ensembles of long simulations, there is significant potential to adopt architectures that are amenable to exascale application of basic single program, multiple data (SPMD) techniques. The most scale-efficient way to exploit exascale computing resources in the context of identifying ice sheet stability thresholds will be to advance the ability to efficiently perform large numbers of side-by-side Earth system model simulations with standard core counts instead of developing fewer individual simulations with ultra-high core/GPU counts.

Top three computing ecosystem aspects to accelerate or impede progress

| Accelerate | Why? |
|---|-----------------------------------|
| 1. increased core counts | more ensembles/longer simulations |
| 2. HPC climate analysis | quicker scientific discovery |
| 3. consolidated super-computing resources | less workflow fragmentation |

| Impede | Why? |
|--|--|
| 1. fundamental re-codes of climate models | slow scientific progress with existing models (but is not a bad thing in its own right!) |
| 2. continued priority for single ultra-high-core number of simulations on HPC machines | multiple low-core-count (~1000 CPU) ensemble-based experiments |
| 3. lack of low-resolution coupled model support | software engineering tasks left to ill-equipped scientists, not competent software engineers |

Software Applications, Libraries, and Tools: Based on continued adoption of Fortran as the primary language for operational climate models, it seems likely that over the timeframe described above, ice sheet stability simulations will continue to use this language and associated compilers/libraries. Regarding analysis: tools that presume a particular scientific interpretation and provide fixed sets of plots (e.g., diagnostic packages, which are critical for validation purposes) are likely to not be particularly useful for final coupled ice-sheet/carbon/climate assessments because interesting results will almost certainly arise from unpredictable cross-component analyses. There will therefore likely be a need for increasing need for easily accessible, low-level Python/R utilities for parallelizing researcher-specific custom data analyses across PB-scale datasets consisting of 10⁵–10⁷ output files. -_-

HPC Services: Guidance from HPC experts on best practices for exploiting exascale-level computational resources to run and analyze large ensembles of relatively low-core-count, long-wall clock-time simulations will be critical.

References

Archer, D., 2005. *Fate of fossil fuel CO₂ in geologic time*. Journal of Geophysical Research, doi: 10.1029/2004JC002625.

Clark, P.U., J.D. Shakun, S.A. Marcott, A.C. Mix, M. Eby, S. Kulp, A. Levermann, G.A. Milne, P.L. Pfister, B.D. Santer, D.P. Schrag, S. Solomon, T.F. Stocker, B.H. Strauss, A.J. Weaver, R. Winkelmann, D. Archer, E. Bard, A. Goldner, K. Lambeck, R.T. Pierrehumbert, and G.-K. Plattner, 2016. *Consequences of twenty-first-century policy for multi-millennial climate and sea-level change*. Nature Climate Change, doi: 10.1038/nclimate2923.

Fyke, J., and H.D. Matthews, 2015. *A probabilistic analysis of cumulative carbon emissions and long-term planetary warming*. Environmental Research Letters, doi: 10.1088/1748-9326/10/11/115007.

Fyke, J.G., A.J. Weaver, D. Pollard, M. Eby, L. Carter, and A. Mackintosh, 2011. *A new coupled ice sheet/climate model: description and sensitivity to model physics under Eemian, Last Glacial Maximum, late Holocene and modern climate conditions*. Geoscientific Model Development, doi: 10.5194/gmd-4-117-2011.

Fyke, J.G., M. Eby, A. Mackintosh, and A. Weaver, 2014a. *Impact of climate sensitivity and polar amplification on projections of Greenland Ice Sheet loss*. Climate Dynamics, doi: 10.1007/s00382-014-2050-7.

Fyke, J.G., M. Vizcaíno, W. Lipscomb, and S. Price, 2014b. *Future climate warming increases Greenland ice sheet surface mass balance variability*. Geophysical Research Letters, doi:10.1002/2013GL058172.

Fyke, J.G., M. Vizcaíno, and W.H. Lipscomb, 2014c. *The pattern of anthropogenic signal emergence in Greenland Ice Sheet surface mass balance*. Geophysical Research Letters, doi: 10.1002/2014GL060735.

Fyke, J.G., W.J. Sacks, and W.H. Lipscomb, 2014d. *A technique for generating consistent ice sheet initial conditions for coupled ice sheet/climate models*. Geoscientific Model Development, doi: 10.5194/gmd-7-1183-2014.

Gregoire, L.J., A.J. Payne, and P.J. Valdes, 2012. *Deglacial rapid sea level rises caused by ice-sheet saddle collapses*. Nature, doi: 10.1038/nature11257.

IPCC (Intergovernmental Panel on Climate Change), 2013. "Climate change beyond 2100, commitment, stabilization and irreversibility" Section 12.5.1 in *Working Group I Contribution to the IPCC Fifth Assessment Report Climate Change 2013: The Physical Science Basis*, September.

Price, S.F., A.J. Payne, I.M. Howat, and B.E. Smith, 2011. *Committed sea-level rise for the next century from Greenland ice sheet dynamics during the past decade*. Proceedings of the National Academy of Sciences, doi: 10.1073/pnas.1017313108.

Robinson, A., R. Calov, and A. Ganopolski, 2012. *Multistability and critical thresholds of the Greenland ice sheet*. Nature Climate Change, doi: 10.1038/NCLIMATE1449.

Winkelmann, R., A. Levermann, A. Ridgwell, and K. Caldeira, 2015. *Combustion of available fossil fuel resources sufficient to eliminate the Antarctic Ice Sheet*. Science, doi: 10.1126/sciadv.1500589.

Ice Sheet Uncertainty Quantification for Sea Level Rise

Daniel Martin (DFMartin@lbl.gov),¹ Charles Jackson,² Esmond Ng,¹
Stephen Price,³ Andrew Salinger⁴

(1) LBNL, (2) University of Texas, Austin, (3) LANL, (4) SNL

Sea level rise (SLR) due to climate change will almost certainly impact coastal regions severely in the United States and globally in the course of the next 50–100 years. Determining likely upper and lower bounds for SLR in short-, medium-, and long-term time frames is of crucial importance, both for its scientific interest and for its societal impact. Currently, the largest source of uncertainty in SLR projections is the response of the Antarctic ice sheet.

Roughly a third of currently observed sea level rise is thought to come from the continental-scale ice sheets of Antarctica and Greenland (the remainder being equally divided between thermal expansion and contributions from the thousands of small ice caps and glaciers), and this contribution (both in amount and proportion) is expected to increase as the ice sheets respond dynamically to intensifying atmospheric and oceanic forcing. Recent observations have shown that these ice sheets are losing mass at a rate of 500 km³/year and could potentially provide the dominant contribution to 21st-century sea-level rise. Ice sheets are difficult to model and the processes important to SLR are difficult to observe. Multi-fidelity models must resolve and simulate complex flows over a wide range of scales (1 km or less in places), solving large, nonlinear systems of partial differential equations (PDEs) that currently challenge the best available computational methods. Ice flow is sensitive to poorly understood and sparsely observed boundary processes, particularly at basal and lateral interfaces, and poorly understood, simply modeled (or absent) physical processes. Ice sheets are closely coupled to the atmosphere and ocean, and this coupling is only beginning to be included in ice sheet and climate models. Despite significant recent improvements, there is generally still a dearth of observational data needed for model initialization, calibration, and validation, which is a primary cause of uncertainty in model predictions.

In the BER/ASCR-sponsored PISCEES SciDAC application partnership, we continue to significantly improve DOE's ice sheet modeling capability. Improvements include: robust, accurate, and scalable dynamical cores (“dycores”) allowing for large-scale, high-resolution ice sheet modeling on both structured and unstructured meshes with adaptive refinement; evaluation of these models using new tools and datasets for verification and validation (V&V); application of new uncertainty quantification (UQ) methods, and integration of these models and tools into the Community Ice Sheet Model (CISM), the Model for Prediction Across Scales (MPAS), and large-scale earth system models (ESMs). A long-term goal is to provide useful, credible predictions, including uncertainty ranges, of future ice sheet mass loss and resulting changes in climate and sea level.

Bounding the spread in future sea-level rise requires estimation of an ice sheet initial state that accounts for uncertainties in observations, boundary conditions, model physics, and model parameters (e.g., ice rheology). The resulting probability distribution for the initial state is then propagated forward in time to determine a distribution of ice sheet mass loss and sea-level rise that reflects uncertainties in model initialization, model physics, and environmental forcing. Model initialization in particular is a challenging UQ problem since the relevant variables are fields (e.g., initial and basal boundary conditions), resulting in high-dimensional parameter vectors. It is well known that the characterization of probability distributions in high dimensions is very challenging—the “curse of dimensionality.” Our approach currently pushes the boundaries of UQ with respect to large parameter spaces. This requires the use of (costly) adjoint and Hessian-based methods.

2. Current and Future Computational and Data Strategies

To accurately model the Antarctic ice sheet with rigorous bounds on uncertainties, high-performance computing at the exascale level will be needed. The dycores we have developed will continue to be engineered to optimize performance on new high-performance computers with heterogeneous architectures, and by 2020 we aim to be simulating ensembles of high-resolution, fully coupled ice sheet

and climate evolutions with uncertainty quantification. We envision both standalone ice-sheet model simulations and model runs fully coupled (ocean-atmosphere-sea ice) to ESMs.

The actual computational problem is the solution of large nonlinear elliptic systems of PDEs. Solution strategies depend on the actual PISCEES-developed ice-sheet models being used, and they make extensive use of computational frameworks and linear and nonlinear solvers developed by ASCR researchers. For example, the BISICLES dycore is based on a finite-volume approach using structured meshes with adaptive refinement. It uses the Chombo framework for block-structured adaptive mesh refinement, relying on nonlinear and linear solvers from PETSc, with coupling to ESM's using CISM. The FELIX dycore is based on a finite-element method using unstructured meshes with variable resolution, using numerous discretization and solver libraries from the Trilinos framework, with coupling to ESMs through the Model for Prediction across Scales (MPAS) Land Ice model.

We plan to bring these solvers into the exascale era by leveraging current and future ASCR-sponsored research to develop algorithmic strategies for solving these elliptic systems of PDEs, which tend to be communication-bound, even on contemporary architectures. At the same time, the use of finer spatial resolution in the ice sheet models threatens to expose stability issues resulting from the parabolic nature of the ice-sheet temporal evolution equation; the natural solution to this stability issue (semi- or fully implicit methods), is challenging in the exascale world. Finally, storing and understanding the results of large ensembles of runs will demand improvements in distributed I/O and analysis of "big data," making use of current ASCR-sponsored research in these areas as well.

In 2015, an example of our finest-resolved coupled full-continent and southern-ocean Antarctic simulations is a 500-m finest-resolution BISICLES AMR ice sheet calculation coupled to a 0.1° ocean model, using approximately 150M CPU-hours on NERSC's Edison for a 100-year run. In 2017, we expect to be doing 125-m and 250-m resolution full-Antarctic runs (with AMR). With the current code, each of these would respectively take approximately 900M and 300M CPU-hours per 100-year run; a suite of these runs is needed to evaluate different climate scenarios and begin to quantify uncertainties in the problem. Data storage for this problem is currently about 5 TB of data for a 20-year run, plus another 4 TB of associated data storage from spin-up runs and visualization output.

Expected improvements in the BISICLES code include a transition from 2.5D to full 3D (roughly a factor of 10) and transition to more-complex physics (a factor of 2–3), which leads to roughly 20M CPU-hours per 100-year run. A suite of 100 runs would then be 2B CPU-hours. The increase is driven somewhat by larger problem size, but mostly by the need for ensembles of runs. Importantly, the need for ensembles will increase significantly as the focus shifts to more rigorous UQ methods.

The FELIX code has leveraged ASCR investments to achieve scalability in solving a steady problem of over 1 billion unknowns in a few thousand CPU-hours. Current, variable-resolution (~ 20 – 1 km) meshes for simulating Antarctic ice-sheet evolution contain $O(20M)$ unknowns. The computational resources greatly expand when the solver is then wrapped in inverse problems to match observational data (hundreds of solves), calibration under uncertainty (hundreds of unlimited ensemble sizes), and transient UQ propagation runs (hundreds of runs of hundreds of implicit time steps).

The number of unknowns per run will be roughly 25 times that at present, a likely indicator of the MPI parallelism we can hope to achieve, which means an MPI parallelism of around 25,000 MPI tasks. We can likely expect $O(10)$ times additional fine-grained parallelism via threading.

Models:

Both stand-alone ice sheet models (ISMs) and fully coupled Earth System Models (ESMs) with active ISM components will be required. The standalone ISMs are currently the Chombo-based BISICLES code and the Albany-Trilinos-based MPAS/FELIX model.

New Capabilities:

We expect new and more-complex physics to be added to both ISMs, including models to better capture the physics of subglacial hydrology and basal friction evolution, improved thermodynamics, ice shelf

fracture, damage evolution, and iceberg calving. For example, expected improvements in the BISICLES code include a transition from 2.5D to full 3D (roughly a factor of 10) and transition to more-complex physics (a factor of 2–3), which leads to roughly 20M CPU-hours per 100-year run. A suite of 100 runs would then be 2B CPU-hours. The increase is driven somewhat by larger problem size, but mostly by the need for ensembles of runs.

Resolution:

In the context of Antarctic ice sheet dynamics, existing studies (e.g., Cornford et al. 2016) have demonstrated the need for sub-kilometer resolution at grounding lines in order to fully resolve the dynamics of marine ice sheet instability. Current fixed-mesh models struggle with this requirement, often using resolutions of 4–10 km. Since such fine resolution is unnecessary in much of the quiescent interior of the continent, this strongly suggests that variable-resolution approaches are essential. Since WAIS collapse scenarios result in grounding lines sweeping much of West Antarctica, adaptive resolution approaches (like the one employed by the BISICLES model) are highly effective for this problem.

I/O:

As mentioned above, our current data needs are around 5 TB for a single 20-year POPSICLES ice-ocean coupled run, along with about 4 TB of associated storage. We expect increased (but variable in the new MPAS-ocean model) resolution to roughly double that need per run, and then we expect to run longer than 20 years (factor of 5 for a century-scale run), and then we expect to run ensembles of these to understand the uncertainties involved. Standalone ice sheet runs will require less storage per run, but we expect to do more of them, so we expect roughly similar usage needs for the uncoupled ice-sheet runs.

Many-Core and/or GPU Readiness:

Neither the BISICLES nor MPAS/FELIX models are currently fully prepared for this system evolution. Since both heavily leverage existing software frameworks (Chombo for BISICLES and Albany/Trilinos for MPAS/FELIX), we expect to be able to rely on these frameworks to contribute the lion's share of the development effort needed for our codes to continue to perform well on existing and emerging architectures. Along the same lines, we expect to heavily leverage programming models developed by and for the frameworks, rather than investing significant energy on our own.

Workflows:

Currently, the POPSICLES coupled ice-ocean model consists of two independent models (the POP2X ocean model and the BISICLES ice sheet model) that are run separately and coupled through NetCDF file I/O and a set of python scripts which manage the coupled runs. Such an approach will have major shortcomings in an exascale world, and we expect a more tightly coupled workflow in the future.

OPTIONAL

3. What top three computing ecosystem aspects will accelerate or impede progress to meeting the challenge described above in the next 5–10 years? Why? Possible topics include the following.

- Application codes (implementation, development, portability, etc.)
- Models and algorithms
- Hardware resources (at all scales) including I/O, memory, etc.
- Data workflow (including sharing, transmitting, archiving, etc.)
- Visualization and analysis resources
- Internal/external libraries/frameworks
- Workforce development
- Other

| Accelerate | Why? |
|---|--|
| 1. Hardware resources | Access to more available CPU resources will be essential |
| 2. Models and algorithms | Will be essential to accurate and efficient model runs on emerging architectures |
| 3. Internal/external libraries/frameworks | All of our codes heavily leverage existing DOE-developed frameworks |

| Impede | Why? |
|------------------------------------|--|
| 1. Hardware architecture evolution | Will require adaptation of current algorithms, codes, and strategies |
| 2. Workforce development | Finding capable people is always difficult, and the expanded skillsets needed to work on these problems in the exascale era (computationally capable glaciologists) will only magnify this |
| 3. | |

4. Software Applications, Libraries, and Tools: If you think your needs for HPC software (applications/libraries/tools/compilers/languages/etc.) will change between now and 2020, 2025, please describe them here. Be sure to consider workflows, analytics, and I/O software.

5. HPC Services: If you anticipate needing additional HPC services not provided today in 2020–2025, what are they? Possibilities might include advanced training, data analytics and visualization assistance, support servers, collaboration tools, web interfaces, federated authentication services, gateways, etc.

More process-inclusive support to support our workflows and development efforts from the LCFs would be welcome as code-base complexity increases, including things like large-capacity head nodes for compiling big codes, up-to-date compilers that support mixed languages, and special queues or a second, small machine with identical environment for regression testing.

6. Additional Needs: If there is anything else important to the success of your project that you have not mentioned, please do so here.

7. References

Cornford, S.L., D.F. Martin, V. Lee, A.J. Payne, and E.G. Ng, “Adaptive mesh refinement versus sub-grid interpolation in simulations of Antarctic ice dynamics.”, *Annals of Glaciology*, to appear.

8. (Optional) Images

Consider submitting one or two already published high-resolution images suitable for inclusion in the report. Please provide the reference(s). Submit these separately from the two-page report; they will not count against the page limit.

Sea Ice Modeling across Scales

Wieslaw Maslowski, Andrew Roberts, and Frank Giraldo, Naval Postgraduate School, Monterey, California;
Elizabeth Hunke, Los Alamos National Laboratory, Los Alamos, New Mexico;
and Michal Kopera, University of California-Santa Cruz

1. Describe a major science challenge expected to be solved in the 2020–2025 time frame that requires using extant computing ecosystems.

The Arctic is undergoing rapid and some of the most coordinated climate changes currently occurring anywhere on Earth, with Arctic sea ice cover changes exceptional in at least the last 1,400 years [1] and related surface temperature extremes unusual in at least the past 600 years [2]. Historical model reconstructions of the high north from global climate and Earth System models (GC/ESMs) are in broad agreement with these changes; however, the rate of change in the GC/ESM predictions remains outpaced by observations [3]. Reasons why models may not be able to simulate rapid environmental change in the Arctic stem from a combination of coarse model resolution, inadequate parameterizations, unrepresented processes and spatio-temporal scales, and a limited knowledge of physical and other real-world interactions.

As the climate and Earth system modeling community approaches the capability of exascale computing and beyond, the overall trend is to resolve physical processes and resulting feedbacks within and across individual model components. This trend increases the urgency of addressing the current limitations of sea ice models in representing both the sub-floe scale (i.e., 1 km or less, sub-kilometer) and multi-floe or continuum scale (i.e., 10 km or more).

Contrary to atmospheric and oceanic models, sea ice models are currently unable to universally simulate the frozen ocean from large scale (10–1,000 km) down to very small, sub-kilometer or sub-floe scale. Continuum sea ice model dynamics are designed to operate at or above the so-called multi-floe scale, that is, at resolutions of 10 km or more [5]. Mean-time finite-element models, resolving individual floes and their interactions, continue to require refinements to simulate mechanical constraints between floes [6]. Neither class of models has the capability to simulate sea ice successfully at both the sub-floe scale and the multi-floe or continuum scale across entire basins and on climatic time scales. However, development of ocean and atmospheric models using quasi-uniform grids with local refinement demand the use of multi-scale sea ice physics in models in order to soon span the sea ice floe-scale divide.

Model limitations are hindering our ability to predict the future state of Arctic sea ice and its impacts on regional and lower latitude climate. These problems are important, because a reduction of perennial sea ice cover exposes open water to direct interactions with the atmosphere, which in turn influences regional atmospheric circulation patterns and temperature profiles, especially along the seasonal marginal ice zone. Sea ice thickness variability in space and time modifies the Arctic-wide atmospheric circulation and appears to impact the troposphere–stratosphere coupling [4]. A more realistic representation of time-dependent conditions of the Arctic sea ice cover and their effect on air-sea interactions is necessary in models, and it requires coupling of the respective model components.

2. Current and future computational and data strategies

While there is currently a strong focus in sea ice modeling research on resolving sea ice at the sub-floe scale [7,8], little attention has been given to the possibility that beneath horizontal scales of about 1 km, sea ice may often be non-hydrostatic. Recent work (in preparation by Roberts et al.) indicates that there is an analytic finite non-hydrostatic horizontal length-scale limit for sea ice, and this may be the most important consideration in modeling sea ice dynamics across scales, rather than the horizontal length scale

of floes. This work is based upon a morphological approach linking small and large-scale sea ice dynamics and may facilitate multi-scale sea ice dynamics models.

In addition, finite-element or discrete granular models [6,7,8], in which interactions between particles are simulated in Lagrangian fashion, are very expensive computationally, in order to search and follow collisions between particles, which leads to highly irregular communication patterns and significantly reduces the computational scalability of the solution. Instead, we argue that when combined with highly scalable discontinuous Galerkin (DG) methods [9], the above morphological approach will be more suitable for use of the exascale computing capability in the next generation of sea ice models. In particular, DG methods could be useful for a very accurate interpolation to an eulerian mesh or adaptive mesh tracking as well as for the solution of velocity field and transport.

Models

While the above approach deals with the sea ice dynamics, the outstanding challenge that remains is to develop an appropriate thermodynamic formulation for a complete multi-scale sea ice model. However, given the outlined advancements and sufficient resources, the thermodynamic formulation is readily achievable within the next 3–5 years with a complete new sea ice model resolving sub-floe and continuum space scales to be ready for exascale production simulations in 2020–2025.

New Capabilities

A new sea ice model that can resolve spatial scales from basin-wide (e.g., the Arctic or Southern oceans) to sub-floe scale (down to order of 0 [100 m]) is expected to increase both computational cost and memory requirements. Additional costs, yet to be determined, will likely be associated with a more complex sea ice thermodynamic formulation. This will allow new capabilities for accurate representation of sea ice deformations, thickness distribution, surface/bottom roughness (for coupling momentum from atmosphere and ocean), fast ice along the coast, and sea ice in narrow fjords (e.g., around Greenland). Overall, modeling of multi-scale sea ice physics will (1) advance understanding of physical processes and feedbacks involved in polar amplification of global climate change and (2) help understand and potentially reduce uncertainty in prediction of polar climate change at seasonal and decadal scales.

Resolution

The resolution of continuum sea ice models is order of 0 (10 km), with some model configurations approaching grid cell size of a few kilometers. In order to realistically represent many of the above additional capabilities, a resolution of about 250 m could be required. While sea ice models are applicable to polar and subpolar regions only, locally refined and/or adaptive resolution might be required in some cases, such as in fjords or along the shallow coastal areas.

I/O

Overall, compared to today's models, additional I/O will be required because of higher resolution, some additional variables, and higher frequency coupling. However, the amount of output to be saved/archived will depend on the problem being addressed; for example, for process studies involving inertial oscillations hourly output would be required, while for CMIP-type analyses weekly or monthly mean data might be sufficient.

Many-core and/or GPU readiness

DG methods are highly scalable to tens of thousands of cores, and they were successfully applied to atmospheric [9] and shallow-water [10] models, including the extended capabilities of those models to nonconforming dynamically adaptive meshes. Hardware accelerators should also be applicable to such codes, but further efforts by computational experts are needed to take full advantage of such capabilities.

Work flows

The theoretical solution of the dynamical sea ice model is completed and should soon be published. Testing of this new formulation in the Regional Arctic System Model (RASIM) will follow in the next 12 months or so. An expansion of this approach with DG methods should be achievable within the next 3 years, while the development and testing of a sea ice thermodynamic component should be possible in 3–5 years. Hence by 2025 a completely new sea ice model resolving multi-scales could be in common use for production simulations.

References

- [1] Kinnard, C., Zdanowicz, C. M., Fisher, D. A., Isaksson, E., de Vernal, A., and Thompson, L. G. 2011. Reconstructed changes in Arctic sea ice over the past 1,450 years, *Nature*, 479(7374).
- [2] Tingley, M. P., and Huybers, P. 2013. Recent temperature extremes at high northern latitudes unprecedented in the past 600 years, *Nature*, 496(7444), doi:10.1038/nature11969.
- [3] Stroeve, J. C., Kattsov, V., Barrett, A., Serreze, M., Pavlova, T., Holland, M., and Meier, W. N. 2012. Trends in Arctic sea ice extent from CMIP5, CMIP3 and observations, *Geophys. Res. Lett.*, 39(16), doi:10.1029/2012GL052676.
- [4] Krinner, G., Rinke, A., Dethloff, K., and Gorodetskaya, I.V. 2010. Impact of prescribed Arctic sea ice thickness in simulations of the present and future climate. *Climate Dynamics* 35:4, 619–633.
- [5] Hunke, Elizabeth C., Lipscomb, William H., and Turner, Adrian K. 2010. Sea-ice models for climate study: retrospective and new directions. *J. Glaciol.*, V. 56, N. 200, pp. 1162–1172(11).
- [6] Hopkins, Mark A. 2004. A discrete element Lagrangian sea ice model, *Engineering Computations*, Vol. 21 Iss 2/3/4 pp. 409–421 (<http://dx.doi.org/10.1108/02644400410519857>).
- [7] Bouillon, S., and Rampal, P. 2015. Presentation of the dynamical core of neXtSIM, a new sea ice model, *Ocean Model.*, 91, 23–37, 5890, 5891, 5892, 5914.
- [8] Baohui, Li, Li Hai, Liu Yu, Wang Anliang, and Ji Shunying. 2014. A modified discrete element model for sea ice dynamics. *Acta Ocean. Sinica*, 33(1): 56–63, doi: 10.1007/s13131-014-0428-3.
- [9] Restelli, M., and Giraldo, F.X. 2009. A Conservative Semi-Implicit Discontinuous Galerkin Method for the Navier-Stokes Equations in Nonhydrostatic Mesoscale Atmospheric Modeling, *SIAM J. Sci. Comp.*, Vol. 31, 2231–2257.
- [10] Giraldo, F.X., and Restelli, M. 2009. High-Order Semi-Implicit Time-Integration of a Triangular Discontinuous Galerkin Oceanic Shallow Water Model, *Int. J. Num. Meths. Fluids*, Vol. 63, 1077–1102.

BER/CESD Exascale Facilities Workshop Case Study: Ocean Mesoscale Eddies

What will probably be solved in the next 5–10 years? The confidence in the simulations of both the long-term climate trajectory and its decadal variability is severely limited by our lack of direct simulations of ocean mesoscale eddies. These eddies are the mechanism by which anthropogenic carbon dioxide and its associated heat are moved from the atmosphere into the deep ocean. Ocean mesoscale eddies are, in many respects, the weather of the ocean. However, mesoscale eddies occur on significantly smaller scales than the atmosphere weather and, thus, are much more computationally demanding to directly simulate. These eddies have length scales between 10 and 100 km that require grid resolutions of ~ 1 km to fully resolve. Current-generation petaflop-scale machines are challenged when presented with global 10-km ocean grids. Upgrading global ocean meshes from 10 km to 1 km will require an order 1000x more computing capacity, putting such an initiative firmly into the exascale computing era.

Models: We envision global ocean models for the exascale era to be broadly similar to those we are currently using in the petascale era when we integrate a set of PDEs with accompanying parameterizations forward in time. An ensemble of possible future ocean states is obtained by conducting a set of simulations, each perturbed in initial conditions and/or parameter settings.

New Capabilities: The fidelity of global ocean models will likely be improved by moving toward high-order finite-volume numerical methods. These high-order methods also tend to be more computationally dense, making them better suited for accelerators than our current-generation low-order methods. The use of explicit-in-time methods will become increasingly untenable as we move to exascale systems. A 10x increase in resolution will require a 10x decrease in time step. Global ocean simulations with 1-km resolution will require $O(10^9)$ explicit time steps in order to complete 100 years of simulation. Yet, as a community, we continue to struggle to find implicit solvers of comparable computational efficiency.

Resolution: Multi-resolution global ocean models will be relatively mature when exascale architectures begin to come online. Multi-resolution meshes will certainly be viable options for reducing the overall computing burden, maybe by a factor of 10 or more. At present, the multi-resolution global ocean models that are most rapidly maturing use static mesh refinement. Although it is certainly possible, the existence of adaptive mesh refinement for global ocean models is uncertain for the beginning of the exascale computing era.

I/O: We expect that I/O will grow more slowly than the increase in degrees of freedom. Such restraint will require continued investment in in-situ (during simulation) processing approaches. The upside of continued growth of in-situ processing is that it exposes additional and easy-to-capture parallelism.

Many-Core and/or GPU Readiness: Global ocean models have started to migrate certain computing kernels to many-core and accelerator architectures. Significant algorithm development and restructuring will be required to make efficient use of heterogeneous computing architectures.

Workflows: We envision that a significant effort toward in-situ analysis will result in workflows that are more closely tied to actual simulation(s). This will place increased emphasis on back-end visualization and analysis platforms that are well connected to the actual exascale platform.

Computing ecosystem aspects that will accelerate or deceleration:

1. Accelerate: Hardware resources: Ocean mesoscale simulations will benefit from an increased computing capability.

2. Accelerate: Application codes: Optimization of existing ocean models will accelerate the efficient use of exascale systems.

Accelerate: Model and algorithms: The redesign of high-computing-density kernels will accelerate the throughput obtained on exascale systems.

Coupled System Integration – Earth System Models (corresponds to Section 3.2.4)

Community Earth System Model

Susan Bates, National Center for Atmospheric Research

1. Describe a major science challenge expected to be solved in the 2020–2025 time frame that requires using extant computing ecosystems.

The contribution of the CESM (Community Earth System Model), employed by the NCAR–DOE Cooperative Agreement (DOE–CA), to DOE science initiatives is model hierarchy, by running a variety of resolutions and model complexities to address compelling science problems. The DOE–CA recently submitted a 5- to 10-year program vision that pinpointed five science topics: near-term decadal climate predictability, prediction, and long-term climate projections; predicting extremes; comprehensive characterization of uncertainty using a hierarchy of models; land surface processes and biogeochemistry (BGC) feedbacks; and atmospheric model improvements. The USGCRP has four global change strategic goals: (1) advance science, (2) inform decisions, (3) sustain assessments, and (4) communicate and educate. The primary focus of the DOE–CA is on (1); however, through that process, we also inform decisions (2) and sustain assessments (3).

The challenge faced by the CESM project both in the next five years and then again later in the upcoming decade is the participation in CMIPs (Coupled Model Intercomparison Projects) as well as DOE–CA climate change simulations using the most up-to-date CESM version under various scenarios and configurations. The CESM version 2 (CESM2) is scheduled to be released to the community in December 2016 and will be the version used for CMIP6 simulations (described below). As with any newly released model version and especially the version contributing to a CMIP, the DOE–CA group will conduct numerous climate experiments in addition to those required by CMIP6. Within the 10-year vision, the DOE–CA defined 3- to 5-year actionable items including contributing simulations to many of the CMIP6-endorsed MIPs: the Decadal Climate Prediction Project (DCPP), Scenario MIP, cloud forcing MIP (CFMIP), detection and attribution MIP (DAMIP), Land Surface, Snow and Soil Moisture MIP (LS3MIP), and Land-Use MIP (LUMIP). This vision document also outlined simulations to address uncertainty quantification with large ensembles and perturbed physics ensembles; decadal climate variability; predictability and prediction using ~1,000-year-long control runs; initialized decadal hindcasts; model sensitivity experiments such as regionally forced (“pacemaker”) and single-forcing experiments; and simulations that push model resolution. The need for so many simulations necessitates access to a large computer allocation appropriate for jobs using a modest node count that run for long periods (i.e., months) as well as storage space for model output. Ensembles of these simulations, or multiple stand-alone independent simulations, are performed easily by bundling several model executables into one job requiring a large node count.

CMIPs have gained in complexity and significance since inception, now being an international, multi-model effort and the central piece to many national and international climate change assessments (i.e., IPCC). We are currently in CMIP6 (CMIP, phase 6). Using the World Climate Research Programme (WCRP) Grand Science Challenges as the motivation, the questions to be addressed by CMIP6 are: (1) How does the Earth system respond to forcing? (2) What are the origins and consequences of systematic model biases? and (3) How can we assess future climate changes given climate variability, predictability, and uncertainties in scenarios (Eyring et al., 2015)? The CMIP6 simulations to be conducted by each modeling center are (at the minimum): a set of common DECK (Diagnostic, Evaluation and Characterization of Klima) experiments and one Historical Simulation (1850 to near present). The DECK experiments include one historical AMIP (Atmospheric Model Intercomparison Project) simulation, one pre-industrial control simulation (PI Control), one abrupt quadrupling of CO₂ (4x CO₂), and one simulation forced by a

1% per year increase in CO₂ (1% CO₂). These simulations common to all modeling centers are designed to bring continuity across all CMIPs and will be examined for common climate signals. CMIP6 also includes a set of optional Model Intercomparison Projects (MIPs) specifically designed for each CMIP phase and to answer specific scientific questions. The individual MIPs will include their own set of model simulations with specific model configurations. It is expected that the CESM will participate in 18 CMIP-endorsed MIPs, 6 of which the DOE–CA will actively be involved in. The CESM project is also planning to run the model for CMIP6 with more model complexity as CESM–WACCM (Whole Atmosphere Community Climate Model) and CESM–BGC (biogeochemistry). The DECK and Historical simulations will be run with these versions, and they will be used in certain MIPs. All the CMIP6-related and additional DOE–CA simulations described here are expected to be conducted over the next 4 to 5 years.

The CESM is targeting two model resolutions for CMIP6, nominal 1-degree (low resolution) and 1/4-degree atmosphere/land coupled to a 1-degree ocean/sea ice (high resolution). The DECK and Historical Simulation will be conducted for both of these resolutions. The plan is to run these simulations, as well as all low-resolution CMIP6-endorsed MIP experiments, on the NCAR–NSF computing system. All high-resolution MIP experiments and the additional simulations proposed by the DOE–CA will be conducted on DOE machines. Currently, the CESM project plans to contribute three ensemble members for the low-resolution model and only one member for the high-resolution model in the CMIP categories that require an ensemble. To reduce uncertainty, the DOE–CA project plans to conduct additional scenarios and increase the number of ensemble members to five for the low resolution and three for the high resolution. We anticipate conducting these additional members on DOE computing resources.

The total number of simulated years planned for the CESM project for CMIP6 is 25,660, with 5,327 of those run with CESM2 low resolution, 11,970 years with CESM2–BGC low resolution, 6,583 years with CESM2–WACCM low resolution, and 1,780 years with CESM2 high resolution. The CESM2 high-resolution simulations are targeted for DOE machines; however, we do not yet have cost estimates on those machines. On the NCAR–NSF Yellowstone system, the current version of CESM high-resolution simulations uses 150K core-hours per model year for a fully coupled case and 120K core-hours for an AMIP case. Thus, using these numbers on this system, the estimated number of core-hours needed to run all high-resolution CESM2 MIP simulations is approximately 260M core-hours.

Additional simulations proposed by the DOE–CA include adding scenarios as well as ensemble members to the required CMIP6 experiments. Some of these simulations will use CESM2–WACCM, which increases the cost of the model fivefold. The following is a preliminary plan that outlines a prototypical set of simulations for the DOE–CA. We plan for two additional ensemble members for the Historical Simulation with the following configurations: low-resolution CESM2, CESM–BGC, CESM–WACCM, and high-resolution CESM2. For the Detection and Attribution MIP (DAMIP), the CESM project plans to contribute two additional members for the four required single-forcing scenarios for both the CESM2 and CESM2–WACCM configurations. The DOE–CA project would also like to add five single-forcing scenarios and have the ensembles contain five members each. This is an additional 33 simulations with the low-resolution model and 14 with CESM–WACCM. For the Scenario MIP, the DOE–CA would like to add two additional ensemble members for the four different scenarios required for both the high and low resolutions of CESM2. This results in an additional 6,455 years of simulation for the low-resolution model at a cost of 32M core-hours, 1,010 years with the high-resolution model at a cost of 152M core-hours, 330 years with CESM2–BGC for a cost of 2M core-hours, and 2,640 years with CESM2–WACCM for a cost of 66M core-hours for a total cost of 252M core-hours.

The total number of model years expected to be conducted on DOE machines is 12,215 years, which would cost 512M core-hours on Yellowstone. Using rough machine comparison factors, these simulations would have a similar cost on Edison (NERSC) and a 20-fold cost on Mira (ALCF). This is the computational need for the DOE-CA over the next 4 to 5 years, with the majority of the simulations conducted in 2017 and 2018.

Looking forward to the latter part of the upcoming decade, we anticipate a follow-on to CMIP6 with a similar list of requirements. We expect that the CESM/DOE-CA will contribute an even higher resolution model with even more complex configurations. For example, we might expect century-long or longer with an eddy-resolving ocean in the fully coupled simulations, cloud resolved atmospheric physics, a significant increase in tracers (specifically in the chemistry or biogeochemistry), and nested climate models or models with a varying grid for focus on a particular geographical region. We expect that computing architectures will have also advanced so that we expect to be in the same situation in that the CESM will use a modest to large number of nodes for long periods of time.

2. Current and future computational and data strategies

We describe our current and future plans for data management strategies, but first want to point out the data storage challenges with the types of computational jobs we run.

1. The current CESM work flow involves running a model simulation to completion before post-processing and saving data to mass storage. The future work flow implemented in the upcoming model release (CESM2) will post-process these data as the model is running but store these data locally until the simulation is complete. Both strategies require large amounts of local storage space. The storage needs are greatly exacerbated when we bundle multiple simulations together to fit into “large job” queues (described in the “HPC Services” section below).
2. It is preferable to maintain all data and publish data from the local center where the simulations ran. This requires long-term mass storage needs.

The current CESM standard procedure is to post-process output data locally and then store only needed data on mass storage systems at the center where the simulations were conducted. Post-processing may include the transposition of time slice output (one file with all variables for each model output time step) to time series files (one file containing all output time steps for each variable) as well as the standardizing of the output data for continuity across all CMIP6 models (i.e., as CMOR, Climate Model Output Rewriter, was used for CMIP5). The CESM has a new automated end-to-end work flow, to be contained in the CESM2 model version, that will greatly reduce the time to publish data and the human interventions required. This effort is being implemented by the CESM Software Engineering Group (CSEG) and NCAR’s Computational Informatics Systems Lab (CISL). This work flow involves additional capabilities within the CESM model that will create the desired time series files and potentially the standardized files as the model is running. This removes the manual post-processing step; however, the need for large amounts of storage space while a simulation is running remains.

For CMIP6 simulations, the data will be shared with the community, so they will be published to the ESGF from that location. It is preferable not to transfer data generated at other centers to NCAR HPSS; it is best to publish directly from that location. However, this creates the need for model output to remain at the local center indefinitely.

Data needs could significantly change in the next decade since we are running higher resolution and more complex models as standard configurations. Higher horizontal and vertical resolution allows us to explore smaller scale processes, which tend to occur on shorter time scales. This means that we need to

output higher frequency data. Both the BGC and WACCM versions of the model output a significantly larger amount of data than the standard CESM.

Higher resolution, higher frequency output, and more complex model versions cause our data storage needs to increase significantly. As a point of reference, the CCSM generated ~50 TB of data for CMIP3 with ~10 TB published and ~1,600 TB for CMIP5 with ~170 TB CMORized and published on ESG. We do not expect the ~20x increase in data from CMIP5 to CMIP6, but we do expect an increase. The total amount of published output data for all CMIP6 simulations, as estimated by the Working Group on Climate Modeling (WGCM) Infrastructure Panel (WIP, <https://www.earthsystemcog.org/projects/wip>), is 233TB; however, these numbers are a work in progress and do not exactly reflect CESM output. For example, CMIP6 defines a high-resolution model as using a half-degree ocean model, but the CESM uses an ocean model with 1-degree resolution.

Models

The CESM project will use CESM2, to be publicly released in December 2016, for CMIP6-contributed simulations run on DOE machines. These include a hierarchy of model configurations varying in both resolution and complexity. Two model resolutions are targeted, low (nominal 1-degree) and high (1/4-degree atm/land coupled to 1-degree ocn/sea ice) resolutions. In addition to the basic CESM2 model, the configurations of CESM2–BGC and CESM2–WACCM will be used.

For the follow-on CMIP later in the decade, we anticipate higher resolution and even more complex model versions. One known significant change in the CESM will be the addition of a new ocean model. Currently, the CESM uses the POP2 ocean model; however, the POP2 model will not carry over to the next version of the CESM.

New Capabilities

The following new capabilities for the CESM are being considered:

1. Data assimilation may be used to obtain initial conditions, specifically because the ocean takes so long (hundreds of years) to spin up at depth. This could save on computational cost in that a long spin-up run will not be needed, but data assimilation simulations are computationally expensive.
2. The Model for Prediction Across Scales (MPAS, <https://mpas-dev.github.io/>) is being evaluated for the ice, ocean, and atmosphere component models.
3. Three-dimensional (3D) coupling may become necessary as the model becomes more modularized. For example, the chemistry may be removed from the atmosphere model and run concurrently, but this requires that the chemistry data be incorporated back into the atmosphere at individual grid points (horizontal and vertical). This will allow the CESM to scale across more nodes more efficiently, but will increase the overall cost of the model.
4. The need for passing subgrid-scale information between components is increasing.
5. A pause/rewind/resume is under development. This involves components being able to save model states and read them back during run time. This capability will both reduce the cost of running CESM in a data assimilation mode and increase its fault tolerance. Currently CESM has no way of recovering from soft errors without resubmitting a job to the queue system.
6. OpenMP threading within the HOMME dynamical core is currently being redesigned. This redesign should expose approximately 10x the amount of parallelism versus the current implementation. While the resulting efficiency of the implementation will not be known until early summer, we expect that this will enable high-resolution configurations of CAM to be efficiently executed on approximately 4x the number of cores than is currently being used on Mira.

Resolution

The CESM2 low-resolution model (nominal 1 degree) grid is the same as that used in CESM1. The atmosphere is finite volume with a horizontal grid that is latitude/longitude with 288×200 points, resulting in a uniform resolution of $1.25^\circ \times 0.9^\circ$ and 26 layers in the vertical. The nominal 1-degree ocean grid uses spherical coordinates in the southern hemisphere, but in the northern hemisphere the pole is displaced into Greenland. This creates a very fine grid around Greenland where meridians converge, improving the ability to capture deepwater formation. Meridians have the most separation in the Pacific Ocean where a grid box is approximately 0.64° . The horizontal grid has 320×384 points and is uniform at 1.125° in the zonal direction but varies in the meridional direction with the finest resolution of 0.27° occurring at the equator. There are 60 levels in the vertical with layers increasing from 10-m thickness in the upper ocean to 250 m by a depth of about 3,500 m, below which it remains constant. The land model uses the same horizontal grid as the atmosphere and the sea ice uses the same horizontal grid as the ocean.

The high-resolution CESM2 model version employs a spectral element (SE) grid for the atmosphere (CAM-SE). It utilizes an explicit Runge-Kutta time-stepping approach and a continuous Galerkin spectral finite element method in the horizontal directions. The latter is described by Taylor et al. [1997] and Fournier et al. [2004]. The horizontal discretization is built upon unstructured quadrilaterals (a cubed-sphere mesh). Typically, third-order polynomials are selected that provide a fourth-order accurate horizontal discretization. These polynomials make use of 4×4 Gauss-Lobatto-Legendre (GLL) quadrature points within each spectral element.

The SE dynamical core utilizes the hybrid sigma-pressure vertical coordinate and vertical finite-difference discretization; the horizontal diffusion scheme is based on fourth-order hyper-diffusion; and a second-order dissipation provides a sponge at the model top. CAM-SE uses the vector-invariant form of the momentum equation. It conserves the dry air mass and total energy.

The physics-dynamics and tracer coupling strategy follows a hybrid paradigm. This means that the physics and dynamics packages are coupled via a process-split approach, whereas the tracer advection (e.g., for the moisture variable q) is coupled to the dynamics via time-splitting. The tracer transport scheme is built upon the same spectral element method in the horizontal direction and utilizes a remapping algorithm in the vertical direction. A positive-definite constraint is applied to ensure a positive tracer mass.

I/O

We anticipate CESM I/O requirements to grow significantly, perhaps exponentially, in the next 5 to 10 years because of the desire to increase both model resolution and the complexity of the modeled earth systems. Examples include significant increases in complexity of model chemistry, biology, and ice physics. Improving our work flow and avoiding duplicate copies of data will help, but improvements in the I/O software level are also needed.

Many-core and/or GPU readiness

The Application Scalability and Performance (ASAP) group within the Computer Information System Laboratory (CISL) is optimizing code for Xeon Phi architectures. This work, funded as an Intel Parallel Computing Center (IPCC) grant and in collaboration with NERSC through its Exascale Science Application Program (NERSC), has focused on preparing CESM for the KNL-based Cori system, which will be deployed at NERSC in June 2016. This project has developed both methodologies and tools that simplify the

optimization of code for Xeon and Xeon Phi architectures. The tools developed under this project include: the CESM Ensemble Consistency Tool (CECT) [Baker et al., 2015, CECT 2016], which simplifies the verification of non-bit-for-bit changes to source code, and the Kernel Generator (KGEN), which extracts computational kernels [Kim et al., 2016, KGEN 2016]. KGEN has enabled the production of over 30 computational kernels [KGEN kernels 2016], which are freely accessible. Several of these kernels, which represent 50% of the total cost CAM as of January 2015, have been optimized for many-core.

Specifically, the cost of the Morrison Gettelman microphysics code version 2, which represents 10% of the total cost of CAM, has been reduced by 56%. The cost of the HOMME dynamical core when configured in a generic CAM configuration has been reduced by 25%, while a configuration used for atmospheric chemistry has been reduced by 44% on Haswell and 55% on KNC. Several of these code modifications are already included in the current development version of CESM, while others are currently under consideration.

Work flows

CSEG and CISL staff are working on an end-to-end work flow that automates as much as possible and includes new parallel tools written in python. This includes the transformation of time slice history file output, standard in the current version of CESM, to time series files containing all time steps for one variable in parallel using the pyReshaper tool. Staff are currently pursuing the ability of scripts called by the model to standardize model output to conform to CMIP6 requirements (similar to CMOR, which was used for CMIP5) in parallel using the pyConform tool. There are also tools to create diagnostics as part of an overall CESM run work flow in parallel using the pyAverager tool and parallel calls to the NCL plotting routines for each model component. Our current plans include working with the scheduler on the machine and stringing jobs together through dependencies. The lightweight parallel python tools use core python 2.7.x and require external modules mpi4Py, PyNIO, and netcdf4py.

3. What top three computing ecosystem aspects will accelerate or impede progress to meeting the challenge described above in the next 5–10 years? Why?

| Accelerate | Why? |
|---|--|
| 1. Visualization and analysis resources | End-to-end work flows require that HPC centers closely link compute platforms to analysis platforms. They should share access to a common file system so that large data transfers are not required, and they should share access to the job-queuing system so that dependencies can be established that allow analysis tools to run as the data become available. We need the ability to run parallel analysis and visualization packages as part of the work flow. |
| 2. Application codes | We need code that is more modular and portable so that the addition of new science is more flexible. |
| 3. Hardware resources (I/O) | I/O capabilities need to be configured better on future architectures so they are not a bottleneck in performance. |

| Impede | Why? |
|------------------------------|--|
| 1. Hardware resources | We must be able to debug a simulation at full node count and with very fast turnaround (minutes to hours). |
| 2. Local online data storage | We need to keep all data from a running simulation local and accessible until a run is complete, at which point we post-process and archive only the necessary files. |
| 3. Long-term mass storage | We are unable to bring data generated at another computing facility to our local mass store. Thus, we publish data from the local center. We also access the data for scientific analysis and publication for years after a simulation is complete; therefore, we need the data to remain accessible for many years. |

4. Software Applications, Libraries, and Tools

At this time, we do not foresee any new HPC software needs.

5. HPC Services

We use this section as a way to provide HPC centers with feedback (from which to infer needed HPC services) on roadblocks or difficulties that we've encountered because of our method of computing. We point out that we have not encountered all these issues at every DOE computing facility and that some points below detail overall struggles of climate computing on larger and larger machines.

1. *High-performance computing*

Leadership Computing Facilities (LCFs) support scientific research projects that require unusual, large amounts of computing time, which can be defined as either capability or capacity.

Because of the large computational needs, these projects cannot be conducted anywhere other than an LCF. A capability job is one that uses a very large number nodes (thousands to tens of thousands) for short periods of time (minutes to hours), while a capacity job uses a modest number of nodes (hundreds to thousands) for very long periods of time (days to months).

Production climate models fit into the second category, using anywhere from hundreds to a few thousand nodes, depending on machine and complexity of the model. Research projects employ both of these types of jobs, and both require similar total amounts of computing time. The proposal process, though INCITE and ERCAP, is open to projects that have maximum scientific impact and that cannot be undertaken at another institution because of the amount of computing resources needed. Under this definition, both capability and capacity jobs are appropriate; therefore, these jobs should be *viewed equally by the computing center management, staff, and proposal review committees.*

2. Computing center expectations

- a. Many computing centers have a “capability” or “large-job” queue that only jobs over a certain node count can utilize. These queues come with some advantage to the user, either preference in the queue or a charging discount. The larger the node count, the faster a job gets through the queue. In effect, this favors the above-described capability jobs and penalizes the capacity jobs. On systems such as NERSC, one-simulation CESM jobs can run in a timely manner (the penalization is in the charging factor), while on systems such as Mira, it is nearly impossible for one-simulation jobs to get through the queue because of their low priority in the scheduler.
- b. To fit into computing center’s definition of a capability job, we bundle simulations together into one large job. For example, on Mira we bundle four standard high-resolution CESM simulations that run on 2K nodes into one large 8K-node job, and on Edison, we bundle nine standard low-resolution simulations into one large job. Although this is a solution that works for us both at Argonne and NERSC, it means that we need to always have multiple simulations that can be run simultaneously. The obstacles for us in doing this are that we must propose simulations that are independent and will need a large amount of computing time. First, having four independent simulations is not always feasible. Second, again using Mira as the example, if we bundle four standard high-resolution, fully coupled, twentieth century (1850–2012) simulations together, that would require ~335M core-hours, which is on the high end of allocations approved.
- c. Climate-computing campaigns rarely show “constant rate of usage of allocation” behavior. Rather, they have periods of high and low utilization, sometimes lengthy. Some computing centers expect that we’ll start running simulations the day the allocation begins and maintain a constant rate of usage.
- d. Computing centers define the end of the project as the end date of the allocation. This is unrealistic for the climate community. Centers may expect data to be removed from their systems once an allocation is complete. The end of an allocation is the beginning of the scientific analysis, which requires continued access to the data and a small amount of computing allocation. Analysis could continue for years beyond the project allocation. In addition, much of the data produced by the CESM are published to the ESG from the local computing center. The data then cannot be removed without breaking the access to that data through the ESG.

References

Baker, A.H., D.M. Hammerling, M.N. Levy, H. Xu, J.M. Dennis, B.E. Eaton, J. Edwards, C. Hannay, S.A. Mickelson, R.B. Neale, D. Nychka, J. Shollenberger, J. Tribbia, M. Vertenstein, and D. Williamson, “A new ensemble-based consistency test for the community earth system model.” *Geoscientific Model Development*, 8, pp. 2829–2840, doi:10.5194/gmd-8-2829-2015, 2015.

CESM Ensemble Consistency Test (CECT), 2016: <https://github.com/NCAR-CISL-ASAP/PyCECT/releases>.

Eyring, V., Bony, S., Meehl, G.A., Senior, C., Stevens, B., Stouffer, R.J., and Taylor, K.E., “Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organisation,” *Geosci. Model Dev. Discuss.*, 8, 10539–10583, doi:10.5194/gmdd-8-10539-2015, 2015.

Fournier, A., M.A. Taylor, and J.J. Tribbia (2004), The spectral element atmospheric model: High-resolution parallel computation and response to regional forcing, *Mon. Wea. Rev.*, 132, 726–748, doi:10.1175/1520-0493(2004)132<0726:TSEAMS.2.0.CO;2.

Kernel GENERator (KGEN), 2016: <https://github.com/NCAR/KGen>.

KGEN Kernel Repository, 2016: <https://github.com/NCAR/kernelOptimization>.

Kim, Y, J. Dennis, C. Kerr, R. Prasanna, A. Simha, A. Baker, and S. Mickelson, 2016: KGEN: A Python Tool for Automated Fortran Kernel Generation and Verification, *ICCS 2016 Workshop Tools for Program Development and Analysis in Computational Science*. San Diego (to appear).

Taylor, M., J. Tribbia, and M. Iskandarani (1997), The spectral element method for the shallow water equations on the sphere, *J. Comput. Phys.*, 130, 92–108, doi:10.1006/jcph.1996.5554.

Task Parallelism for Earth System Models

Phil Jones (LANL)

Describe a major science challenge expected to be solved in the 2020–2025 timeframe that requires using extant computing ecosystems.

Computer architectures are evolving toward very large numbers of streaming processing elements as a path toward exascale computing. Climate models on current architectures have difficulty exploiting even moderate parallelism using standard domain decomposition techniques. At the same time, climate models are becoming more complex through the inclusion of new processes, through the use of variable resolution to resolve a range of spatial scales, and in addressing multiple timescales. Task parallelism may provide an opportunity to exploit the increasing core counts while also managing the complexity of the models being deployed.

Parallel implementations of climate models generally rely on a two-dimensional decomposition of the model's horizontal spatial domain. While this is a natural decomposition that is relatively straightforward to implement, even very-high-resolution climate simulations have difficulty efficiently exploiting systems beyond 10–100k processing elements. Ideas to further exploit parallelism have tended to rely on ensembles rather than revisiting the approach to parallelism. New approaches to parallelism based on task parallelism have (re-)emerged recently, with some examples including ParSEC [1] and Legion [2]. These approaches rely on lightweight task run-time systems and dataflow analyses to efficiently map tasks and their associated data to the underlying hardware. With enough granularity, such systems may provide many advantages in exploiting exascale hardware, including far greater parallelism, load balancing, fault tolerance, mapping different tasks to different types of processing elements in hybrid systems, and the management of data. From the climate model perspective, a task-based approach can also be used to manage complexity at various levels since each new process or component is the addition of a new task. Tasks and associated data can be defined and distributed on the hardware, managing multiple space and time scales based on when results of a task are needed.

2. Current and Future Computational and Data Strategies

Refactoring climate models to utilize task parallelism will take place on multiple levels. Frameworks like Legion currently manage tasks at the higher (e.g., component or domain decomposition) levels, distributing work across nodes, managing layouts, data, and any required messaging. We can use Legion to explore task layouts and data strategies at these higher levels, in addition to adding new components like analyses or managing separate processes like I/O. At lower levels within components, we can begin to separate different tasks like various tendency terms as kernels and begin to compute these tasks concurrently on different cores or accelerators. This early refactoring will enable us to explore the efficacy of the approach and begin to define the data flow and tasks within our models. As programming models like Legion become more mature, we can rely on more sophisticated dataflow analyses to manage a broader range of tasks and exploit further advantages of the approach. Ultimately, we would view the climate model as simply a continuing list of tasks and data dependencies that can be managed and distributed to the hardware with a run-time system.

As described above, a task-based model has several advantages from both the computing and climate science perspectives. Computational advantages include:

- *Exposing additional parallelism.* As future architectures add more processing elements, we will be able to move beyond the limitations of horizontal domain decomposition.
- *Load balancing.* With a fine enough granularity of tasks, idle processing elements or threads will be able to move on to the next work unit.

- *Fault tolerance.* If a task does not complete, a replacement task can be launched on other resources to recover from the failure.
- *Mapping tasks to most suitable hardware.* For hybrid architectures, tasks that are more suitable to hardware elements (e.g., GPU accelerators, burst buffers, I/O) can be scheduled or mapped to those elements.
- *Optimizing data.* Data management is often the most limiting aspect of current simulations, so explicit data requirements and dataflow analyses will help to optimize dataflow through the complex memory hierarchy.

For climate science, advantages include:

- *Greater modularity and extensibility.* Adding new processes would involve defining the task and the data dependencies and is less dependent on the overall model framework.
- *Varying complexity.* Within a single model, sub-models of varying complexity could be deployed for different time or space domains. For example, deploying active biogeochemical (BGC) models only in selected regions or changing the complexity of those models based on other criterion. Parameterizing processes in some regions while resolving them in others or deploying an IAV model within a particular region would also be possible. Tasks would simply be launched based on data availability and other criteria and scheduled and load balanced by the run-time system.
- *Time scales.* Because tasks are started and stopped based on analyses of dataflow, tasks with different timescales could be launched and run concurrently with shorter time tasks distributed and run while waiting for data from the longer timescale tasks.
- *Simpler programming model.* A task-based programming model would rely more on the task run-time system to distribute work; developers would not need to understand details of a domain decomposition or message passing to contribute to a model.

The above description is obviously an optimistic view and considerations like numerical and physical consistency with other algorithms or parameterizations will impact our ability to define tasks completely independently. In addition, a naïve implementation may result in excessive data motion. However, a careful and staged implementation will allow us to explore the effectiveness of this approach.

References:

- [1] G. Bosilca, A. Bouteiller, A. Danalis, M. Favre, T. Herault, J. Dongarra. "PaRSEC: A programming paradigm exploiting heterogeneity for enhancing scalability," *IEEE Computing in Science and Engineering*, September 2013.
- [2] M. Bauer, S. Treichler and E. Slaughter and A. Aiken. "Legion: Expressing Locality and Independence with Logical Regions," *Proceedings of the International Conference on High-performance Computing, Networking, Storage and Analysis*, IEEE Computer Society Press, p. 66, 2012.

Exascale Computation for Data Assimilation with Earth System Model

Minghua Zhang
Stony Brook University, Stony Brook, New York

Data assimilation has been traditionally performed at operational weather forecasting centers as a way for model initialization. The value of data assimilation, however, goes well beyond that. For example, the meteorological reanalysis products from data assimilation have provided much of the analysis of climate variability and change in the last 60 years.

To date, there has been very little effort to assimilate data with Earth System Models that encompasses different model components, especially data on plants, biogeochemical cycles, soil, and hydrology. However, to develop the predictive capability of the Earth System, to quantitatively describe the past and monitor the current state of the physical Earth System, and to rapidly respond to human and natural events, data assimilation with Earth System Models is needed.

Data assimilation can be considered as an interpolation technique to blend sparse observations with model results. There has been a lot of experience in the community with assimilating heterogeneous data into coupled atmosphere–ocean–land models. The one consistent thing that has emerged from these experiences is the huge demand for data assimilation on computing resources.

In the era of data assimilation with ESM, challenges include the following:

1. Organization of data structures from different sources and processes across vastly different temporal and spatial scales with rapid throughput
2. Interfaces for models to take observations
3. Fast I/O architecture
4. Algorithms to conduct data assimilation with coupled systems of vastly different scales
5. Load balancing of assimilations on massively parallel computers.

The community will need to be engaged to identify the values of ESM data assimilation, the end users of the products, the availability and desirability of data, and the algorithms to assimilate data in coupled models as complex as ESM.

Integrated Assessment Modeling (corresponds to Section 3.2.5)

Computing Requirements for Integrated Assessment Modeling: 2020 – 2025

Kate Calvin and Robert Link (both of Pacific Northwest National Laboratory)

1 Science challenges

One of the major challenges in integrated assessment (IA) research is that there are not laws of physics for human systems. That is, there are no rules for how societal decisions are made, and the societal decisions made in the next 10–20 years could have a profound influence on the world over the next century. As a result, there is significant uncertainty both in what actions are taken and in how the physical system responds to those actions. Quantifying those uncertainties and using that quantification as a means of informing (or guiding) decisions in the future is *the* key challenge for IA modeling over the next 5–10 years.

There are a variety of subchallenges that will have to be met along the way, all of which will require computing resources far beyond what has been traditionally used in IA research. First among these is that we will need to calibrate IA models to produce reliable predictions. IA models have traditionally not emphasized historical back-testing, but many of the questions we are now asking of IA models will require some reasonable predictive power. Developing this capability will involve many Markov Chain Monte Carlo runs over historical periods, comparing to historical data.

The second computational challenge will be defining and running the parameter space search. This will use all of the Monte Carlo sampling strategies common to this sort of work; however, we face an additional complication in that we don't actually know what the most important outputs are. We can make some educated guesses, but limiting ourselves in that way risks missing unanticipated key indicators for scenario outcomes. On the other hand, saving everything, just in case, is probably not scalable. Overcoming these limitations will likely require a combination of an iterative process of discovery (which imposes computational burdens of its own) with an automated in-flight analysis of scenario results.

Beyond their uses for scenario analysis in their own right, IAMs are increasingly being coupled to other models, with higher computational costs. For example, the integrated Earth System Model (Collins, et al., 2015) couples the GCAM IAM to the CESM Earth System Model. IAMs are also being coupled to high-resolution impact models, such as hydrological models and detailed models of buildings. As those models move to exascale computing, it will be important to ensure that the IAM is not a bottleneck in the coupled system. In the current iESM, GCAM runs on a single processor, forcing hundreds to thousands of processors to await its completion. In addition, the coupled system is limited by the computational expense of the climate system, rendering large-scale uncertainty analysis difficult to infeasible. However, given the tremendous uncertainty as to how the human system will evolve in the future, uncertainty analysis in an IAM is critical.

2 Current and future computational and data strategies

Traditionally, IA modeling has not been a compute-intensive discipline, and our community has not given much thought to issues like computational performance or parallel scalability. This has begun to change over the past 5–10 years, but high-performance IA modeling is still in its infancy, and there is a lot of work to be done to get IA models ready for the coming challenges.

Narrowing our focus to the GCAM IA model, we see four primary areas of improvement: software performance, inter-model coupling, resolution, and automated feature discovery. Getting acceptable performance from GCAM will require a substantial restructuring of the code's internal layout. These efforts will allow us to improve cache performance, parallelize parts of the code that are currently serial bottlenecks, and provide for vectorization. This restructuring might also allow us to take advantage of many-core architectures, but that is speculative at this point and would depend greatly on the amount of local cache available on the many-core devices.

Going forward, IA models will be expected to support sophisticated two-way coupling with a variety of other model types, including climate models and their components, detailed sectoral models, and agent-based models. This presents a particular challenge both because of the irregular structure of IA models (i.e., the absence of a “grid”) and because of the diversity of the models on the other end of the coupling. The data produced and consumed by an agent-based migration model are very different from that used by the land component of a climate model. Addressing this challenge will require innovative ways of coding, exchanging, and translating data with a variety of structures (grid, hierarchical, etc.).

Increased resolution in IA models generally comes not in the form of higher geographical resolution, but in the form of enhanced sectoral detail. Some of this will come through coupling to more detailed sector-specific models, and some will be built into the sectoral structure of GCAM natively. Either way, the computational burden increases approximately linearly with the number of sectors; however, the complexity of the analysis of the output increases as at least $O(N^2)$, depending on the type of analysis desired. Therefore, we must devote as much attention to our analysis as to our model runs, anticipating the day when the analysis overtakes the model as the computational driver.

Finally, we come to the question of understanding our model results. As the complexity of the output grows, so do the challenges of finding key indicators for evaluating scenarios.

Moreover, the size of the individual data sets and number of scenarios contemplated for the ensembles make it unlikely that it will be feasible to store everything. Developing automated analysis tools that can digest and winnow the data while they are in flight will be a crucial enabler for achieving the goals described in the previous section.

3 Reference

Collins, W. D. et al. (2015). “The integrated Earth system model version 1: formulation and functionality.” *Geosci. Model Dev.* 8(7): 2203–2219.

BER/CESD Exascale Facilities Workshop—Simulating Urban Climate Systems

Robert Jacob (jacob@anl.gov), Andrew Jones (adjones@lbl.gov), Rao Kotamarthi (vrkotamarthi@anl.gov)

Describe a major science challenge expected to be solved in the 2020–2025 timeframe that requires using extant computing ecosystems.

With urban areas increasing in extent and more and more people living in cities, understanding the meteorology and climate of urban areas and ensuring these areas are appropriately represented in ESMs becomes crucial to understanding the energy demands and the health and welfare of most of the world's population. Because of their materials and geometry, cities do not have the same climate as the surrounding region (and micro-climates within cities can vary [Rosenzweig et al. 2011]). This is most famously seen in the urban heat island (UHI), the tendency for cities to be warmer than their surroundings, especially at night and in winter (NRC 2012). All of the human activity in a city takes place inside a layer of the atmosphere that extends from the surface to slightly above the buildings. The flow within this surface layer is dictated by the morphology of the urban environment (buildings, streets, and parks; i.e., urban canopy) and the local emissions of heat from energy use. The anthropogenic heat emissions from use of air-conditioning, traffic, and other local uses of energy are a distinguishing feature of urban flow. The urban canopy and the emitted heat introduce highly variable, localized flow conditions that are not resolvable with the current generation of atmospheric modeling tools. The coupling of this complex flow in the surface layer with the boundary layer above the urban canopy drives urban climate and the larger effects of urban regions on regional-scale climate. Thus, to model the urban system and its relation to local and regional climate requires advanced numerical modeling tools that operate at meter-scale resolution. The arrival of exascale computing will make it possible to begin simulating these spatial scales using flow and turbulence resolving methods. These new modeling systems will allow us to answer questions such as: What aspect of city features (size, building material, energy use, aspect ratio) has the most impact on its climate? How do urban areas modify the flow of moisture, cloud formation and precipitation patterns over local and regional scales? How does the UHI depend on building and road material, geometry, and the surrounding atmosphere? How will large-scale climate change interact with existing urban microclimates (see also BERAC 2015, pp. 25–27)?

Climate, weather, and LES models are already large consumers of high-end computing, storage, and networking. An LES simulation of a major metropolitan area will by itself be an exascale application requiring similar resources as global climate and limited-area weather models 5–10 years from now.

2. Current and Future Computational and Data Strategies

Models: There is a rich variety of models with heat, flow, and turbulence resolving capabilities that have been shown to scale reasonably well to many of the DOE HPC clusters using DNS, RANS, and LES methods for resolving turbulence flow scales in the atmosphere (NEK5000, PALM, Fluent, OpenFOAM, WRF-LES).

New Capabilities: Highly scalable atmospheric turbulence-resolving codes that allow for multi-physics (radiation and surface fluxes) coupled to similarly resolved models of the urban surface and subsurface will be required and do not yet exist. Mesoscale systems that can reach the necessary coupling resolution (1–3 km) already exist but less ad-hoc and more mathematically sound schemes will be required to couple the LES model with mesoscale-to-global models.

Resolution: Homogeneous regions of the atmospheric boundary layer can learn much from 100–200 m “Giga-LES” simulations (Moeng et al. 2009). The highly heterogeneous urban surface, containing parks, buildings, roads, canyons, and waterways, requires LES simulations over complex terrain with resolutions of 10–1 m. This will be uniform resolution over a 100 km² region.

I/O will increase 10 times over current models because of additional variables (moisture, chemistry). Detailed topography of cities (e.g., from LIDAR mapping) will need to be ingested. Similar to climate models, all simulation data (multivariate 2D and 3D time series) will need to be archived for additional analysis by the community.

Many-Core and/or GPU Readiness: Traditional LES codes have been adapted to co-processor architectures including complete conversion to GPUs (Schalkwijk et al. 2015). The complex topography and physics of urban regions may require more flexible solutions, but the fundamental problem is well suited for many-core or GPU approaches.

Workflows: The current workflow includes gathering initial and boundary data sets, executing the simulations, and analyzing, visualizing, distributing, and archiving the output. While this basic pattern will not change, the practice of “saving everything” will have to change, requiring a mix of in-situ analysis, re-running, and traditional archived output analysis.

3. What top three computing ecosystem aspects will accelerate or impede progress to meeting the challenge described above in the next 5–10 years? Why?

| Accelerate | Why? |
|----------------------------|--|
| 1. Hardware resources | Currently limited by cycles/memory |
| 2. Visualization resources | 3D animation essential for understanding turbulent flows |
| 3. Models and algorithms | Well-developed and strong community |

| Impede | Why? |
|------------------------------|--|
| 1. Validation data for model | Few observation data sets exist (BERAC 2015) |
| 2. Output strategy | Not clear how to balance what can be saved from exascale system vs. what must be saved |
| 3. Data sharing | Users of data extend far beyond the small number of LCF users who generate the data |

4. Software Applications, Libraries, and Tools: To the extent that resilience is a problem, we will need libraries (including fault-tolerant MPI) and OS-level capabilities to signal, detect, and recover from faults. Soft errors may need language-level constructs. We will continue to need scalable solutions for data movement, storage, and analysis. GIS and/or spatial data analytical tools for integrating detailed geographic information regarding urban landscapes in a multi-scale fashion will also be needed.

5. HPC Services: Simulations of urban areas will attract interest from many communities, requiring collaboration tools, web interfaces, and gateways to the data.

References and Figures

BERAC (Biological and Environmental Research Advisory Committee), 2015, IFL (Integrated Field Laboratories) Response, November 4. http://science.energy.gov/~media/ber/berac/pdf/Reports/BERAC_IFL_Response.pdf.

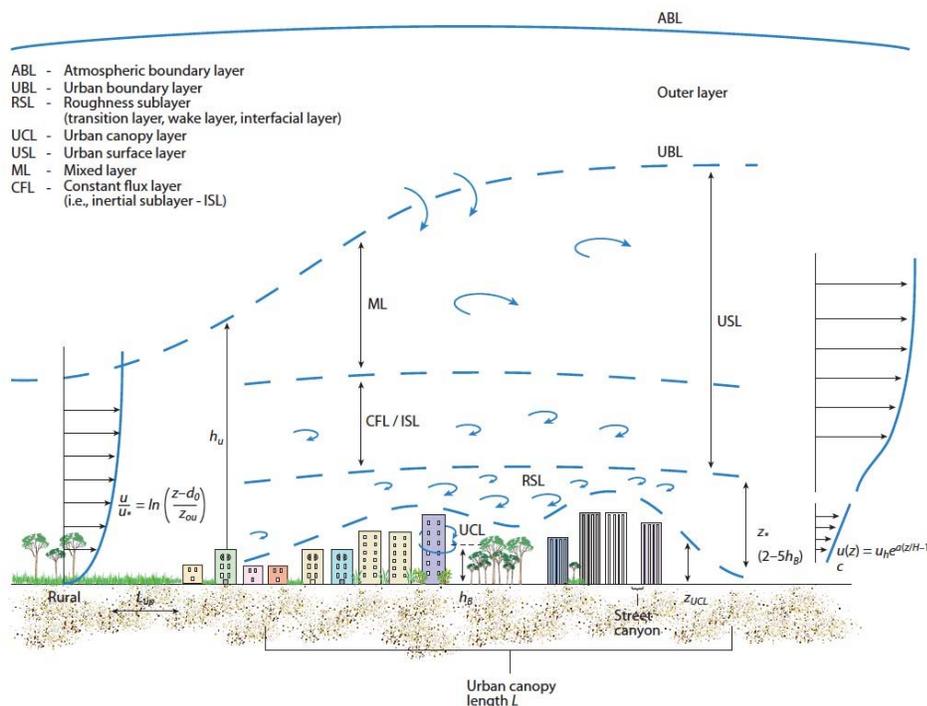
Fernando, H.J.S., 2010, "Fluid Dynamics of Urban Atmospheres in Complex Terrain," *Annu. Rev. Fluid Mech.* 42:365–89.

Moeng, C.-H., M.A. LeMone, M.F. Khairoutdinov, S.K. Krueger, P.A. Bogenschutz, and D.A. Randall, 2009, "The tropical marine boundary layer under a deep convection system: A large-eddy simulation study," *J. Adv. Model. Earth Syst.* 1:16.

NRC (National Research Council), 2012, *Urban Meteorology: Forecasting, Monitoring, and Meeting Users' Needs*. Washington, DC: The National Academies Press. doi: 10.17226/13328.

Rosenzweig, C., W.D. Solecki, S.A. Hammer, and S. Mehrotra (editors), 2011, *Climate Change and Cities*, assessment report, Urban Climate Change Research Network. Cambridge University Press, Cambridge, UK.

Schalkwijk, J., H.J.J. Jonker, A.P. Siebesma, and E. Van Meijgaard, 2015, "Weather Forecasting Using GPU-Based Large-Eddy Simulations," *Bull. Amer. Meteor. Soc.* 96: 715–723. doi: <http://dx.doi.org/10.1175/BAMS-D-14-00114.1>



Elements of the urban boundary layer (from Fernando 2010)

Multi-Scale Modeling of Integrated Urban Systems for Climate Resilience

Andrew Jones (LBNL), Robert Jacob (ANL), Haruko Wainwright (LBNL), Zenaida Mourao (U Cambridge)
Rebecca Pass (LBNL)

Cities—home to more than 50% of the world’s population and more than 80% of the U.S. population—face a number of interconnected challenges in adapting to climate change and also in mitigating its impact, including land-use planning, public health, emergency preparedness, infrastructure resiliency, and responsible management of water and energy systems. Predicting the urban microclimate is a major scientific challenge in its own right given the importance of both large-scale climate dynamics and small-scale microclimatic features driven by urban topography, built structure, boundary layer dynamics, and turbulence, as well as mesoscale phenomena such as convection, fog, and sea/lake breeze (see related whitepaper by Jacob et al.). The resulting urban microclimate drives other aspects of the urban atmosphere such as air quality and the urban hydrologic cycle (NRC Report 2012; Fiore et al. 2015).

Translating climate predictions into adaptation strategies involves additional complexity as the dynamics of exposure and response to climate variability and extremes often plays out at fine scales within the heterogeneous structure of the urban environment. For example, although urban hydrology models for flood protection and other purposes have evolved significantly in the past decade, the existing models have not been able to capture the fine-scale, three-dimensional configuration of urban topography and surface/subsurface built structures necessary for optimizing various management strategies (Fletcher et al. 2013). Likewise, targeted interventions to reduce energy and water demands and/or improve resiliency are becoming possible with the advent of models that consider detailed characteristics of individual land parcels and their placement within the three-dimensional morphology of cities (Salat 2009; Halu et al. 2016; Polebitski et al. 2011).

The interrelated nature of multiple human and natural systems within the urban environment has inspired calls for multi-scale data management and modeling platforms that can be deployed in unique configurations to address specific scientific and urban planning challenges (Catlett et al. 2016). Below, we highlight four applications of such a system relevant to BER science priorities. The requirements for each application differ, but three elements potentially requiring exascale computing platforms are present to varying degrees across the four applications: dynamics at high temporal and/or spatial resolution, process-rich interactions among physical and human systems, and a need for uncertainty quantification and multiple scenario analysis to inform decision-making.

- **Urban microclimates** – How do urban topography, built structures, and relationships with surrounding environments shape the climatic conditions experienced by urban inhabitants, and how will these microclimates evolve under changing climate and urban development?
- **Air quality** – How will climate change interact with alternative scenarios of energy system development to affect urban air quality and exposure to pollutants through changes in meteorological conditions, circulation patterns, boundary layer volume, and chemistry?
- **Urban hydrology** – How do precipitation extremes interact with complex topography and surface/subsurface structures to determine flood and associated risks (e.g., landslides) in urban environments, and how can those risks be mitigated through infrastructure upgrades and various management strategies (e.g., porous pavements, retention ponds)?
- **Energy and water demand** – How does climate variability and change interact with, building technologies, landscape characteristics, and the behavior of building inhabitants to determine the seasonality and timing of energy and water demands? What are the implications of these demand shifts under changing climate for the design of energy and water supply systems?

Current and Future Computational and Data Strategies

Models: Multiple, sector-specific urban models and data platforms exist today, but few (with the exception of atmospheric models and hydrological models) have been optimized for leadership class computing facilities. The evolving landscape of urban data streams and modeling capabilities requires that a flexible, multi-scale data platform be coupled with a highly modular, general modeling framework that can be exercised in a number of unique configurations to address specific applications. Such a coupled model framework could function in much the same way that today's Earth system models link multiple atmosphere, ocean, land, and sea ice models in unique configurations at multiple scales. Data assimilation and inverse modeling capabilities are required to take full advantage of emerging sensor networks and other "big data" streams such as social media. Data assimilation applications include pollution source identification, operational forecasting of air quality events, and tuning of energy–water demand parameterizations. Furthermore, in order to facilitate decision-support, this framework must easily support uncertainty quantification and multiple scenario analysis, as well as data visualization and integration with existing planning tools and platforms.

New Capabilities: Coupling existing models in a flexible and modular framework represents a new a new capability unto itself. Meanwhile, component models will continue to evolve. For instance, the ability to represent 3---dimensional surface and subsurface built structure in urban hydrological models is an emerging capability. Reduced---form models of particular components will permit a large number of scenarios to be explored for uncertainty quantification and decision support.

Resolution: Required model resolutions will vary depending on the application. For this reason, it is essential that the general framework accommodate multi---scale models (e.g., variable resolution atmospheric models, adaptive mesh, etc.), as well as advanced spatial statistics and scaling capabilities for integration of diverse data streams and models. The highest resolutions (tens of meters) are required for applications that assess risk and resilience for individual buildings, city blocks, or infrastructure elements. Adaptive mesh refinement and reduced order models will play an important role in deploying computational resources wisely.

I/O: At the core of the proposed urban systems model framework is a multi-scale geographic data management system that can integrate diverse data sources and translate across scales as needed. Thus, management of initial condition and boundary data sets will represent a significant I/O challenge. Very high spatial and temporal resolutions for component models imply that it will be impractical to save all of the time-evolving state variables to disk, requiring careful consideration of which data and diagnostic quantities are to be saved. One solution is to track and update probability distributions for diagnostic quantities (e.g., using Bayesian updating) as the model runs. This represents a significant departure from current practice of saving high temporal resolution state variables to disk and computing diagnostics after the completion of model runs. However, an exception to this is operational forecasting, in which time-evolving predictions of state variables are themselves the outcome of interest.

Many-Core and/or GPU Readiness: Atmospheric and hydrological models have received the most attention. Buildings, agent-based models, and data-scaling algorithms could be improved in this regard.

Workflows: Two important changes in workflows will be required: (1) real-time computation of the statistical properties of diagnostics as described above, and (2) more streamlined integration of multiple scenarios into UQ and decision-support frameworks.

References

Catlett, C. and M. Piette. (2016). Multiscale Coupled Urban Systems, DOE Exascale Applications Whitepaper.

Fiore, A. M., V. Naik, and E. M. Leibensperger. (2015). Air Quality and Climate Connections, *Journal of Air and Waste Management Association*. 65(6), 645–685, doi: 10.1080/10962247.2015.1040526.

Fletcher, T.D., H. Andrieu, P. Hamel. (2013). Understanding, management and modelling of urban hydrology and its consequences for receiving waters: A state of the art. *Advances in Water Resources*. 51, 261–279, doi: [10.1016/j.advwatres.2012.09.001](https://doi.org/10.1016/j.advwatres.2012.09.001).

Gröger, G., L. Plümer. (2012). CityGML – Interoperable semantic 3D city models. *Journal of Photogrammetry and Remote Sensing*, 71, 12–33, doi: 10.1016/j.isprsjprs.2012.04.004.

Jacob, R., A. Jones, and R. Kotamarthi. (2016). Simulating Urban Climate Systems. Companion whitepaper for the BER/ASCR Exascale Facilities Workshop, March 29–30, Washington D.C.

National Research Council. (2012). *Urban Meteorology: Forecasting, Monitoring, and Meeting Users' Needs*. Washington, D.C.: The National Academies Press. doi:10.17226/13328.

Polebitski, A.S., R.N. Palmer, M. ASCE, P. Waddell. (2011). Evaluating Water Demands under Climate Change and Transitions in the Urban Environment, *Journal of Water Resources Planning and Management*. 137(3): 249–257, doi: 10.1061/(ASCE)WR.1943-5452.0000112.

Salat, S. (2009). Energy loads, CO₂ emissions and building stocks: morphologies, typologies, energy systems and behavior. *Building Research & Information*. 37(5–6), 598–609.

Halu, A., et al. (2016). Data-driven modeling of solar-powered urban microgrids. *Science Advances*. 2(1), doi: 10.1126/sciadv.1500700.

***Transforming Science through Exascale Capabilities: Model-Data Fusion and Testbeds
(corresponds to Section 3.2.6)***

Large-Scale Heterogeneous Data Management

Large-Scale Climate Data Analytics

Forrest M. Hoffman, Jitendra Kumar, and Sarat Sreepathi, Oak Ridge National Laboratory, Oak Ridge, Tennessee; William W. Hargrove, USDA Forest Service; and Richard T. Mills, Intel Corp.

The volumes of climate and related Earth science data are rapidly growing as model resolutions increase and observing networks and satellites collect data at higher spatial and temporal resolutions. New data analytics approaches are required on high-performance computing platforms to synthesize and analyze these data. To realize the promise of new discoveries in Earth science over the next decade, an effort is required to develop and extract key analytics methods useful in climate research, to optimize these methods for existing Leadership Computing platforms using large climate data sets, and to develop benchmark problems for co-design of future platforms with data analytics in mind.

1. Describe a major science challenge expected to be solved in the 2020–2025 time frame that requires using extant computing ecosystems.

Observational and modeled data acquired or generated by the various disciplines within the realm of the Earth sciences encompass temporal scales of 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). Because of rapid technological advances in sensor development, computational capacity, and data storage density, the volume, complexity, and resolution of Earth science data are increasing at a commensurate rate. Moreover, combining, integrating, and synthesizing data across Earth science disciplines offers new opportunities for scientific discovery that are only beginning to be realized. The rise of data-intensive scientific pursuits, in Earth sciences and other disciplines, has led some visionaries to proclaim it the fourth paradigm of discovery alongside the traditional experimental, theoretical, and computational archetypes (Hey et al., 2009).

The promise of scientific advances in sustainability and environmental change research has stimulated an enormous increase in the volume of model and observational data. Earth system model (ESM) simulations are being conducted at progressively higher resolutions, generating growing volumes of data and requiring improved hardware and software infrastructure and more efficient post-processing and data-mining analysis techniques. Organized global climate modeling activities, like the Coupled Model Intercomparison Project (CMIP) that coordinates simulations in support of the United Nations' Intergovernmental Panel on Climate Change (IPCC) assessment reports, can generate tens of terabytes to several petabytes of simulation results in raw form (Overpeck et al., 2011). Observational data pose their own challenges. Satellite remote sensing data tend to be very large, and their size has grown as spatial and temporal resolutions have increased.

While great strides are being made in the quantitative assessment of model fidelity through comparison with benchmark observational data sets (Randerson et al., 2009; Luo et al., 2012) and uncertainty quantification (UQ) methods (Ricciuto et al., 2012), today's large and complex Earth science data often cannot be synthesized and analyzed using traditional methods or on individual workstations. Data-mining, machine learning, and high-performance visualization approaches are increasingly filling this void and can often be deployed only on parallel clusters or supercomputers. However, supercomputer architectures designed for compute-intensive simulations, usually containing large numbers of cores with

high-speed interconnects between nodes, are not typically optimal for large-scale analytics. Instead, such applications demand large and fast on-node memory, high-bandwidth input/output (I/O), and fast access to large local disk volumes. Most Earth scientists are ill-equipped to develop analytics codes for these architectures, while system vendors have largely focused on compute-intensive applications and must acquire representative analytics benchmarks and scientific expertise to design systems for geospatial big data analytics.

2. Current and future computational and data strategies

To realize the promise of new scientific discovery from very large, long time series Earth science data, a distinct balance of increasing computational, storage, and bandwidth capacity from high-performance computing resources is required. Traditional analysis methods and algorithms are insufficient for analyzing and synthesizing such large data sets, and those algorithms rarely scale out onto distributed-memory parallel platforms. Therefore, new analysis techniques and scalable algorithms and software tools must be developed to enable analysis, exploration, and visualization of today's Earth science data. Highly scalable parallel k-means cluster analysis ([Hoffman and Hargrove, 1999](#)) and global domain traversal using a parallel MapReduce approach ([Kendall et al., 2011](#)) are two techniques used for spatio-temporal analysis that exploit massively parallel, high-performance computing resources. These algorithms are generalized, flexible, and increasingly employed in Earth science, making them good candidates for analytics cores in a co-design process with system architects and integrators directed at large-scale climate data reduction and analysis.

[Hoffman and Hargrove \(1999\)](#) developed a parallel k-means clustering algorithm, which they implemented on an early Beowulf-style parallel cluster computer they constructed from surplus PCs ([Hargrove et al., 2001](#)), useful for segmentation, feature extraction, network analysis, change detection, model intercomparison, and model-data comparison in a number of Earth science applications ([Hoffman et al., 2008](#)). Recent improvements to that code, including adoption of a triangle-inequality-based acceleration technique and “warping” of unassigned/empty cluster centroids, have significantly reduced the time to solution ([Hoffman et al., 2008](#)), and a new technique for initial centroid determination has improved the statistical performance of the clustering result. These enhancements have enabled the analysis of large satellite data sets for identification of forest disturbances, a key component in a national-scale early warning system for detection of threats to forest health ([Hargrove et al., 2009](#); [Hoffman et al., 2010](#)). This well-instrumented code could be encapsulated and combined with similar machine learning methods in a library, then optimized on DOE's existing Leadership Computing platforms for real, large-scale climate and satellite remote sensing analytics. Moreover, benchmark tests and data sets could be designed for exercising and interrogating the performance of different potential future system designs, possessing various balances of CPUs, GPUs, memory hierarchies, I/O capacities, and volatile and nonvolatile storage options.

Development and application of data analytics methods connects mathematicians and statisticians performing algorithm design, with computer scientists implementing and optimizing algorithms for target hardware platforms, with climate scientists applying the tools to answer new questions and test hypotheses. Stand-alone and flexible implementations of these algorithms could be used by computer scientists and computer system vendors to inform the design of future Leadership Computing platforms suitable for the large-scale data analytics needs of climate researchers.

3. What top three computing ecosystem aspects will accelerate or impede progress to meeting the challenge described above in the next 5–10 years? Why?

| Accelerate | Why? |
|---|---|
| 1. Analytics methods and algorithms | Highly scalable analytics algorithms are needed. |
| 2. High-bandwidth I/O | Fewer computes on high-resolution, long time series |
| 3. Larger, faster caches | Fewer computes on high-resolution, long time series |
| Impede | Why? |
| 1. Very large core counts | Reduced memory bandwidth for data-intensive tasks |
| 2. Data distribution, provenance, and archiving | Need reliable and fast subsetting, versioning. |
| 3. Visualization resources | Need scalable systems designed around GIS tools. |

References

- Hargrove, W. W., F. M. Hoffman, and T. Sterling (2001), The do-it-yourself supercomputer, *Sci. Am.*, 265(2), 72–79.
- Hargrove, W. W., J. P. Spruce, G. E. Gasser, and F. M. Hoffman (2009), Toward a national early warning system for forest disturbances using remotely sensed phenology, *Photogramm. Eng. Rem. Sens.*, 75(10), 1150–1156.
- Hey, T., S. Tansley, and K. Tolle (Eds.) (2009), *The Fourth Paradigm: Data-Intensive Scientific Discovery*, 284 pp., Microsoft Corporation, Redmond, Washington, USA.
- Hoffman, F. M., and W. W. Hargrove (1999), Multivariate geographic clustering using a Beowulf-style parallel computer, in *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA '99)*, vol. III, edited by H. R. Arabnia, pp. 1292–1298, CSREA Press.
- Hoffman, F. M., W. W. Hargrove, R. T. Mills, S. Mahajan, D. J. Erickson, and R. J. Oglesby (2008), Multivariate Spatio-Temporal Clustering (MSTC) as a data mining tool for environmental applications, in *Proceedings of the iEMSs Fourth Biennial Meeting: International Congress on Environmental Modelling and Software Society (iEMSs 2008)*, edited by M. Sánchez-Marrè, J. Béjar, J. Comas, A. E. Rizzoli, and G. Guariso, pp. 1774–1781.
- Hoffman, F. M., R. T. Mills, J. Kumar, S. S. Vulli, and W. W. Hargrove (2010), Geospatiotemporal data mining in an early warning system for forest threats in the United States, in *Proceedings of the 2010 IEEE International Geoscience and Remote Sensing Symposium (IGARSS 2010)*, pp. 170–173, doi: 10.1109/IGARSS.2010.5653935, invited.
- Kendall, W., J. Wang, M. Allen, T. Peterka, J. Huang, and D. Erickson (2011), Simplified parallel domain traversal, in *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis, SC '11*, pp. 10:1–10:11, ACM, New York, NY, USA, doi: 10.1145/2063384.2063397.
- Luo, Y. Q., J. T. Randerson, G. Abramowitz, C. Bacour, E. Blyth, N. Carvalhais, P. Ciais, D. Dalmonech, J. B. Fisher, R. Fisher, P. Friedlingstein, K. Hibbard, F. Hoffman, D. Huntzinger, C. D. Jones, C. Koven, D. Lawrence, D. J. Li, M. Mahecha, S. L. Niu, R. Norby, S. L. Piao, X. Qi, P. Peylin, I. C. Prentice, W. Riley, M. Reichstein, C. Schwalm, Y. P. Wang, J. Y. Xia, S. Zaehle, and X. H. Zhou (2012), A framework for benchmarking land models, *Biogeosci.*, 9(10), 3857–3874, doi:10.5194/bg-9-3857-2012.
- Overpeck, J. T., G. A. Meehl, S. Bony, and D. R. Easterling (2011), Climate data challenges in the 21st century, *Science*, 331(6018), 700–702, doi:10.1126/science.1197869.
- Randerson, J. T., F. M. Hoffman, P. E. Thornton, N. M. Mahowald, K. Lindsay, Y.-H. Lee, C. D. Nevison, S. C. Doney, G. Bonan, R. Stöckli, C. Covey, S. W. Running, and I. Y. Fung (2009), Systematic assessment of terrestrial biogeochemistry in coupled climate-carbon models, *Global Change Biol.*, 15(10), 2462–2484, doi:10.1111/j.1365-2486.2009.01912.x.
- Ricciuto, D. M., M. L. Branstetter, D. Wang, J. Mao, J. Kumar, F. Hoffman, H. Najm, K. Sargsyan, C. Safta, and P. E. Thornton (2012), CSSEF progress towards uncertainty quantification in CLM-CN, DOE BER Poster Session, Oak Ridge National Laboratory, Oak Ridge, Tennessee.

*Observational Data Processing: Retrieval Algorithms and Instrument Network Simulation***Synthesis and Interpretation of 3D Atmospheric Measurements from a multi-dimensional, multi-sensor field laboratory**Pavlos Kollias^{1,2} and Edward Luke²

1. Stony Brook University, Stony Brook, NY
2. Brookhaven National Laboratory, Upton, NY

Urbanized centers, major ports of entry (airports, harbors), coastal areas, and DOE climate research facilities are current and future examples of locations where siting integrated high-resolution field observations (see 2013 BERAC report: BER Virtual Laboratory: Innovative Framework for Biological and Environmental Grand Challenge) is envisioned. These field laboratories and their counterpart modeling frameworks are required to: (i) advance our current understanding about the structure, variability, and interactions of complex physical processes that act across a wide range of scales and (ii) improve our ability to predict changes and evaluate the impact of decision making on human–nature interactions.

An example of an integrated field laboratory (IFL) is the U.S. Department of Energy (DOE) Atmospheric Radiation Measurement (ARM) program Southern Great Plains (SGP) megasite (<https://www.arm.gov/news/facility/post/29985>). Past observational strategies at ARM facilities were limited to profiling the atmosphere with a set of sensors that were capable of observing only parts of the atmospheric system. The new ARM radar facilities are capable of providing holistic multi-scale observations of clouds and precipitation to better understand the water cycle in our climate system. The new ARM program facilities operate radars at five different frequencies covering a wide range of scattering mechanisms, thus improving the information content of collocated multi-wavelength observations. This is particularly true in the vertical column; this provides enhanced, calibrated, multi-parametric measurements of clouds and precipitation. In addition, the ARM radar facilities feature scanning polarimetric Doppler radar observations over a wide range of climatological conditions. The main objective is to provide a revolutionary characterization of a complete volume of the cloudy atmosphere, over a long period of time, to act as a natural laboratory for the modeling community—both for testing their models and for improving the parameterizations of clouds in climate models. The configurable ARM radar facilities employ adaptive scanning strategies that enable focused experiments to study critical aspects of the water cycle at a range of spatial scales from the inner scale (30–50 m) to the outer scale (50–100 km).

In parallel with preparations for the SGP megasite, the ARM program recently kicked off a pilot study (LES ARM Symbiotic Simulation and Observation [LASSO]) that aims to provide a framework for performing routine high-resolution modeling and model-data integration at the SGP site (<https://www.arm.gov/science/themes/lasso>). It is conceivable that, in addition to the SGP, we will see similar activities at other ARM sites in the future (www.arm.gov).

However, it is only recently that the atmospheric research community has started considering observational and modeling facilities that produce this type of information. The sheer volume of observations from sensor networks and model output and the need to derive understanding from it suggests that extant computing resources are required. This applies to the following areas: (i) interpretation, visualization, and analysis of the natural processes as manifested through our observing sensors; (ii) application of inversion methods to extract physical parameters of interest; (iii) use of instrument simulators to convert model output to observation space; and (iv) application of machine learning methods to discover useful patterns and parameterizations in high-dimensional dataspace.

Models: A combination of low-resolution regional models and large eddy simulation (LES) models is required.

New Capabilities: Data assimilation techniques for assimilating at high resolution physical properties that traditionally have not been considered (e.g., cloud field properties, mass flux) need to be developed. Instrument network simulators that can convert the numerical model output to an integrated field laboratory view (instrument view) are needed. These network simulators will enable an apples-to-apples comparison of observations and models, and will allow the inclusion of new, advanced measurements in model evaluations, with less uncertainty due to forward models, rather than attempts at unstable, highly uncertain inversions (retrievals). Data synthesis techniques are another area where new capabilities are needed. Tens of instruments observing different, limited parts of the atmospheric state generate the need to merge their information into a best 3D/4D estimate of the atmospheric state.

Resolution: Focusing on boundary layer issues (urban heat island, renewable energy facilities) will require a very high horizontal (~25–50-m) and vertical (~20–30-m) resolution over a large horizontal domain (30–50 km) that captures the main modes of variability.

I/O: A significant increase in the I/O compared to today's models is required (~2 orders of magnitude). The model output at the proposed domain size and resolution is expected to grow exponentially over the next 5–10 years. It is important for us to develop appropriate operators that reduce the dimensions of the model output and capture its key structural information. This will align with what high resolution observations can provide: a basic, statistical description of the atmospheric state. On the other hand, observations, in particular radar observations at the ARM facilities, have grown exponentially already. The estimated ARM radar data volumes are on the order of few petabytes every year. At these volumes, the ability of PIs in universities to store and analyze such large datasets is limited. Large computational facilities at DOE laboratories will have to perform the bulk of the computations needed to analyze these data.

Many-Core and/or GPU Readiness: Our group, recognized as a leader in developing methods for the analysis of millimeter (cloud) and centimeter (precipitation) wavelength radar observations in atmospheric remote sensing, already routinely uses software developed in house for execution on GPUs. The datasets involved, radar Doppler spectra in particular, are information rich and notoriously large; in fact, they constitute the majority of data managed by the DOE ARM Program. While achieving performance gains on the order of 2 orders of magnitude with respect to equivalent CPU implementations, from our experience, external I/O has been the factor that ultimately limits throughput. We embrace machine learning, having extensively applied neural networks as a workhorse for extracting information from radar Doppler spectra. In addition, on the machine learning front, we see potential in evolutionary computing (genetic algorithms, genetic programming, etc.), an area of continuing active research, leading to discovery of new atmospheric process parameterizations valuable to climate modeling in an exascale computing ecosystem.

Uncertainty Quantification and Analytics of Observations and Model Results: Uncertainty Quantification, Statistics, Emulators, and Analytics of Model-Observation Comparisons

Calibration and Comparison of Climate Models: Accounting for Structural and Discretization Error

Bert Deusschere, Habib Najm, Khachik Sargsyan, and Kenny
Chowdhary
Sandia National Laboratories, Albuquerque, New Mexico

1. Describe a major science challenge expected to be solved in the 2020–2025 timeframe that requires using exascale computing ecosystems.

A major science challenge in climate modeling is accounting for structural and discretization error to improve predictive fidelity. Climate models are complex as they are made up of many discrete formulations of resolved or under-resolved physical processes pertaining to the atmosphere, the land and its ecosystems, the ocean, and other systems. Formal methods for model calibration have been developed in the statistics and uncertainty quantification (UQ) communities, often relying on Bayesian concepts, for example, Kennedy and O’Hagan [2001]. The application of these concepts to coupled climate system models is still in its infancy, however, and is limited by the high computational cost of climate simulations. Moreover, while parameter tuning generally improves the agreement with available observations, the predictive skill of the climate model is often limited by structural error.

To improve the predictive fidelity of climate models, it is essential to take into account both the effect of structural error as well as spatial and temporal discretization errors. Both types of error can be significant, have historically not been accounted for, and can be tightly linked as process parameterizations are strongly influenced by the time and length scales that can be resolved.

Novel UQ methods are able to account for structural error by embedding statistical model error terms in model parameters, such that the resulting calibrated uncertain model prediction is representative of the impact of structural uncertainty in the model. This approach has been demonstrated, for example, in the calibration of chemical models [Sargsyan et al., 2015] and tends to be more effective than additive model error terms, especially for assessing the uncertainty in quantities of interest (QoIs) for which no observational data are available. Applying these methods to climate models will require very close collaboration between climate scientists and UQ practitioners to determine the best place(s) to embed model error terms. Methods that account for discretization error as a source of uncertainty and incorporate it accordingly in making uncertain predictions are also emerging in UQ [Conrad et al., 2015]. These methods are still in their infancy, particularly in regard to large-scale computational models, but offer a promising framework. They will require significant development, indeed, to be evaluated and applied to climate models.

The challenges are formidable, but the impact on climate modeling would be enormous. Being able to account for structural and discretization error in calibration would allow for automated and robust tuning of climate models in a way that takes advantage of all available observational data in a consistent way, is not disrupted by changes in grid or temporal resolution, can readily handle updates in model components, and is able to readily incorporate new data as they become available. Such a framework would be invaluable for determining where models need to be

improved and where more data would be the most beneficial for enhancing the overall predictive fidelity of the climate model.

2. Current and future computational and data strategies

Models

The advances described here are germane to all climate models in operation. It is a cross-cutting capability. It will be important to have more data available, especially data that allow differentiation between the contributions of different model components. As such, rigorous quantification of the observational data uncertainty and development of metrics for comparing model outputs to observational data will be essential.

New capabilities

New algorithmic capabilities will need to be developed to account for model and discretization error. Their application is likely to require tremendous computational resources in terms of very large ensembles to account for the associated uncertainties. To mitigate these computational costs, new approaches for surrogate models, multi-fidelity calibration, Markov Chain Monte Carlo methods, and optimal experimental design will be required, as well as better methods for choosing ensemble member configurations and durations.

Resolution

There is a very interesting interplay between model error and discretization error, since the model error is often tied to a level of approximation that is tied to the scale on which physical phenomena are resolved. Theoretical and algorithmic advances will be needed to rigorously tie these together. From a computational efficiency point of view, a multi-fidelity, multi-resolution approach will be required to enable the calibration. As such, solid support for adaptive and regionally refined resolution in climate models will be important. Also, being able to run at resolutions much finer than today's resolutions will aid in the characterization of the model and discretization error at coarser resolutions.

I/O

I/O should be similar compared to today's models for each run, but with higher requirements for storage as the grid resolutions increase. The necessity to run many more ensembles will require much more I/O for the set of simulations as a whole, which will require a lot of temporary storage. Accounting for model and discretization error should not lead to more permanent data storage.

Many-core and/or GPU readiness

Large-scale ensembles typically enable a lot of coarse-grained parallelism, but not fine-grained parallelism. To get more fine-grained parallelism, which allows one to take advantage of many-core and GPU systems, one approach is to run embedded ensembles, that is, rely on intrusive or hybrid methods rather than strictly on sampling-based [Phipps et al., 2013]. This capability will need to be further developed and applied to climate models.

Work flows

A typical UQ work flow involves determining key model parameters, performing sensitivity analysis to assess their importance, followed by a dimensionality reduction and surrogate construction. Calibration against available data sets then proceeds, followed by a forward propagation of the inferred uncertainties to assess the predictive fidelity of the model. Sensitivity analysis and surrogate construction typically require large ensembles.

3. What top three computing ecosystem aspects will accelerate or impede progress to meeting the challenge described above in the next 5–10 years? Why?

| Accelerate | Why? |
|---|--|
| 1. Internal/external libraries/frameworks | Will provide tools for effective use of next-generation architectures. |
| 2. Models and algorithms | Will provide needed functionality for handling model and discretization error. |
| 3. Hardware | Allow for larger ensembles. |

| Impede | Why? |
|--------------------------------|---|
| 1. Hardware power restrictions | Will reduce robustness of computing systems and will require more resilience or fault mitigation. |

References

Conrad, P. R., Girolami, M., Särkkä, S., Stuart, A., and Zygalakis, K. (2015). Probability Measures for Numerical Solutions of Differential Equations. *Annals of Applied Probability*, 23(4), 1437–1468. Retrieved from <http://arxiv.org/abs/1506.04592>.

Kennedy, M.C., and O’Hagan, A. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B*, 63(3): 425–464.

Phipps, E., Edwards, H. C., Hu, J., and Ostien, J. T. (2014). Exploring emerging manycore architectures for uncertainty quantification through embedded stochastic Galerkin methods. *International Journal of Computer Mathematics*, 91(4), 707–729. <http://doi.org/10.1080/00207160.2013.840722>.

Sargsyan, K., Najm, H. N., and Ghanem, R. (2015). On the Statistical Calibration of Physical Models. *International Journal of Chemical Kinetics*, 47(4), 246–276. <http://doi.org/10.1002/kin.20906>.

Uncertainty Quantification in Climate Prediction and Projection

Charles Jackson

The University of Texas at Austin

1. Science challenge

Over the next 5–10 years, we can anticipate there will be a clearer understanding of how the representation of uncertain processes are impacting key vulnerabilities that society has to a changing climate. At present, there are only vague, hypothesis-driven notions of how observations are used to test the adequacy of climate models for their intended uses. The ability to make significant headway on this problem will depend on the level of investment that is made now to sort out how observations can be used to inform parameterizations within a hierarchy of test beds that can cheaply and effectively connect observations to the representation of uncertain processes.

Uncertainty quantification involves an exploration of the space of model plausibility. At present, establishing such plausibility for a coupled climate system model for a single-model version involves significant human and computational resources. Thus, to make uncertainty quantification remotely feasible, we have to establish a work flow (if that is the correct word) that can take advantage of model hierarchies and other shortcuts that can more cheaply link model behavior on short time scales with limited interactivity with the behavior of the model within a coupled climate system model. Uncertainty quantification provides formalism to this work flow that leverages computational resources to help connect decisions that need to be linked, such as determining the values for two or more parameters that affect the long-term behavior of an important phenomenon.

2. Current and future computational and data strategies

I am a strong advocate for sampling-based approaches to uncertainty quantification in part because it is well paired with the hypothesis-driven approach scientists take to make use of carefully constructed experiments and data. Within a statistical framework, one may use information about the rate of change of model plausibility measures to infer the relative density of solutions. The computational and data requirements for this approach correspond to dimensionality of the uncertainty space. The requirements are sobering. I have based my estimates of computational requirements on the results of Jackson et al. (2004) and Villagran et al. (2008) and have been borne out in a number of applications that range from idealized (cheap) models, land surface models, the MITgcm, and CAM3.1. Typically I have focused on problems with ~6 dimensions, but have also been successful with the MITgcm and CAM3.1 exploring 16-dimensional parameter spaces. The 6-dimensional spaces require ~500 experiments, whereas exploration of 16-dimensional spaces has made do with ~3,000 experiments. The length of each experiment should be long enough to see the effects of the parameter change relative to the model's internal variability. For AMIP-style experiments with CAM in which SSTs and sea ice are prescribed, one needs at least 3–4 years of model integration. This equates to the equivalent of a ~10,000-model year integration for estimating uncertainties in 16 parameters of CAM.

Models/Capabilities

In a hierarchical work flow, one would attempt to use test beds and a hierarchy of model versions to estimate model response sensitivity to parameter perturbations so that one could limit the search space. There are also strategies one could employ to estimate the equilibrated response of a model to a change in parameters more cheaply than using time integration. Use of model adjoints (if they exist) or an ensemble of very short experiments could perhaps reduce the costs by an order of magnitude or more. The use of adjoints is a feasible approach for ice sheet models. It may be

possible for other components as well, although the approach has its limits when the equilibrated state of a system is a result of high-frequency variability (e.g., the ocean).

Resolution

In some circumstances, one can make use of lower resolution versions of a model to anticipate model plausibility at higher resolutions. However, the extent to which new phenomena occur at higher resolutions and not at lower resolutions limits the applicability of multi-fidelity methods.

I/O

When one is conducting thousands of experiments, it is easy to drown in model output. Since many science drivers are driven by resolution in both time and space, I/O can be a big issue. UQ calculations can anticipate the data that are needed to establish plausibility and run jobs in parallel to post-process the data. At present, since it is not yet clear what observables best test model plausibility, the norm has been to save everything. This is a safe approach and scientifically useful; however, it will require large disk spaces to keep the results of all the experiments easily accessible.

Work flows

I have developed MECS, an open-source python set of scripts, for managing the UQ work flow over HPC resources. The UQ work flow involves running a set of experiments in parallel, running scripts for post-processing the data and for generating metrics for establishing model plausibility, and either using those results to determine a new set of experiments (i.e., an adaptive approach) or cycling through a list of predetermined experiments according to a design (i.e., Latin hypercube). Both strategies typically run ~10 experiments in parallel (depending on the number of uncertain parameters). This is because Latin hypercube (or polynomial Chaos) typically need the same number of experiments as adaptive approaches. This implies that there are ~100 experiments in each sequence, which means that the computation has to be sustained for about 3 months. This assumes each experiment takes 24 hours of wall-clock time and there is fast turnaround between experiments.

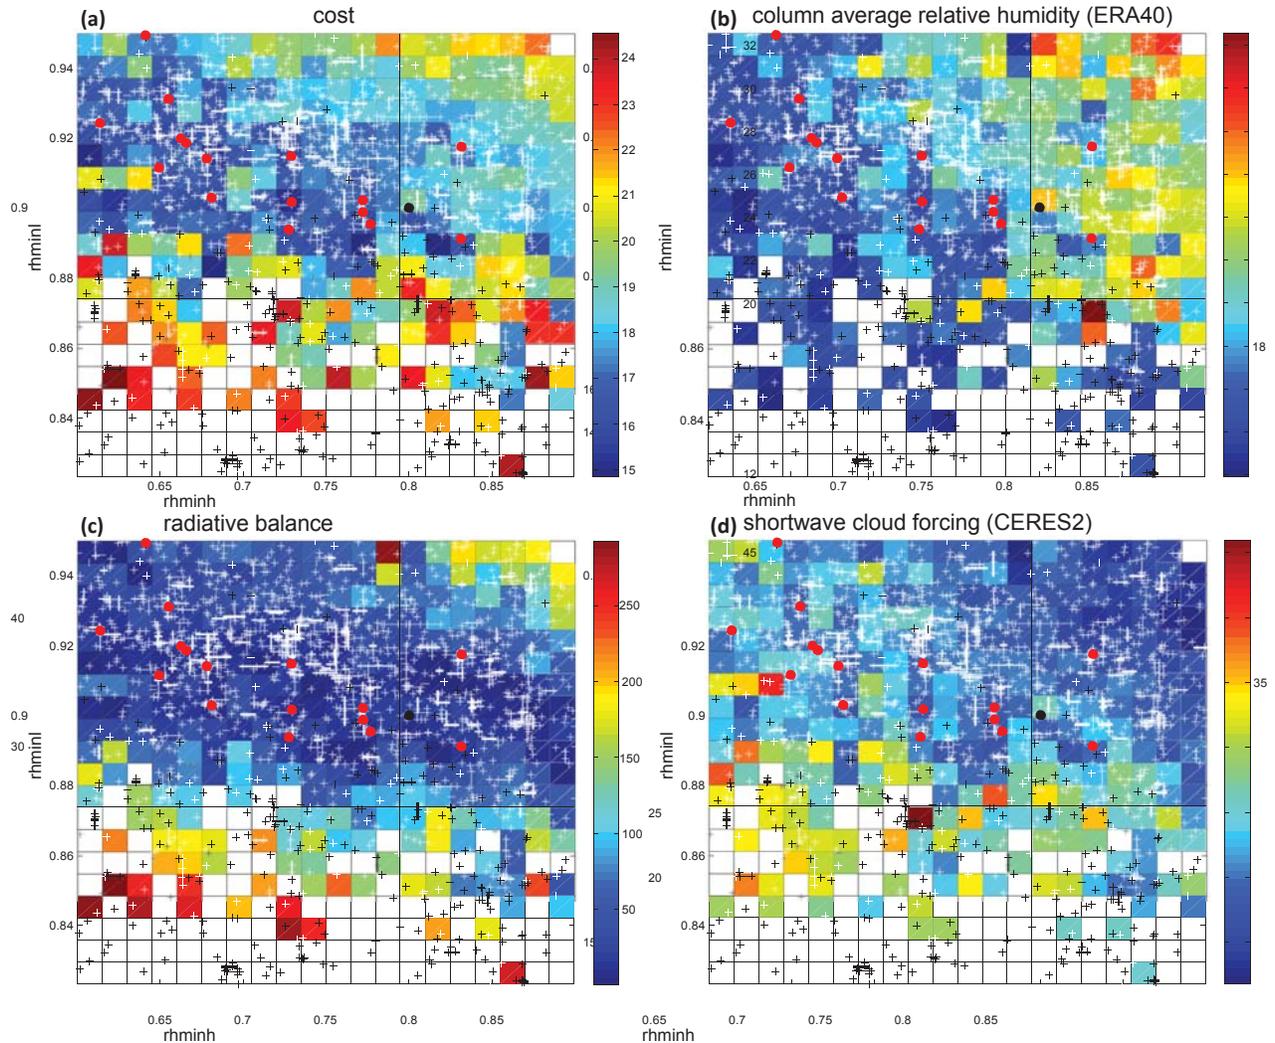
3. Priorities

With adequate computing, uncertainty quantification has the potential to transform climate model development, particularly the way this process integrates observations. For this to work well, *a significant amount of resources needs to be available* to learn how to make this process work more cheaply. One does not know a shortcut works until one can demonstrate its effectiveness. The second most important development is the *availability of a hierarchy of application codes*, such as a single-column model or experiment frameworks such as CAPT, which bridge the gap between model physics packages and observations. Third, there has been a relatively slow adoption of UQ in climate model development. This is likely because there are limited resources, both computational and human, to pursue these topics.

References

Jackson, C., M. Sen, and P. Stoffa (2004) An Efficient Stochastic Bayesian Approach to Optimal Parameter and Uncertainty Estimation for Climate Model Predictions, *Journal of Climate*, 17(14), 2828-2841.

Villagran, A., G. Huerta, C. S. Jackson, and M. K. Sen (2008) Parameter Uncertainty Estimation in Climate Models. *Bayesian Analysis*, 3(4), 823–850.



Accelerating the Uncertainty Quantification of Climate Models through Surrogate-Based Approaches

George S. H. Pau, Lawrence Berkeley National Laboratory

1. Describe a major science challenge expected to be solved in the 2020–2025 time frame that requires using extant computing ecosystems.

Within the next 5–10 years, we expect to provide a robust quantification of the impact of climate change on humans, and vice versa. Most existing integrated assessments are scenario-driven and cannot fully explore the interactions between the human and natural drivers of climate change. In addition, a quantitative, probabilistic justification of risks and benefits will provide a rigorous basis for policy- and decision-making.

My research focuses on developing efficient approaches that accurately quantify and reduce uncertainties in human-climate system (abbreviated henceforth to climate) models. Direct application of robust uncertainty quantification techniques, such as Monte Carlo methods, is typically infeasible even with existing high-end computing ecosystems since it involves $>10^5$ evaluations of high-resolution climate models, each requiring significant computational resources. My research relies on developing rigorously certified reduced order models or surrogate models that emulate the climate models at a significantly reduced cost, allowing uncertainty analysis to be performed at the desired spatial and temporal scales. Construction of these models relies on an offline-online computational framework that allows efficient and judicious use of high-end computing ecosystems. Data storage and retrieval are crucial as well since these surrogate models are typically built using data-intensive techniques that assimilate both simulated and measured data.

2. Current and future computational and data strategies

My computational strategy is based on an offline-online computational framework that allows uncertainty quantifications to be performed efficiently using surrogate models (online stage) through an amortization of the construction cost of these models (offline stage). The offline stage is computationally intensive because of the need to obtain field solutions from a large number of climate model evaluations. The construction of the surrogate models from these solutions can also be computationally and memory intensive. An additional advantage of this computational framework is its efficient utilization of heavily shared high-performance computing resources. By executing the offline stage during the off-peak cycles, we are able to execute the online stage even during peak cycles. We can also execute the online stage on smaller machines with a smaller user base and thus better throughput.

The above computational strategy will remain relevant within the next 10 years for two reasons. First, advances in computing technology and numerical methods during that period of time are unlikely to lead to a speedup needed to make a direct utilization of climate models in Monte Carlo techniques feasible. Second, the complexity and thus the computational cost of climate models will continue to rise as the modeling community aims to improve these models by including more processes at the desired spatial and temporal scales. The latter will remain an important focus for most scientists within the modeling community.

Models

Dimensional reduction and statistical techniques are used to construct the surrogate models. We use proper orthogonal decomposition and different response surface approaches, such as Gaussian process

regression. These techniques train a surrogate model based on solutions obtained from running the high-resolution models at samples judiciously selected from a space spanned by the parameters of interest. The accuracies of the models thus depend on the training sample set, as well as on whether the chosen techniques are appropriate for describing the dynamics of the approximated variables.

New capabilities

New requirements for the surrogate models will depend on the new capabilities of the climate models. If the accuracy of the climate models increases, the surrogate models must also be constructed more accurately by increasing the number of samples in the training sample set and optimizing the hyperparameters of the surrogate models using a more exhaustive approach. New methods may be needed to handle new variables that have different dynamics than existing variables. Solutions with higher resolution mean it is no longer feasible to load all training data into the memory; efficient I/O capabilities are thus needed. The computational cost and memory requirements for both the offline and online stages are expected to increase.

Resolution

To explore regional scale effects of global climate change, models need to be simulated at 0.1° in order to begin studying localized impact at the human level. While this is only an order of magnitude smaller than the resolution of existing models, the computational and development efforts will increase exponentially. For surrogate models, while global uniform resolution and regionally refined resolution are easier to handle, adaptive resolution can be easily accommodated by mapping the resulting solution onto a globally refined mesh.

I/O

Higher spatial and temporal resolutions of the climate models will lead to an increase in I/O. In addition, construction of the surrogate models will require frequent I/O since the on-board memory cannot hold all the training data. These data will need to be permanently archived to ensure that the provenance of the surrogate models is documented.

Many-core and/or GPU readiness

For the construction of the surrogate models, we have not looked into the use of many-core and/or GPU. Since the online stage involves many independent evaluations of the surrogate models, transitioning the code to many-core or GPU architecture should not be too difficult. However, the offline stage will require more in-depth study to understand how the new architecture can be efficiently utilized.

Work flows

- a. Identify the climate models and the parameter space that we would like to study.
Requirements: As the climate models become more complex, the cost of evaluating these models should not increase dramatically more than the increase in the computational capacity.
- b. Choose a set of samples from the parameter space, either statistically or adaptively.
- c. Evaluate the climate models corresponding to the above training sample set. This step is the most computational intensive step but can easily be run in parallel.
Requirements: Able to run hundreds of the chosen climate model within a month in 2020, and thousands in 2025.
- d. Construct surrogate models based on the simulations results obtained in the previous step. *Requirements:* Able to use the new computing architecture to speed up the construction procedure and handle a large training data set.
- e. Repeat steps b through d until surrogate models have the desired accuracy.
- f. Use the surrogate models in uncertainty analyses.

Requirements: Able to use the new many-core/GPU technologies to run many parallel evaluations of the surrogate models.

3. What top three computing ecosystem aspects will accelerate or impede progress to meeting the challenge described above in the next 5–10 years? Why?

| Accelerate | Why? |
|--|---|
| 1. Hardware resources (at all scales) including I/O, memory, etc. | Greater capacity will allow more simulations to be executed, thus increasing the accuracy of the surrogate models. |
| 2. Models and algorithms | Better dimension reduction and statistical techniques can lead to more accurate surrogate models. |
| 3. Data work flow (including sharing, transmitting, archiving, etc.) | Better sharing and archiving will allow data from different projects to be utilized for constructing the surrogate models, thus reducing the computational cost of the offline stage. |

| Impede | Why? |
|---|--|
| 1. Application codes (implementation, development, portability, etc.) | If the application codes for the climate models cannot be implemented efficiently for the new computing architecture, the computational cost of the offline stage will increase. |
| 2. Data work flow (including sharing, transmitting, archiving, etc.) | Inconsistencies in how data are shared and archived will severely impede the ability to use data from different projects for constructing the surrogate models. |
| 3. Workforce development | We need more scientists with the ability to understand, modify, or improve codes that are becoming increasingly complex. |

4. Software applications, libraries, and tools

I do not anticipate any changes in the above between now and 2020 or 2025.

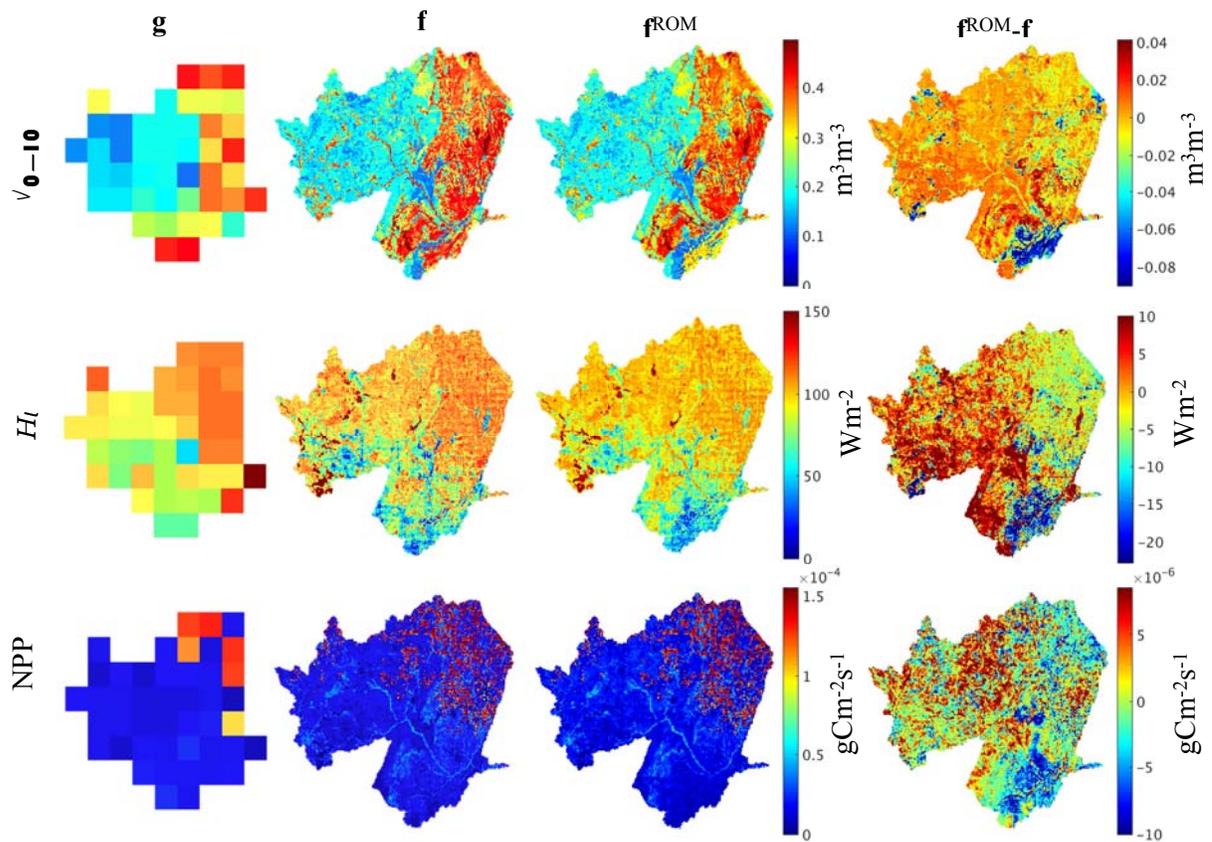
5. HPC Services

Better support for database management will be very helpful.

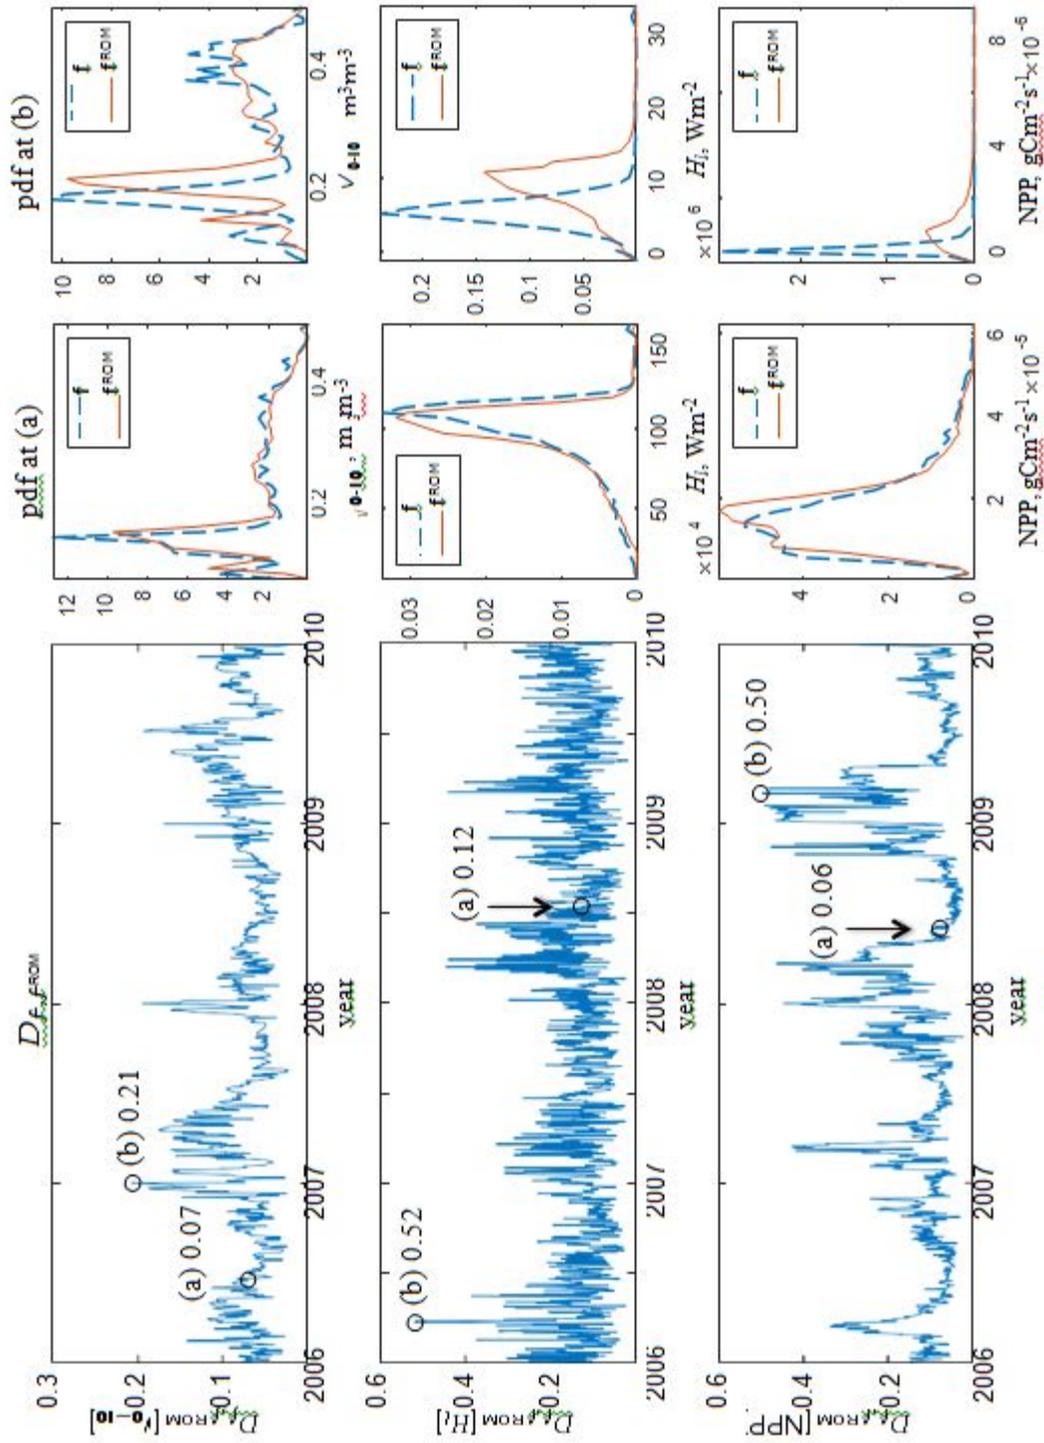
6. References

- a. Pau, G. S. H., C. Shen, W. J. Riley, and Y. Liu (2016), "Accurate and efficient prediction of fine-resolution hydrologic and carbon dynamic simulations from coarse-resolution models," *Water Resour. Res.*, 52, doi:10.1002/2015WR017782.
- b. Pau, G. S. H., G. Bisht, and W. J. Riley (2014). "A reduced-order modeling approach to represent subgrid-scale hydrological dynamics for land-surface simulations: application in a polygonal tundra landscape." *Geosci. Model Dev.* 7(5): 2091–2105.
- c. Liu, Y., G. Bisht, Z. Subin, W. J. Riley, and G. S. H. Pau (2015). "A hybrid reduced-order model of fine-resolution hydrologic simulations at a polygonal tundra site." *Vadose Zone Journal*.
- d. Pau, G. S. H., Y. Zhang, and S. Finsterle (2013). "Reduced order models for many-query subsurface flow applications." *Computational Geosciences* 17(4): 705-721.

7. Images



Source: Figure 1 of Pau, G. S. H., et al. (2016).



Source: Figure 6 of Pau, G. S. H., et al. (2016).

“Quantifying Uncertainty in the Antarctic Contribution to Sea Level Rise” BER/CESD Exascale Facilities Workshop White Paper/Case Study

Nathan Urban, Los Alamos National Laboratory

1. Describe a major science challenge expected to be solved in the 2020–2025 timeframe that requires using extant computing ecosystems.

The ocean-driven instability of the Antarctic ice sheet is responsible for large uncertainties in future sea-level rise projections over the next century (Joughin and Alley, 2011). This uncertainty has critical implications for the climate resilience of coastal populations. Uncertainties arise from global and regional atmospheric-ocean processes that set the properties of relevant water masses such as the circumpolar deep water (CDW); finer-scale eddy-driven transport of heat from the open ocean across the continental shelf; and mechanical, thermodynamic, and topographic bed uncertainties in the ice sheet response. Few of these uncertainties have been explored deeply in numerical simulations and none have been explored in combination. As is common in uncertainty quantification (UQ), ensembles of high-fidelity simulations are needed to explore the space of possible climate futures. Large ensembles of high-resolution, coupled climate models can easily tax the capacities of even exascale computing facilities, and “big data” approaches will be needed to develop statistical/machine learning approaches to emulate the behavior of computationally expensive numerical codes.

2. Current and Future Computational and Data Strategies

To address the UQ problem of Antarctic ice sheet retreat, it will be necessary to devise a computational strategy to combine all available sources of information, including high-resolution coupled global simulations, larger ensembles of lower-resolution/regional simulations, offline model component simulations, idealized process studies, non-DOE models, and observational data (Kopp et al., 2014). This information fusion process may take the form of a graphical network model of linked statistical emulators or reduced-order dynamical models trained to the available data (Kyzurova et al., 2016). Simulation strategies will be required that may not correspond directly to the transient forced model configurations usually used to project sea level rise, as discussed below.

It is particularly important to develop a means to address the under-explored problem of model structural uncertainties. The standard UQ paradigm is perturbed-parameter ensembles, but varying parameters within a single version of a single (e.g., DOE) climate model is unlikely to span the space of relevant structural uncertainties. For example, updating structural parameterizations in a model component greatly impacts the model projections (Kay et al., 2012), to the extent that uncertainties found in a perturbed-parameter ensemble clearly under-represent uncertainties within a multi-model ensemble (Yokohata et al., 2011). There are large between-model differences in, for example, the CDW ocean changes (Little and Urban, 2016). Model structural uncertainties may be addressed through reduced-order models whose free parameters can be tuned to reproduce structural differences (Jonko and Urban, 2016), which may require new model reduction approaches that are non-intrusive (do not require code access/knowledge of the governing equations), given the complexity of modern numerical codes (Chen, 2012).

Models: It will not be possible to robustly quantify all relevant uncertainties in Southern Ocean circulation, eddy transport, and ice sheet dynamics using only high-resolution coupled climate model ensembles. A model hierarchy is needed, combining (1) a limited number of high-resolution simulations with (2) greater number of lower-resolution simulations to explore uncertainties and (3) statistical approaches to bridge results from (1) and (2) across the consequent resolution gap (Goh et al., 2013).

Other ways to explore uncertainties will include running short integrations of coupled models (Wan et al., 2014) or offline model component ensembles (e.g., an ice sheet model forced with prescribed ocean boundary conditions). Most UQ will be done with statistical models calibrated to emulate simulation and observational data. To train those models it may be preferable to run simulations designed to maximize the signal-to-noise ratio (e.g., with a variety of step-function or impulse forcings). Idealized process studies or other simplified configurations can inform statistical models when an adequate range of fully coupled simulations may not be possible (e.g., to quantify uncertainty in fine-scale eddy-driven heat transport).

New Capabilities: Many of the advances needed to constrain potential future sea level rise due to Antarctic ice loss will be determined by what we choose to simulate, rather than by modifications to the simulation codes themselves. However, given the increasing simulation output volumes associated with Southern Ocean/Antarctic-centered UQ activities, *in situ* data assimilation/calibration capabilities within standalone and coupled climate component code would provide more efficient UQ. Alternatively, new *in situ* model reduction methods could be developed; the resulting reduced models would then be used offline in a UQ calibration setting. In addition to reducing the need to access the filesystem, *in situ* model reduction algorithms that can access state variables and intermediate calculations at the time-step level could potentially also be made more intrusive, and more accurate, than traditional statistical emulators that operate only on aggregated black-box output. To facilitate *in situ* UQ algorithms, it would be helpful to have an analysis framework that could “plug in” routines from external statistical codes and libraries to sample and manipulate state variables, including those written in high-level languages rather than just Fortran/C++. Embedded ensemble propagation techniques exploit redundant code parallelism to efficiently run ensembles within a single simulation instead of in multiple processes (Phipps et al., 2015). To explore multi-model uncertainties, it may be desirable to develop versions of DOE model components (e.g., ocean models) that can accept boundary conditions derived from non-DOE coupled models (e.g., Shchepetkin and McWilliams, 2005).

Resolution: Basic uncertainties in CDW formation could be explored with ensembles of existing 100-km global ocean models, but 10-km models (in the Southern Ocean, and possibly globally) will be needed to estimate biases in the coarser models. In addition, to resolve cross-shelf heat transport, ocean resolutions of <1 km near the Antarctic coastline may be needed (Stewart and Thompson, 2013).

I/O: I/O will scale with the number of ensemble members run (and that number is typically “as large as is practical”). Typically only monthly fields will need to be archived, but they may be 3D ocean fields at high regional resolution. *In situ* reduction methods can reduce the need for large storage capabilities.

Many-Core and/or GPU Readiness: The UQ will primarily leverage whatever climate simulation codes are available (many-core/GPU ready or not). *In situ* model and data reduction capabilities will require designs specific to heterogeneous architectures, which could be done from scratch instead of porting existing algorithms as these capabilities do not currently exist; embedded ensemble propagation, or running large numbers of “embarrassingly parallel” independent processes, will also exploit parallelism.

Workflows: Current workflows fit non-intrusive reduced models in an offline setting to archived CMIP simulation output (e.g., fitting the parameters of simplified or intermediate-complexity mechanistic models or semi-empirical data-driven models to complex numerical simulation data). To address the full UQ challenge, a wide variety of new ensemble simulation experiments must be designed, and statistical emulators/reduced models must be constructed either offline or online. These emulators must then be combined into a hierarchical/graphical statistical model that fuses different information sources.

3. What top three computing ecosystem aspects will accelerate or impede progress to meeting the challenge described above in the next 5–10 years? Why?

| Accelerate/Impede | Why? |
|--------------------------|---|
| 1. Workforce development | The labor required to do large-scale UQ is comparable to the labor required to do large-scale scientific studies or model development, and would benefit from UQ teams operating in a modular way analogous to how different model components are developed and integrated. There is not yet enough staffing, and no cross-institution teams exist, to address comprehensive decision-relevant UQ problems spanning the full Earth system (as opposed to the status quo, small scientific studies of individual uncertainties). |
| 2. Models and algorithms | New methods are needed to deal with the problems of model reduction and system identification, <i>in situ</i> statistical analysis, structural/multi-model UQ, and fusion of heterogeneous information sources. Few of these capabilities currently exist, although there is a strong foundation in the literature from which to develop new capabilities. |
| 3. Hardware resources | As many cores as possible are needed for ensemble studies, particularly to resolve uncertainties in eddy heat transport near the Antarctic coastline. The cores are needed for large ensembles of lower-fidelity/offline/idealized/etc. simulations, in addition to ultra-high resolution “hero”-class simulations. |

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.

5. HPC Services: Ability to work with HPC professionals and model developers to interface statistical codes with simulation codes.

References

Chen, H., “Blackbox stencil interpolation method for model reduction,” S.M. thesis, MIT (2012).

Goh, J., D. Bingham, J.P. Holloway, M.J. Grosskopf, C.C. Kuranz, and E. Rutter, "Prediction and computer model calibration using outputs from multifidelity simulators," *Technometrics* **55**, 501 (2013).

Jonko, A.K., and N.M. Urban, "A hierarchical Bayesian method for multi-model combination and calibration of climate models," in preparation (2016).

Joughin, I., and R.B. Alley, "Stability of the West Antarctic ice sheet in a warming world," *Nature Geoscience* **4**, 506 (2011).

Kay, J.M., M.M. Holland, C.M. Bitz, E. Blanchard-Wrigglesworth, A. Gettelman, A. Connley, and D. Bailey, "The influence of local feedbacks and northward heat transport on the equilibrium Arctic climate response to increased greenhouse gas forcing," *Journal of Climate* **25**, 5433 (2012).

Kopp, R.E., R.M. Horton, C.M. Little, J.X. Mitrovica, M. Oppenheimer, D.J. Rasmussen, B.H. Strauss, and C. Tebaldi, "Probabilistic 21st and 22nd century sea-level projections at a global network of tide gauge sites," *Earth's Future* **2**, 383 (2014).

Kzyurova, K.N., J. Berger, and R.L. Wolpert, "Coupling computer models through linking their statistical emulators," *SIAM Conference on Uncertainty Quantification*, Lausanne, Switzerland (2016).

Little, C.M., and N.M. Urban, "CMIP5 temperature biases and 21st century warming around the Antarctic coast," *Annals of Glaciology*, in revision (2016).

Phipps, E., M. D'Elia, H.C. Edwards, M. Hoemmen, J. Hu, S. Rajamanickam, "Embedded ensemble propagation for improving performance, portability and scalability of uncertainty quantification on emerging computational architectures," arXiv:cs/1511.03703 (2015).

Shchepetkin, A.F., and J.C. McWilliams, "The regional oceanic modeling system (ROMS): a split-explicit, free-surface, topography-following-coordinate oceanic model," *Ocean Modelling* **9**, 347 (2005).

Wan, H., P.J. Rasch, K. Zhang, Y. Qian, H. Yan, and C. Zhao, "Short ensembles: an efficient method for discerning climate-relevant sensitivities in atmospheric general circulation models," *Geoscientific Model Development* **7**, 1961 (2014).

Yokohata, T., J.D. Annan, M. Collins, C.S. Jackson, M. Tobis, M.J. Webb, and J.C. Hargreaves, "Reliability of multi-model and structurally different single-model ensembles," *Climate Dynamics* **39**, 599 (2011).

*Model Development Testbeds***Diagnosis and Evaluation of Climate Model Biases and Parameterizations with the Coupled CAPT (Cloud-Associated Parameterizations Test Bed)**

Hsi-Yen Ma
Lawrence Livermore National
Laboratory
Livermore, California

1. Describe a major science challenge expected to be solved in the 2020–2025 time frame that requires using extant computing ecosystems.

The CAPT (Cloud-Associated Parameterizations Test Bed) approach or the Transpose AMIP approach is now a widely used hindcast technique to diagnose the contribution of fast physical processes in the atmosphere to long-term errors in climate simulations, as well as to evaluate cloud-associated parameterizations in global climate models (GCM).

In the next 5 to 10 years, the LLNL CAPT project will build a modeling framework with a hindcast technique for evaluating cloud-associated parameterizations in a fully coupled high-resolution GCM employing the most advanced cloud process-level observations. Since cloud parameterizations must perform well in a fully coupled climate system, it is logical and desirable to perform parameterization testing with a coupled climate system model.

In order to perform routine parameterization testing in a fully coupled climate system, a significant amount of computing time to efficiently perform extensive ensemble coupled model hindcasts, as well as large data storage (both scratch and long-term space) to store model output, is required.

2. Current and future computational and data strategies*Models*

We currently test atmospheric parameterizations mostly for the CAM at 1-degree horizontal resolution with 30 vertical levels. In the future, we will test the fully coupled CESM and ACME with higher horizontal and vertical resolutions.

New capabilities

A highly scalable, coupled model with the capability to automatically run ensemble hindcasts and post-processing programs for model evaluation will have to be developed for routine parameterization testing.

Resolution

In the next 5 to 10 years, we anticipate routinely evaluating models at $\sim 1/4$ - or $1/8$ -degree horizontal resolution for atmosphere and land model components and $\sim 1/10$ -degree horizontal resolution for the ocean and sea ice model components. We will use *either* a globally uniform, high-resolution model or a regionally refined model for a targeted area.

I/O

Additional I/O will definitely be required, and it will depend on model resolutions, number of output variables, and frequency of output variables.

Many-core and/or GPU readiness

The current ACME or CESM is not ready to utilize GPUs. As most climate scientists do not have extensive knowledge about computer hardware, computer scientists, programmers, or software engineers who are familiar with computer architectures and climate model source codes are necessary for the transition to the utilization of GPUs.

3. What top three computing ecosystem aspects will accelerate or impede progress to meeting the challenge described above in the next 5–10 years? Why?

| Accelerate | Why? |
|---|--|
| 1. Hardware resources | Faster and more CPU cores, more memory |
| 2. Models and algorithms | More computation-efficient codes |
| 3. Internal/external libraries/frameworks | More computation-efficient libraries |

| Impede | Why? |
|--|--|
| 1. Data work flow (sharing, transmitting, archiving) | Increased data volume due to high-resolution model output |
| 2. Models and algorithms | More climate processes will be included in the models. |
| 3. Hardware resources | More sophisticated computer architecture will make code implementation and model development/debugging more challenging. |

4. Software applications, libraries, and tools

Highly parallelized post-processing software applications for model output are critical because of the data volume of high-temporal-frequency and high-spatial-resolution model output.

5. HPC services

Real-time data analytics and visualization with web interfaces will halve the evaluation process. However, this will require additional HPC services for data processing and calculation.

*Transforming Science through Exascale Capabilities: Algorithms and Computational Science
(corresponds to Section 3.2.7)*

Software Engineering for Portability

**Flexible Looping and Directives for Portable Performance on Multiple Types of Accelerators
(Matt Norman, ORNL)**

- 1. Describe a major science challenge expected to be solved in the 2020–2025 timeframe that requires using extant computing ecosystems.**

Regardless of the problems we solve in the future, we are practically guaranteed that future computing will have some form of accelerator, whether it be a GPU, a MIC, or something else. Accelerators are marked by increased compute throughput relative to aggregate memory bandwidth, and they are marked by increasingly complex memory hierarchies. We need a flexible and portable strategy to obtain good performance on different accelerators moving forward.

- 2. Current and Future Computational and Data Strategies**

There have been certain difficulties in portably using accelerators. For instance, K20x GPUs require explicit use of CUDA shared memory in order to use the per-SM L1 cache, and in OpenACC, this requires manual creation of an ad-hoc-sized temporary local array that is passed down from the gang level loops (highest level loops), making portability difficult. However, K40 and above will automatically cache any GPU DRAM access through the L1 cache without having to use shared memory. What this means is that for future-generation GPUs (i.e., those newer than what we have on OLCF's Titan), one can code based on existing caching constructs by blocking (also called tiling) loops at a certain level and still gain good performance.

Automatic loop blocking will likely not be available when exascale computing rolls out on the floor. At the very least, directives will need to be used to guide the compiler as to when and how to block loops. There is currently a tile directive in the OpenACC standard; this is at least a start. It gives the compiler information on how to tile the loop instead of having to use automatic tiling, which muddies and complicates the array indexing and is prone to bugs. Having directives for tiling aids performance portability. However, caching is increasingly hierarchical, especially on the MIC, and having the ability to tile loops at multiple levels with directives is a must.

In addition, most codes are structured with many calls to reusable subroutines like “divergence,” “curl,” or viscosity,” each relating to a physically or algorithmically motivated computation to be done on the model state data. Often, for CPU optimization, a relatively small amount of looping is exposed to these routines so that they can each operate on a certain (hopefully cache-friendly) block of data. It will be necessary for performance portability to have the ability to flexibly push looping down the call stack or pull it back out. For example, the GPU needs roughly 128–1024 threads (on current architectures) available inside these calls, but current CPUs generally need significantly fewer than this. Thus, having the ability to push more looping to expose more threading at the “vector” or “innermost cache” level (depending upon the architecture in question) is important. Therefore it is likely that the programmer may have to do at least one level of manual loop tiling for this purpose.

In the ACME atmosphere code, for instance, the parallelism comes in “chunks,” as it will in all models. We have the 4×4 basis functions that provide 16 loop indices. Then we have vertical levels that provide a factor of 72 more loop indices. For tracer transport, we have up to 50 tracers as the

next loop index. Finally, we have the element loop, which usually has no more than 64 elements in it for a node. We chose to manually block the vertical level loop, which flexibly allows the inner routines to either operate on 16 independent loop indices or more than 1,000. This should provide adequate caching and vectorization flexibility for the GPU and MIC.

Another part of this is that GPUs need to be able to be handled in a manner similar to loop blocking. No formal experiments have been performed at present, but it is likely that the K40 and above architectures can be effectively utilized by loop blocking in an intelligent manner. The major difference between GPUs and other architectures is the nature of the vectorization. On the GPU, the moment you are inside an SM, you are already technically vectorized. Every instruction is accessed in a breadth-wise manner, meaning the SM takes a single instruction and performs it on multiple data, and this is (as far as the programmer is concerned) performed across the entire vector loop (using OpenACC terminology). On a CPU or MIC, however, the formal vector units are significantly smaller. It is likely that most aspects of performance portability can be handled via tiling (mostly directive-based) and flexible exposing of independent loop indices through the callstack.

Models: OpenACC, OpenMP

New Capabilities: An effective, hierarchical OpenACC and OpenMP tiling clause

Resolution: N/A

I/O: N/A

Many-Core and/or GPU Readiness: Codes are partially ready for this. We can manually do some blocking in our looping structure to expose parallelism intelligently to reusable subroutines and keep data in cache for multiple types of accelerators. However, within routines, it would be nice to have the ability to tile loops with directives. OpenACC currently has the restriction that loops may only be tiled once (not multiple times or hierarchically). In addition, it has the restriction that the tiled loops will map to gangs (outer) and workers (inner). This is not flexible enough to work on a MIC, which has multiple levels of cache that could benefit from tiling.

Workflows: N/A

4. Software Applications, Libraries, and Tools: If you think your needs for HPC software (applications/libraries/tools/compilers/languages/etc.) will change between now and 2020–2025, please describe them here. Be sure to consider workflows, analytics, and I/O software.

5. HPC Services: If you anticipate needing additional HPC services not provided today in 2020–2025, what are they? Possibilities might include advanced training, data analytics and visualization assistance, support servers, collaboration tools, web interfaces, federated authentication services, gateways, etc.

6. Additional Needs: If there is anything else important to the success of your project that you have not mentioned, please do so here.

7. References (*please keep the reference length to no more than 10 critical references*)

New Numerical Methods for Efficient Exascale Simulation with Time-Explicit Codes (Matt Norman, ORNL)

There are many instances where the physics of the problem are relatively non-stiff for which time-explicit simulation is more efficient. Still, time-explicit methods have the well-known problem of time step reduction with added Degrees of Freedom (DOFs). Time step reduction is, for instance, the cause for the throughput problems in ACME atmosphere and ocean—both of which are now time-explicit. In 2-D, a 2x refinement leads to 8x more work, but there's only 4x more data. Thus, to keep throughput the same, the amount of data per node (also workload per node) reduces by at least 2x for every 2x improvement in spatial refinement. This leads to relatively higher MPI overheads. As an example, to run ACME atmosphere with 28-km globally averaged grid spacing between DOFs, we already have to spend roughly half of the time in MPI waiting to reach about 2 Simulated Years Per Day (SYPD) throughput. We need some means of increasing resolution while decreasing the workload per node less severely. There are two ways to do this: (1) change the approach to the science problem or (2) change the underlying mathematical approximation (or algorithm).

This time step reduction, however, differs widely depending upon (1) the spatial operator, (2) the temporal operator, and (3) whether the DOFs are added via h-refinement or p-refinement. Exascale computing rewards compute intensity, or the number of computations performed per amount of data moved. So, algorithms with higher compute intensity are desirable. One obvious way to obtain this is by using higher-order accuracy. The core computational cost of nearly all time-explicit methods is a matrix-vector product, usually performed in dimensional sweeps for efficiency. For Finite-Difference (FD) and Finite-Volume (FV) methods, this is reconstruction (also called recover). For Galerkin methods, this is quadrature summation. However, higher-order accuracy usually causes problems either in terms of time step or in terms of expense.

The problem with higher-order accuracy for FD and FV is prohibitive expense per time step. Unlike Galerkin and other multi-moment (MM) methods, FD and FV methods must reconstruct *every* DOF, whereas Galerkin methods merely reconstruct every element and every DOF in that element shares the same reconstruction. In D dimensions, Galerkin methods are N^1 times less expensive per time step than FV methods. The benefit for FV methods, however, is that they suffer *no* time step reduction with increasing order, though it does not appear to be enough to offset the extra cost compared to Galerkin methods, even when using GPUs. The problem for Galerkin methods is that the time step reduces nearly quadratically (relative to element size) with the order of accuracy, also acting as a barrier to higher-order accuracy. Galerkin methods do exhibit parallel benefits in that they require minimal parallel communication, needing only to communicate element boundary values. Within the realm of currently implemented Eulerian Galerkin methods, Spectral Element (SE) methods exhibit the best time steps, related to the averaging and sharing of element boundary DOFs rather than fluxing.

There are alternatives to the traditional Galerkin methods that are also multi-moment, however; and these alternatives sometimes give much larger time steps with the same amount of work and similar parallel data transfer properties. Examples of these are the various Multi-Moment Finite-Volume (MMFV) methods, which generally include Multi-Moment Constrained Finite-Volume (MCV) methods and others of the sort. These methods are FV methods with additional constraints, whether they be cell-averaged derivatives or point-wise derivatives or values. Often, many of these are shared from one element, not unlike the continuity imposed by SE methods. It is this *sharing* of boundary data that seems to lead to larger time steps.

However, many of these methods cannot feasibly operate beyond, say, fifth-order accuracy (obtained experimentally) because of oscillations. When the data is too non-smooth, many of these methods are too oscillatory to be useful. Sure, one can limit these oscillations, but when the original method is too oscillatory, resolution of non-smooth features does not improve with increasing orders of accuracy. This is because the basis functions (or effective basis functions) are themselves too

oscillatory. However, for a Galerkin method with a nodal Lagrange basis formed on a GLL grid, the bases are remarkably well-bounded, no matter the order of accuracy. For these methods, an increased order of accuracy always improves resolution of non-smooth features, and this is an attribute needed for most geophysical flows, which abound with non-smooth features.

For these reasons, and from much experimentation, some good properties of exascale algorithms are that they should (1) be multi-moment, (2) share boundary values between elements, (3) have well-bounded basis functions, and (4) use upwinding to determine the shared boundary values between elements. For instance, if using the “strong” form of SE (meaning, *without integration by parts*), one can then use upwinding to determine the shared element boundary DOFs. This gives a time step that is 2x larger than traditional SE. Also, one can overlap the bounding two DOFs instead of only one, and this, along with upwinding, gives a 3x larger time step than traditional SE. How conservation is maintained with these methods is open research.

Also, there are temporal algorithms that can perform higher-order simulation without stages, while maintaining non-oscillatory properties of an underlying limited interpolation. These are called ADER methods, and there are various methods of implementing them as well. With larger time steps than limited Runge-Kutta methods, they are ideal for exascale computing. Regarding limiting, Weighted Essentially Non-Oscillatory (WENO) ideology in limiting is preferred for exascale because it can, in a single application, remove oscillations to acceptable levels. WENO requires significantly less parallel communication than hyperdiffusion.

Clearly, there is much to be gained by investigating new algorithms, and there are significant unexplored areas for algorithms, both spatial and temporal. It would be well worth investment to try to obtain more efficient approaches to numerical approximation to better utilize the coming exascale machines for BER projects.

Describe a major science challenge expected to be solved in the 2020-2025 timeframe that requires using extant computing ecosystems.

New algorithms enable higher spatial resolution with better model throughput.

2. Current and Future Computational and Data Strategies

Models: N/SA

New Capabilities: New multi-moment algorithms

Resolution: Higher resolution

I/O: N/A

Many-Core and/or GPU Readiness: New high-order algorithms lead to more efficient GPU usage.

Workflows: N/A

4. Software Applications, Libraries, and Tools: N/A

5. HPC Services: Efficient GPU Direct or its equivalent for multi-accelerator nodes will be necessary

6. Additional Needs: N/A

References

- “Hermite WENO limiting for multi-moment finite-volume methods using the ADER-DT time discretization for 1-D systems of conservation laws,” Matthew Norman, *Journal of Computational Physics*, **282**, 2015, 381–396.
- “A WENO-limited, ADER-DT, finite-volume scheme for efficient, robust, and communication-avoiding multi-dimensional transport,” Matthew Norman, *Journal of Computational Physics*, **274**, 2014, 1–18.
- “Algorithmic Improvements for Schemes Using the ADER Time Discretization,” Matthew Norman, *Journal of Computational Physics*, **243**, 2013, 176–178.

APPENDIX D: BIOLOGICAL AND ENVIRONMENTAL RESEARCH (BER) CASE STUDIES

The following case studies were submitted by the following authors in advance of the BER Exascale Requirements Review to guide both the agenda and review discussions.

D.1 Case Studies Addressing Biological Research

Page No. ***Multiscale Biophysical Simulation from Molecules to Cells (corresponds to Section 3.1.1)***

D-3 X. Cheng (Oak Ridge National Laboratory)

D-7 Loukas Petridis and Jeremy Smith (both of Oak Ridge National Laboratory)

Page No. ***Microbes to the Environment (corresponds to Section 3.1.3)***

D-10 Senghwa Kang (Pacific Northwest National Laboratory)

D-14 R. Mahadevan (University of Toronto) and
C. Maranas (Pennsylvania State University)

D-18 C. Maranas (Pennsylvania State University)

Page No. ***Biological Big Data Challenges (corresponds to Section 3.1.4)***

D-23 Srinivas Aluru (Georgia Institute of Technology)

D-26 Ben Bowen, Oliver Ruebel, and Aydin Buluc
(all of Lawrence Berkeley National Laboratory)

D-31 Chirag Jain and Srinivas Aluru (both of Georgia Institute of Technology)

D-34 Nikos Kyrpides, Dan Rokhsar, Aydin Buluc, Rob Egan, Evangelos Georganas,
Leonid Olikier, Steven Hofmeyr, and Katherine Yelick
(all of Lawrence Berkeley National Laboratory)

D-38 Bernard Ng and Sara Mostafavi (both of Center for Molecular Medicine and
Therapeutics)

D-40 Chongle Pan (Oak Ridge National Laboratory)

D-43 H. Steven Wiley and Ronald Taylor (both of Pacific Northwest National Laboratory)

D.2 Case Studies Addressing Environmental Research

Page No. ***Software Engineering for Portability (corresponds to Section 3.2.7.3)***

D-47 Glenn Hammond (Sandia National Laboratories) and
Jeff Johnson (Lawrence Berkeley National Laboratory)

This page is intentionally left blank.

D.1 Case Studies Addressing Biological Research

Multiscale Biophysical Simulation from Molecules to Cells (corresponds to Section 3.1.1)

Case Study:

Lead Author(s): X. Cheng, Oak Ridge National Laboratory

1. Description of Research

1.1 Overview and Context: Please give a high-level description of your research and the roles played by high-end computing, storage, and networking.

My research is focused on computational biochemistry and biophysics, with applications in bioenergy and microbial-plant interfaces. In particular, molecular dynamics simulations are the main tool we use to probe the structure and dynamics of biological systems such as proteins, lignocellulosic biomass, and biomembranes. Molecular dynamics (MD) evolves the molecular systems by numerically solving Newton's equation of motion. Limited by the fastest vibration frequency, ~ 1 fs, MD simulation must use a time step comparable to or shorter than this vibration timescale to ensure the stability of the system's numerical integration. Therefore, to reach biologically relevant time scales, MD must be propagated for billions of steps and beyond, which requires each step to be completed within 1 ms, thus posing a great challenge for us. In this regard, high-end computing with massive parallelism plays an increasingly important role in research. Multi-level parallelism and fast inter-node communication will be key barriers to be overcome.

1.2 Research Objectives for the Next Decade: What are the scientific goals for the next decade in this area of research? What are the computational and data analysis/processing goals related to achieving the scientific goals?

In my view, the scientific goals for the next decade in my area of research are (1) a predictive (quantitative) understanding of biomolecular structure, dynamics and interaction; and (2) a systems-level understanding of biological systems from biomolecules to cells and beyond, including an understanding of the function of proteins in their native context. The associated computational and data analysis/processing goals are (1) the development of enhanced sampling techniques; (2) convergence quantification; (3) multiscale modeling, integrating modeling across different length and time scales; (4) parallel and automatic event detection and pattern recognition; and (5) combining items 4 and 2 and then with item 1.

2. Computational and Data Strategies

2.1 Approach: Give a short, high-level description of your computational and data problems and the strategies used to solve them today and how those might change in the next decade.

Multiscale and multiphysics models—The strategy is to bridge scales by describing different phenomena using models at different resolutions and/or using different physics laws. This strategy will continue to advance because this may be the only viable way more realistic modeling of an entire cell will become feasible in computer architectures of the next decade.

Scalable algorithms—The strategy is to develop numerical algorithms that can scale better on massively parallel and heterogeneous computing architectures. I expect to see more progress in this aspect in the next decade because MD simulation is inherently a fine-granularity parallelization (strong scaling) problem. New parallelization strategies are always needed to overcome the scaling barrier in future computer architectures.

Enhanced sampling techniques—The strategy is to enhance the barrier crossing rates by developing new enhanced sampling techniques. This area witnessed remarkable advances in late 1990s and early 2000s, but the pace has slowed down in recent years. In the future, a promising strategy will be to combine enhanced sampling techniques with scalable algorithms, and to enhance sampling through learning from historic data using machine learning concepts.

2.2 Codes and Algorithms: Please briefly describe your current codes and workflows and the algorithms that characterize them.

GROMACS/NAMD: Numerical solution of a second-order differential equation. The long-range electrostatic interaction is computed with an FFT-based Particle mesh Ewald (PME) method. The parallelization is based on a spatial decomposition algorithm.

3. Current and Future HPC Needs

Quantitative information requested in Section 3 must be put in a table following the template below.

3.1 Computational Hours: How many hours on conventional cores (no accelerators) do the codes involved in the study use now? If you use GPUs or other accelerators, use the next row for that answer. How do you expect your computational requirements to increase through 2020 and 2025? Include all hours you will need to reach the scientific goals you listed in Section 1.2 above.

3.2 Parallelism: Please describe the scale at which your codes use coarse-grained (multi-node) and fine-grained (on-node or accelerators) parallelism today. Please describe current plans to increase either level of parallelism.

MPI for internode communication and OpenMP for on-node communication.

3.3 Memory: Describe your current and future memory requirements in terms of the minimum shared memory pool (node) and aggregate memory required for you to run. Note that future systems may have much less memory as a function of peak performance than systems have today. Note also that the memory hierarchy will be potentially complex (on-chip fast memory and significantly slower to off-chip memory).

Memory is not a limiting factor for MD applications.

3.4 Scratch Data and I/O: How much online scratch storage space do you need for your runs (current/future), including checkpoint/restart data? Please estimate your I/O bandwidth requirement (bandwidth = data read or written/time to read or write). What percentage of your total runtime are you willing to devote to I/O?

We can devote <1%. This can also be controlled by the user to have less frequent I/O if I/O overhead is too large.

3.5 Long-term and Shared Online Data: How much active, online long-term storage do you need today and in 2020 and 2025? Please describe any requirements for sharing or accessing the data.

3.6 Archival Data Storage: Archival data is accessible online, but may involve a delay in accessing it (e.g., data stored on HPSS tapes). How much data do you have stored in a data archive currently? How much will you need in 2020 and 2025?

3.7 Workflows: Please briefly describe your current workflows and requirements for 2020 and 2025.

3.8 Many-Core and/or GPU Readiness: Future systems will contain “lightweight” cores and/or hardware accelerators (e.g., GPUs) with deepening memory hierarchies. Are your codes ready for this? If yes, please explain your strategy for exploiting these technologies. If not, what are your plans for dealing with such systems and what do you need to help you successfully transition to them?

3.9 Software Applications, Libraries, and Tools: If you think your needs for HPC software (applications/libraries/tools/compilers/languages/etc.) will change between now and 2020, 2025, please describe them here? Be sure to consider workflows, analytics, and I/O software.

3.10 HPC Services: If you anticipate needing additional HPC services not provided today in 2020–2025, what are they? Possibilities might include advanced training, data analytics and visualization assistance, support servers, collaboration tools, web interfaces, federated authentication services, gateways, etc.

3.11 Additional Needs: If there is anything else important to the success of your project that you have not mentioned, please do so here.

4 Requirements Summary Worksheet

| Code: _____ | Column 1: Current Usage | Future Usage: 2020 (As a factor of column 1)**** | Future Usage: 2025 (As a factor of column 1)**** |
|--|-------------------------------|---|---|
| Computational core hours (Conventional)* | 40 M | 100 M | 300 M |
| Computational node hours (Homogeneous many-core)** | N/A | | |
| Computational node hours (w/GPU or accelerator)*** | 20 M | 40 M | 100 M |
| Memory per node | 32 GB | 32 GB | 32 GB |
| Aggregate memory | 32 TB | 32 TB | 32 TB |
| Data read and written per run | 1 TB | 2 TB | 4 TB |
| Maximum I/O bandwidth needed | 100 GB/sec | 100 GB/sec | 100 GB/sec |
| Percent of runtime for I/O | <1% | | |
| Scratch file system space needed | 20 TB | 40 TB | 80 TB |
| Permanent online data storage | 100 TB | 200 TB | 400 TB |
| Archival data storage needed | 1000 TB | 2000 TB | 4000 TB |

*Please use “core hours” for “conventional” processors (i.e., node-hours * cores_per_node). Intel “Ivy Bridge” is an example conventional processor.

**Please use “node hours” for homogenous many-core architectures. A self-hosted Intel Xeon Phi “Knights Landing” is an example.

***Please use “node hours” for “GPU or accelerator” usage.

****E.g., 32× column 1.

Case Study Title: Atomistic simulation scaling to the size of a living cell

Lead Author(s): Loukas Petridis and Jeremy Smith, Oak Ridge National Laboratory

1. Description of Research

1.1 Overview and Context: High-performance computer simulation has a significant role to play in obtaining an understanding of the physical processes that lead to biological function. Understanding how structure and dynamics give rise to function requires spatiotemporal characterization spanning decades of time and length scales. The overarching aim of our research is to employ simulation and neutron-scattering techniques to obtain high-resolution spatial and temporal information on biological processes, and thus to demonstrate the role interactions between the members of complexes play in defining their function. Specifically, molecular dynamics (MD) simulations have provided atomic-level insight into processes at the core of bioenergy research.

1.2 Research Objectives for the Next Decade: Our research objective is to elucidate the complex interplay between the molecular systems within cells. This can be achieved by simulating, at the atomic scale, an entire living cell. To make this possible, two challenges must be addressed: building a realistic starting model of the cell and extending the timescales accessible to all-atom MD simulation.

2. Computational and Data Strategies

2.1 Approach: Molecular dynamics (MD) simulations involve stepwise integration of the equation of motion for a system of classical particles through an empirically derived potential energy function. A singular challenge is improving the strong scaling of MD simulations, which will allow extending the timescales accessible to all-atom MD simulation. This limitation can be addressed by both hardware (lower latency network) and software (improved task parallelism) improvements.

2.2 Codes and Algorithms: Current codes: GROMACS, NAMD. Algorithms: velocity Verlet, neighbor list; particle mesh Ewald (FFT).

3. Current and Future HPC Needs

Quantitative information requested in Section 3 must be put in a table following the template below.

3.1 Computational Hours: 100 M TITAN hours (3.3 node hours) for 2016. Require a factor of 400 increase by 2025.

3.2 Parallelism: Both multi-node and thread parallelism. Currently simulations scales up to 50% of TITAN (9,000 nodes)

3.3 Memory: Memory requirements are relatively small for biological MD.

3.4 Scratch Data and I/O: Currently need 1 TB scratch, check-pointing is about 1 GB. I/O bandwidth 10 MB/sec. I/O is currently not a limiting factor, but this may change with larger systems.

3.5 Long-term and Shared Online Data: Need about 2 TB currently, will increase by factors of 20 and 400 by 2020 and 2025, respectively.

3.6 Archival Data Storage: About 20 TB currently. This will increase by factors of 20 and 400 by 2020 and 2025, respectively.

3.7 Workflows: Please briefly describe your current workflows and requirements for 2020 and 2025.

3.8 Many-Core and/or GPU Readiness: FCodes are already using GPUs and there is ongoing development to use Intel MIC.

Answers to the following questions are not required as these issues will be dealt with more globally in the white papers. However, if you have a unique situation that you feel will need to be addressed, please do so. Keep it short!

3.9 Software Applications, Libraries, and Tools: If you think your needs for HPC software (applications/libraries/tools/compilers/languages/etc.) will change between now and 2020, 2025, please describe them here? Be sure to consider workflows, analytics, and I/O software.

3.10 HPC Services: If you anticipate needing additional HPC services not provided today in 2020–2025, what are they? Possibilities might include advanced training, data analytics and visualization assistance, support servers, collaboration tools, web interfaces, federated authentication services, gateways, etc.

3.11 Additional Needs: If there is anything else important to the success of your project that you have not mentioned, please do so here.

4 Requirements Summary Worksheet

| Code: ___GROMACS_____ | Column 1: Current Usage | Future Usage: 2020 (As a factor of column 1)**** | Future Usage: 2025 (As a factor of column 1)**** |
|--|-------------------------------|---|---|
| Computational core hours (Conventional)* | | | |
| Computational node hours (Homogeneous many-core)** | | | |
| Computational node hours (w/GPU or accelerator)*** | 3.3 M | 20 | 400 |
| Memory per node | <1 GB | 5 GB | 25 GB |
| Aggregate memory | TB | TB | TB |
| Data read and written per run | 0.1 TB | 2 TB | 40 TB |
| Maximum I/O bandwidth needed | 0.01GB/sec | 0.2GB/sec | 4 GB/sec |
| Percent of runtime for I/O | <1% | | |
| Scratch file system space needed | 1 TB | 20 TB | 400 TB |
| Permanent online data storage | 5 TB | 100 TB | 1000 TB |
| Archival data storage needed | 20 TB | 400 TB | 8000 TB |

*Please use “core hours” for “conventional” processors (i.e., node-hours * cores_per_node). Intel “Ivy Bridge” is an example conventional processor.

**Please use “node hours” for homogenous many-core architectures. A self-hosted Intel Xeon Phi “Knights Landing” is an example.

***Please use “node hours” for “GPU or accelerator” usage.

****E.g., 32× column 1.

Microbes to the Environment (corresponds to Section 3.1.3)

Case Study Title: Microbial Community Formation, Maintenance, and Disruption in Complex Natural and Artificial Environments

Lead Author(s): Senghwa Kang, Pacific Northwest National Laboratory

1. Description of Research

1.1 Overview and Context: Please give a high-level description of your research and the roles played by high-end computing, storage, and networking.

This case study aims to mathematically model and simulate microbial communities in complex natural (e.g., soil aggregates) and artificial (e.g., bioreactors) environments. Researchers can currently simulate an individual cell (or a few cells) in high resolution and high fidelity, or they can simulate a large number of cells in low resolution while adopting computationally inexpensive mathematical models and sacrificing fidelity. This case study aims to integrate high-fidelity individual cell models and direct (via physical contact) and indirect (via complex environments) cell interaction models to model and simulate microbial communities of billions of cells or more. Even the run-time of a single cell model (on a single core and using the code written by domain experts) ranges from minutes to hours (and sometimes days). Simulating billions of cells and their interactions requires performance-optimized software and high-performance computers that exceed the capacity of current leadership-class supercomputers.

1.2 Research Objectives for the Next Decade: What are the scientific goals for the next decade in this area of research? What are the computational and data analysis/processing goals related to achieving the scientific goals.

We aim to model and simulate microbial community formation, maintenance, and disruption to understand key driving mechanisms, predict their behaviors in reaction to external perturbations, and discover techniques to manipulate their behaviors.

To achieve this goal, we need to simulate billions of cells and their interactions with high fidelity. Individual cell models and interaction models with varying temporal and spatial scales are highly coupled with complex interdependencies as well. Scaling this multiscale computation to large supercomputers is a significant computational challenge.

2. Computational and Data Strategies

2.1 Approach: Give a short, high-level description of your computational and data problems and the strategies used to solve them today and how those might change in the next decade.

We should simulate individual cell behaviors, run range-limited N-body computations (to simulate interactions via direct contact), and solve multiple partial differential equations (e.g., to simulate diffusion and advection of various molecular species in complex environments). These sub-models are highly coupled. As modeling further matures over the next 10 years, we need to simulate individual cells with more diverse and complex mathematics, and solve a larger number of partial differential equations to incorporate a larger number of molecular species. This will increase compute density (which lowers the communication burden) and irregularity (which challenges load balancing).

2.2 Codes and Algorithms: Please briefly describe your current codes and workflows and the algorithms that characterize them.

We have developed Biocellion to simulate a large number of cells and their interactions. Biocellion currently simulates billions of cells (but with simple single cell models) and their interactions. Biocellion integrates individual agent simulations, range-limited N-body interactions, and partial differential equations. Paraview is used for visualization.

3. Current and Future HPC Needs

Quantitative information requested in Section 3 must be put in a table following the template below.

3.1 Computational Hours: How many hours on conventional cores (no accelerators) do the codes involved in the study use now? If you use GPUs or other accelerators, use the next row for that answer. How do you expect your computational requirements to increase through 2020 and 2025? Include all hours you will need to reach the scientific goals you listed in 1.2 above.

3.2 Parallelism: Please describe the scale at which your codes use coarse-grained (multi node) and fine-grained (on-node or accelerators) parallelism today. Please describe current plans to increase either level of parallelism.

Biocellion exploits both coarse-grained (runs multiple MPI processes) and fine-grained (using Intel Thread Building Block) parallelism. Currently, single-cell models are expected to run in sequential order. We plan to further increase parallelism by parallelizing expensive solvers to simulate individual cell behaviors.

3.3 Memory: Describe your current and future memory requirements in terms of the minimum shared memory pool (node) and aggregate memory required for you to run. Note that future systems may have much less memory as a function of peak performance than systems have today. Note also that the memory hierarchy will be potentially complex (on-chip fast memory and significantly slower to off-chip memory).

3.4 Scratch Data and I/O: How much online scratch storage space do you need for your runs (current/future), including checkpoint/restart data? Please estimate your I/O bandwidth requirement (bandwidth = data read or written / time to read or write). What percentage of your total runtime are you willing to devote to I/O?

3.5 Long-term and Shared Online Data: How much active, online long-term storage do you need today and in 2020 and 2025? Please describe any requirements for sharing or accessing the data.

3.6 Archival Data Storage: Archival data is accessible online, but may involve a delay in accessing it (e.g., data stored on HPSS tapes). How much data do you have stored in a data archive currently? How much will you need in 2020 and 2025?

3.7 Workflows: Please briefly describe your current workflows and requirements for 2020 and 2025.

Biocellion runs simple data analytics online while relying on Paraview for visualization. A larger-scale visualization and data analytics capability is necessary because the scale and fidelity of the simulation increases over time. Its visualization and data analytics can be performed online without requiring files to interface; this will significantly reduce I/O requirements.

3.8 Many-Core and/or GPU Readiness: Future systems will contain “lightweight” cores and/or hardware accelerators (e.g., GPUs) with deepening memory hierarchies. Are your codes ready for this? If yes, please explain your strategy for exploiting these technologies. If not, what are your plans for dealing with such systems and what do you need to help you successfully transition to them?

We are currently updating Biocellion to run on Intel Xeon Phi processors.

Answers to the following questions are not required as these issues will be dealt with more globally in the white papers. However, if you have a unique situation that you feel will need to be addressed, please do so. Keep it short!

3.9 Software Applications, Libraries, and Tools: If you think your needs for HPC software (applications/libraries/tools/compilers/languages/etc.) will change between now and 2020, 2025, please describe them here? Be sure to consider workflows, analytics, and I/O software.

3.10 HPC Services: If you anticipate needing additional HPC services not provided today in 2020–2025, what are they? Possibilities might include advanced training, data analytics and visualization assistance, support servers, collaboration tools, web interfaces, federated authentication services, gateways, etc.

3.11 Additional Needs: If there is anything else important to the success of your project that you have not mentioned, please do so here.

4. Requirements Summary Worksheet

Please fill out the following table to the best of your ability prior to the meeting. If you are not able to make an estimate, leave the entry blank and we will discuss it at the meeting.

| Code: Biocellion | Column 1: Current Usage | Future Usage: 2020 (As a factor of column 1)**** | Future Usage: 2025 (As a factor of column 1)**** |
|--|----------------------------|--|--|
| Computational core hours (Conventional)* | 100,000 | 10,000,000 | 1,000,000,000 |
| Computational node hours (Homogeneous many-core)** | | | |
| Computational node hours (w/GPU or accelerator)*** | | | |
| Memory per node | 32 GB | 256 GB | 1024 GB |
| Aggregate memory | 32 TB | 256 TB | 1024 TB |
| Data read and written per run | 10 TB | 80 TB | 320 TB |
| Maximum I/O bandwidth needed | 5 GB/sec | 20 GB/sec | 50 GB/sec |
| Percent of runtime for I/O | Less than 5% | Less than 1% | Less than 1% |
| Scratch file system space needed | 10 TB | 80 TB | 320 TB |
| Permanent online data storage | 0 TB | 0 TB | 0 TB |
| Archival data storage needed | 0 TB | 0 TB | 0 TB |

*Please use “core hours” for “conventional” processors. (i.e., node-hours * cores_per_node). Intel “Ivy Bridge” is an example conventional processor.

**Please use “node hours” for homogenous many-core architectures. A self-hosted Intel Xeon Phi “Knights Landing” is an example.

***-Please use “node hours” for “GPU or accelerator” usage.

****E.g., 32× column 1.

Case Study Title: Integrated Analysis and Design of Metabolism in Cells and Communities

Lead Author(s): R. Mahadevan, University of Toronto; and C. Maranas, Pennsylvania State University

1. Description of Research

1.1 Overview and Context: Please give a high-level description of your research and the roles played by high-end computing, storage, and networking.

Genome-scale metabolic models and communities have enabled the simulation of metabolism and communities. Typical workflows for such simulations involve the following steps: (1) significant sequence analysis to develop models (enabled by KBase), (2) model validation workflows by integrating genome-wide data, (3) optimization methods for redesigning metabolism in cells, and (4) dynamic multi-level optimization for analysis of microbial community metabolism in homogenous environments, and (5) spatial-temporal modeling of metabolism in microbial communities (e.g., biofilm, field-scale simulations). In addition, genome-scale models have been extended to include transcriptional and translation (ME models which require significant additional computational requirements).

1.2 Research Objectives for the Next Decade: What are the scientific goals for the next decade in this area of research? What are the computational and data analysis/processing goals related to achieving the scientific goals.

Section 1 should be about half a page (12pt).

An important extension in the future will be the development of computational efficient multi-scale modeling methods for describing microbial community dynamics in order to simulate multiple species (>10) in spatially heterogeneous environments. This advance will require close integration with other disciplines that are developing advanced reactive transport modeling codes. Although methods exist for the analysis and design of metabolism in cells and communities, most of these methods use a simple description of metabolism alone in cells and communities. In contrast, there have been significant advances in describing metabolism including transcription, translation, and other cellular process, as well as kinetics [PMID:24084808, 22817898, 26474788]), which require significant computational and data requirements for single cellular level simulations. Integrating such representations with the design and multi-scale methods for optimization will require the development of novel computationally efficient methods. In addition, integrating these models with genome-wide data will also require new computational methods in data mining and analysis.

2. Computational and Data Strategies

2.1 Approach: Give a short, high-level description of your computational and data problems and the strategies used to solve them today and how those might change in the next decade.

Computational problems: Combinatorial growth associated with integer-based optimization formulation, as well as kinetics and model-reduction

Solution Approach: Parallel computing for reactive transport models

Data problems: Standards for biological entities

2.2 Codes and Algorithms: Please briefly describe your current codes and workflows and the algorithms that characterize them.

Sections 1 and 2 together **must be no more than 1 page in length.**

Codes: LP Solvers (CPLEX)

Workflows: KBase: (a) genome-scale model development, (b) metabolic network design methods, (c) multi-level optimization of microbial communities, (d) integrated reactive transport modeling with metabolic models.

Algorithms: Lots here ???

3. Current and Future HPC Needs

Quantitative information requested in Section 3 must be put in a table following the template below.

3.1 Computational Hours: How many hours on conventional cores (no accelerators) do the codes involved in the study use now? If you use GPUs or other accelerators, use the next row for that answer. How do you expect your computational requirements to increase through 2020 and 2025? Include all hours you will need to reach the scientific goals you listed in 1.2 above.

3.2 Parallelism: Please describe the scale at which your codes use coarse-grained (multi-node) and fine-grained (on-node or accelerators) parallelism today. Please describe current plans to increase either level of parallelism.

3.3 Memory: Describe your current and future memory requirements in terms of the minimum shared memory pool (node) and aggregate memory required for you to run. Note that future systems may have much less memory as a function of peak performance than systems have today. Note also that the memory hierarchy will be potentially complex (on-chip fast memory and significantly slower to off-chip memory).

3.4 Scratch Data and I/O: How much online scratch storage space do you need for your runs (current/future), including checkpoint/restart data? Please estimate your I/O bandwidth requirement (bandwidth = data read or written/time to read or write). What percentage of your total runtime are you willing to devote to I/O?

3.5 Long-Term and Shared Online Data: How much active, online long-term storage do you need today and in 2020 and 2025? Please describe any requirements for sharing or accessing the data.

3.6 Archival Data Storage: Archival data is accessible online, but may involve a delay in accessing it (e.g., data stored on HPSS tapes). How much data do you have stored in a data archive currently? How much will you need in 2020 and 2025?

3.7 Workflows: Please briefly describe your current workflows and requirements for 2020 and 2025.

3.8 Many-Core and/or GPU Readiness: Future systems will contain “lightweight” cores and/or hardware accelerators (e.g., GPUs) with deepening memory hierarchies. Are your codes ready for this? If yes, please explain your strategy for exploiting these technologies. If not, what are your plans for dealing with such systems and what do you need to help you successfully transition to them?

Answers to the following questions are not required as these issues will be dealt with more globally in the white papers. However, if you have a unique situation that you feel will need to be addressed, please do so. Keep it short!

3.9 Software Applications, Libraries, and Tools: If you think your needs for HPC software (applications/libraries/tools/compilers/languages/etc.) will change between now and 2020, 2025, please describe them here? Be sure to consider workflows, analytics, and I/O software.

3.10 HPC Services: If you anticipate needing additional HPC services not provided today in 2020-2025, what are they? Possibilities might include advanced training, data analytics and visualization assistance, support servers, collaboration tools, web interfaces, federated authentication services, gateways, etc.

3.11 Additional Needs: If there is anything else important to the success of your project that you have not mentioned, please do so here.

4. Requirements Summary Worksheet

Please fill out the following table to the best of your ability prior to the meeting. If you are not able to make an estimate, leave the entry blank and we will discuss it at the meeting.

| Code: _____ | Column 1: Current Usage | Future Usage: 2020 (As a factor of column 1)**** | Future Usage: 2025 (As a factor of column 1)**** |
|--|-------------------------------|---|---|
| Computational core hours (Conventional)* | | | |
| Computational node hours (Homogeneous many-core)** | | | |
| Computational node hours (w/GPU or accelerator)*** | | | |
| Memory per node | GB | GB | GB |
| Aggregate memory | TB | TB | TB |
| Data read and written per run | TB | TB | TB |
| Maximum I/O bandwidth needed | GB/sec | GB/sec | GB/sec |
| Percent of runtime for I/O | | | |
| Scratch file system space needed | TB | TB | TB |
| Permanent online data storage | TB | TB | TB |
| Archival data storage needed | TB | TB | TB |

*Please use “core hours” for “conventional” processors. (i.e., node-hours * cores_per_node). Intel “Ivy Bridge” is an example conventional processor.

**Please use “node hours” for homogenous many-core architectures. A self-hosted Intel Xeon Phi “Knights Landing” is an example.

***Please use “node hours” for “GPU or accelerator” usage.

****E.g., 32X column 1.

Case Study Title: Design and analysis of metabolic pathways

Lead Author(s): Costas D. Maranas

1. Description of Research

1.1 Overview and Context: At Maranas group, we are working on the development of modeling and optimization techniques for the analysis and (re)design of metabolic networks in various microorganisms. To this end, we have developed a number of algorithms such as optKnock, optForce, k-optForce, optStoic etc, and metabolic model reconstruction workflows for seamless synthetic design of microbial factories.

1.2 Research Objectives for the Next Decade: Construction of high resolution whole cell models that enable (i) accurate predictions of metabolic phenotypes upon genetic/environmental perturbations (ii) identification of practical genetic and protein engineering strategies for the production of “designer molecules”.

2. Computational and Data Strategies

2.1 Approach: In our exercise to predict a number of metabolic phenotypes, the computational complexity is highest for predicting enzyme kinetic parameters. The runtime scales quadratically to the fitted omics data. With the exponential increase in measured omics, a move towards exascale computing is ever more apparent.

Moving beyond the scope of biological functions catalogued in existing data sources, pathway-prospecting tools would be used to pinpoint desirable enzyme substrate and/or cofactor activity changes, and identify engineering strategies for the production of any target “designer molecule” in the future. This challenge expands the combinatorial complexity and search space of pathway design, and would therefore require weeks, if not months of simulation runtime to identify veritable solutions.

2.2 Codes and Algorithms:

| Tool | Algorithm | Software | Predictions/Design objective |
|-------------------------|-----------------------------|----------|--|
| Metabolic flux analysis | NLP/Active-set optimization | Matlab | Flux spans consistent with 13-C labeling |
| Kinetic modeling | NLP/Genetic algorithm | Matlab | Fluxomics consistent kinetic parameters |
| OptForce | MILP/Branch and Bound | Cplex | Worst case genetic interventions |
| k-OptForce | MILP/Branch and Reduce | BARON | OptForce interventions w/ kinetic parameters |
| optStoic/minRxn | MILP/Branch and Bound | Cplex | Stoichiometry consistent metabolic pathway |

3. Current and Future HPC Needs

Quantitative information requested in Section 3 must be put in a table following the template below.

3.1 Computational Hours: All the specifics below are hours per core, on a multicore cpu e.g. Intel "Ivy Bridge."

| Tool | 2016 | 2020 | 2025 | notes |
|-------------------------|--------|---------|-----------|--|
| Metabolic flux analysis | 300 | 2,000 | 5,000 | Compute time proportional to size of model |
| Kinetic modeling | 60,000 | 500,000 | 1,000,000 | Compute time proportional to # of omics measurements |
| OptForce | 48 | 168 | 672 | Compute time proportional to size of model |
| k-OptForce | 60,000 | 500,000 | 1,000,000 | Compute time proportional to # of omics measurements |
| optStoic/minRxn | 48 | 100 | 200 | Compute time proportional to size of database |

3.2 Parallelism: Coarse-grained (multi node) and fine-grained (on-node or accelerators) parallelism of current procedures.

| Tool | Multi-node | Accelerators |
|-------------------------|------------|--------------|
| Metabolic flux analysis | no | None |
| Kinetic modeling | yes | None |
| OptForce | no | None |
| k-OptForce | yes | None |
| optStoic/minRxn | no | None |

3.3 Memory: The current and future memory requirements in terms of the minimum shared memory pool (node) and aggregate memory required per run as of 2016.

| Tool | Node memory | Aggregate memory |
|-------------------------|-------------|------------------|
| Metabolic flux analysis | 12 GB | 12 GB |
| Kinetic modelling | 10 GB | 600 GB |
| OptForce | 12 GB | 12 GB |
| k-OptForce | 10 GB | 600 GB |
| optStoic/minRxn | 20 GB | 20 GB |

3.4 Scratch Data and I/O: None of the procedures we mention are network I/O intensive. The only parallel components, in kinetic modelling and k-OptForce procedures are SIMD (Single instruction multiple data). Based on current usage statistics recorded on Penn State HPC, a 1 GBPS or higher, low-latency Infiniband network would be sufficient. The checkpoint and restart data is comparable to the snapshot memory requirement of around 1.5 TB aggregate. Scratch data for backup and logging is expected to be between 20 to 30 TB

3.5 Long-term and Shared Online Data: Online long-term in 2020 and 2025 is around 30TB and 100 TB respectively. A secure online FTP service will be required to provide public access to shared data.

3.6 Archival Data Storage: Archival data on HPSS tapes in 2020 and 2025 is around 30TB and 100 TB respectively.

3.7 Workflows:

Kinetic modeling workflow is based on a genetic-algorithm (GA) based strategy that selects the optimal combination of the sampled parameters in the ensemble. We parallelized the fitness evaluation of the chromosomes in each generation. A slave processor is assigned to each chromosome to calculate its goodness-of-fit by integrating the system of ordinary differential equations (ODEs) of mass conservation. The fitness of each chromosome is passed to a master processor to identify the elite chromosome for populating the next generation. We believe more efficient parallel implementation of both gradient and machine learning-based algorithms will be required to accelerate the convergence.

In the current scope of “optStoic”, we first search through a dataset of metabolites (~30,000 metabolites) to identify optimal combination of carbon substrates and products that maximize a target objective (e.g., maximize overall profit of production). Subsequently, for this feasible overall conversion, we search through a database of reactions (~40,000 reactions) to find a minimum network of reactions that converts the substrates to the target metabolites. These MILP problems are solved by deterministic optimization solvers, which in future for large metabolite and reaction datasets may become computationally intractable. Alternate evolutionary algorithms will be required to handle problem complexity in such cases.

The current metabolic flux analysis (MFA) workflow employs a gradient-based optimization method to estimate flux spaces. At each iteration, the goodness-of-fit is evaluated using a variance-weighted sum of square of deviation of predicted and experimentally measured metabolite labeling patterns. The predicted metabolite labeling patterns are obtained as the solution to a square system of equations for steady-state MFA and a system of ODEs for instationary MFA for a given flux distribution. We believe that the scalability issue can be addressed by leveraging network topological features and model simplification to minimize the number of fluxes to be evaluated. We will use faster LP-based methods to evaluate the remaining fluxes.

3.8 Many-Core and/or GPU Readiness: None of the procedures use hardware accelerators (e.g., GPUs). This is primarily due to the underlying software package, which have been designed based on the current CPU architecture. Physiologically relevant fluxes in the range of 10^{-9} to 10^{-10} need to be resolved along with other fluxes of order 1. The difficulties become prevalent when the fast-timescale reactions in metabolism are coupled to slow timescale reactions in macromolecular synthesis. Most CPU architectures with IEEE 64-bit fps, with double precision, unlike most single precision GPU's are capable of addressing such scaling issues. In addition, GPU based gradient search methods are still rudimentary. In alliance with Penn State HPC, one of the lab members, worked on an implementation of the gradient-based Newton-Raphson optimization procedure using CUDA libraries. The application initially parses the conservation of mass equations using the Boost::Spirit C++ framework, finds an analytic Jacobian \mathbf{J} , and then iteratively updates the best solution with Δ by solving the linear system $\mathbf{J} \cdot \Delta = -\mathbf{F}$, where \mathbf{F} corresponds to the numerical evaluation of the set of equations. The linear system is solved using the GMRES algorithm from the CUSP library, which is dedicated to computations and algorithms for sparse matrices on GPU. Successive updates of the parameter set, Jacobian matrix, and functions, as well as the system solver are all implemented on GPU. Similar effort will be carried out for GPU implementation of GA algorithm.

Once developed, such GPU methods is envisaged to greatly reduce the runtime of kinetic modeling and k-OptForce procedures.

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*.

3.10 HPC Services: Hardware and software maintenance services is required.

3.11 Additional Needs: none

Requirements Summary Worksheet

| Code: _____ | Column 1: Current Usage | Future Usage: 2020 (As a factor of column 1)**** | Future Usage: 2025 (As a factor of column 1)**** |
|---|-------------------------------|---|---|
| Computational core hours (Conventional)* | 120,000 | 1,200,000 | 2,500,000 |
| Computational node hours (Homogeneous many-core)** | 10,000 | 100,000 | 200,000 |
| Computational node hours (w/GPU or accelerator)*** | unknown | unknown | unknown |
| Memory per node | 70GB | 100GB | 200GB |
| Aggregate memory | 1.5TB | 2TB | 4TB |
| Data read and written per run | 1TB | 1TB | 1TB |
| Maximum I/O bandwidth needed | 1GB/sec | 1GB/sec | 1GB/sec |
| Percent of runtime for I/O | ~1% | ~1% | ~1% |
| Scratch file system space needed | 10TB | 20TB | 30TB |
| Permanent online data storage | 10TB | 30TB | 100TB |
| Archival data storage needed | 10TB | 30TB | 100TB |

*-Please use “core hours” for “conventional” processors. (i.e., node-hours * cores_per_node). Intel “Ivy Bridge” is an example conventional processor.

** -Please use “node hours” for homogenous many-core architectures. A self-hosted Intel Xeon Phi “Knights Landing” is an example.

***-Please use “node hours” for “GPU or accelerator” usage.

****- e.g., 32X column 1.

Biological Big Data Challenges (corresponds to Section 3.1.4)

Case Study Title: Reverse Engineering Genome-Scale Networks

Lead Author(s): Srinivas Aluru, Georgia Institute of Technology

1. Description of Research

1.1 Overview and Context

Reverse engineering genome-scale networks is a grand challenge in systems biology. A vast and growing number of gene expression datasets available in public repositories through reported microarray or RNA-seq experiments provide a rich resource for genome-scale network construction and analysis. Several mathematical modeling techniques have been developed and applied for constructing gene networks. These include simple pairwise correlation techniques such as Pearson correlation, more advanced techniques such as Gaussian Graphical Modeling, and machine-learning approaches such as mutual information and Bayesian networks. In general, high-quality models tend to be more complex and involve a higher computational burden. In particular, Bayesian network reconstruction is an NP-hard problem, but Bayesian networks can effectively model combinatorial interactions.

1.2 Research Objectives for the Next Decade

The broad goal of this research is to develop parallel Bayesian network structure learning methods that can model genome-scale networks. Because complex organisms can contain a few tens of thousands of genes, and the number of gene expression datasets already number in the tens of thousands, this is a computationally demanding challenge. The run-time of exact Bayesian learning grows exponentially as the number of genes increases, and the scoring methodology used determines the computational cost with respect to the number of experiments. The need to construct such networks for various categories of experiments, the availability of multiple scoring functions, parameterization of the scoring functions, and statistical techniques such as bootstrapping all point to the need for multiple runs, even when building genome-scale networks of a single organism.

2. Computational and Data Strategies

2.1 Approach: Because Bayesian network construction is NP-hard, only heuristic methods are feasible for large-scale networks. Our approach to this problem consists of constructing a sparse undirected skeleton network with the goal of achieving sparsity while not missing any gene interactions, and computing the optimal parents for each gene in the network while constraining the set of choices by the skeleton network. For sets of a size below a threshold, exhaustive enumeration is used to select the optimal parent sets, while heuristics must be relied on for larger sets. Massive parallelization can be used to push the threshold below which exhaustive enumeration can be carried out. For example, we were able to exactly solve the problem for set sizes up to 32 on the Tianhe-2 supercomputer.

2.2 Codes and Algorithms: We have implemented the parallel software in C+ MPI, with the capability to use one or more Xeon Phi accelerators per node (Misra et al., Supercomputing 2014 best paper finalist). We also developed a dynamic load-balancing scheme that was spun off as standalone software DTree.

3. Current and Future HPC Needs

3.1 Computational Hours: We constructed genome-scale networks of the model plant *Arabidopsis thaliana* from 11,760 microarray experiments. We reported a strong scaling efficiency of 75% on 1.57 million cores of the Tianhe-2 supercomputer (counting both Xeon CPU and Xeon Phi cores). The software can use the maximum available compute and time resources by dialing up the set sizes for which exhaustive enumeration is used. Thus, the quality of the network can be improved with available resources.

3.2 Parallelism: The software can exploit both distributed and shared memory parallelism effectively.

3.3 Memory: The input to the algorithm is a gene expression matrix with dimension sizes corresponding to the number of genes and number of experiments. The matrix may have up to a few hundred million expression values, but can still be easily stored within the main memory of each node. The algorithm explores an exponential search space for selecting optimal parents for each gene, but only stores the best-scoring sets. Thus, memory footprint is not a bottleneck. If multiple high-scoring sets need to be stored, this can increase the memory requirement.

3.4 Scratch Data and I/O: The gene expression matrix is created by extracting data from public repositories, and carrying out preprocessing steps including data normalization. The matrix is read by the network construction software at the beginning of its execution.

3.5 Long-term and Shared Online Data: Gene expression data is already available in public repositories. Constructed networks can be easily shared through a Web portal with minimal backend resources.

3.6 Archival Data Storage: None needed, as expression data is already archived.

3.7 Workflows: None.

4 Requirements Summary Worksheet

The code is highly scalable, and can use higher compute capacity to improve the quality of answer. The present code can utilize Xeon Phi co-processors, but GPU versions can be easily developed. Exascale systems can be used to increase the threshold within which the problem is exactly solved, beyond which heuristics methods will be applied.

Case Study Title: From alchemy to chemistry: an extreme case of creation and searching chemical networks.

Author(s): Ben Bowen, Oliver Ruebel, and Aydin Buluc (Lawrence Berkeley National Laboratory)

1. Description of Research

1.1 Overview and Context: Because of the sensitivity of mass spectrometry (MS), tandem mass spectra are often the first (or only) data obtained on unknown compounds in complex samples. Tandem mass spectrometry involves fragmentation of analyte-derived ions via a sudden infusion of energy into the molecule. MS can detect hundreds of compounds at high sensitivity from complex mixtures, and mass spectra are often the first data available on unknown samples in microbial ecology, metabolomics, and analysis of metabolism for synthetic biology. Although the metabolome has been an extensive target of study since the dawn of biochemistry, the number of natural products on Earth is unknown; likewise, methods for definitively identifying them are lacking. Ideal approaches for the exhaustive mapping of relationships within and between molecules are likely not readily achievable with today's computing technology.

1.2 Research Objectives for the Next Decade: The most reliable approach to matching unknown spectra to compounds relies on spectral databases. In this approach, the tandem mass spectra of a large number of pure authentic standards is acquired and compared to unknown spectra. This approach is useful only when pure standards are available. **Unfortunately, these standards are typically not available commercially.** Many tools have been developed to infer the structural features of unknown compounds on the basis of tandem mass spectra. These include (i) calculations based on chemical principles such as quantum chemistry, (ii) spectral tree methods, and (iii) databasing of spectra of pure compounds. Each of these three tools is amenable to being efficiently applied in an integrated manner, but computational approaches for efficiently achieving this are lacking.

Objective 1: Tree generation. *In silico* fragmentation trees are the complete enumeration of all fragmentation possibilities and the lineage relating each fragment back to the original structure. Consequently, generating trees is compute intensive for large molecules. To overcome the barrier associated with tree generation for large molecules, new parallel algorithms are needed. For example, one possible approach to reduce redundant calculations would identify common substructures of molecules and reuse them to create other fragmentation subtrees via tree crafting. Improving scalability will require parallelization of the calculation of fragmentation trees for individual molecules, for example, by calculating subtrees for different fragments independently and merging the subtrees into complete trees.

Objective 2: Create the graph to maintain knowledge of molecules. A graph relating all molecular fragmentation trees would group molecules and their fragments around common substructures. Consequently, this graph would exchange an increase in edges for a decrease in nodes. Due to the decrease in the number of nodes (substructures), this is attractive for performing theoretical calculations and QC molecular simulations to describe bond energies, protonation affinities, and other molecular properties for molecules and their substructures. In addition, common substructures would yield an easy to use metric for molecular similarity. Last, this graph would be a place to maintain fragmentation observations from mass spectrometry measurements.

Objective 3: Interpretation and improvement of scoring results. Scoring requires algorithms, methods for visualization of fragmentation graphs in conjunction with scoring results, and a framework for efficiently querying large graphical databases. The algorithms would use theoretical estimates of bond energies, rearrangements, and protonation energies along with learning real/bogus likelihoods from measurements. Visualizations would allow scientists to traverse the fragmentation tree in concert with the measured fragmentation spectra, and high--performance computing frameworks and interfaces would lower the barrier for querying and appending the database.

2. Computational and Data Strategies

The Metabolite Atlas and OpenMSI projects use the high--performance computing at NERSC to accelerate compound identification for mass spectrometry. The objectives described above are in development, but require significantly more research and computational resources than what is currently in use.

2.1 Current Approach: For our naive approach, identification is based on comparisons of measured spectra to theoretically possible fragmentation paths for known molecular structures. Defining these templates, fragmentation path references are major computational challenges; to date, we have computed and stored complete fragmentation trees to a depth of five consecutive bond disassociations for more than 11,000 compounds. Metabolite Atlas and OpenMSI users are now searching their raw spectra against these trees and getting results in minutes. Without supercomputing, these tasks would take months or would not be performed at all. In addition, the real-time queue enables users to make better use of their time by avoiding the highly variable wait times previously experienced on the normal queue.

2.2 Codes and Algorithms: To enable the efficient calculation of large numbers of fragmentation trees, we parallelized the computation using MPI. We here use a dynamic, work-request scheduling scheme to load balance calculations and account for the large variability in time to generate fragmentation trees. By sorting compounds in ascending order of number of bonds and/or dividing compounds into groups of similar numbers of bonds, we can further improve the ability of the dynamic scheduler to load balance calculations and improve parallel efficiency and utilization of compute resources. Using NERSC allowed us to generate all fragmentation trees in parallel using >1000 cores in less than a day.

3. Current and Future HPC Needs

Current usage has been moderate, largely due to the fact that early efforts have focused on small molecules (~50 atoms) and on the calculation of fragmentation trees only. As we move to larger molecules with hundreds of atoms, compute requirements are expected to increase dramatically by several orders of magnitude. In addition, the addition of first-principle calculations and QC molecular simulations to compute bond energies, protonation affinities, and other molecular properties and *in silico* fragmentation spectra is expected to further dramatically increase compute, memory, and storage needs.

| Code: Network-based compound identification | Column 1: Current Usage | Future Usage: 2020 (As a factor of column 1)**** | Future Usage: 2025 (As a factor of column 1)**** |
|---|-------------------------|---|---|
| Computational core hours (Conventional)* | O(100K) | 400X | 800X |
| Computational node hours (Homogeneous many-core)** | -- | -- | -- |
| Computational node hours (w/GPU or accelerator)*** | -- | -- | -- |
| Memory per node | 20 GB | 10X | -- |
| Aggregate memory | -- | -- | -- |
| Data read and written per run | 1 GB | 10X | -- |
| Maximum I/O bandwidth needed | -- | -- | -- |
| Percent of runtime for I/O | <5% | | |
| Scratch file system space needed | 1 TB | 10X | 10X |
| Permanent online data storage | 10 TB | 10X | 10X |
| Archival data storage needed | -- | -- | -- |

*Please use “core hours” for “conventional” processors (i.e., node-hours * cores_per_node). Intel “Ivy Bridge” is an example conventional processor.

**Please use “node hours” for homogenous many-core architectures. A self-hosted Intel Xeon Phi “Knights Landing” is an example.

***Please use “node hours” for “GPU or accelerator” usage.

****E.g., 32× column 1.

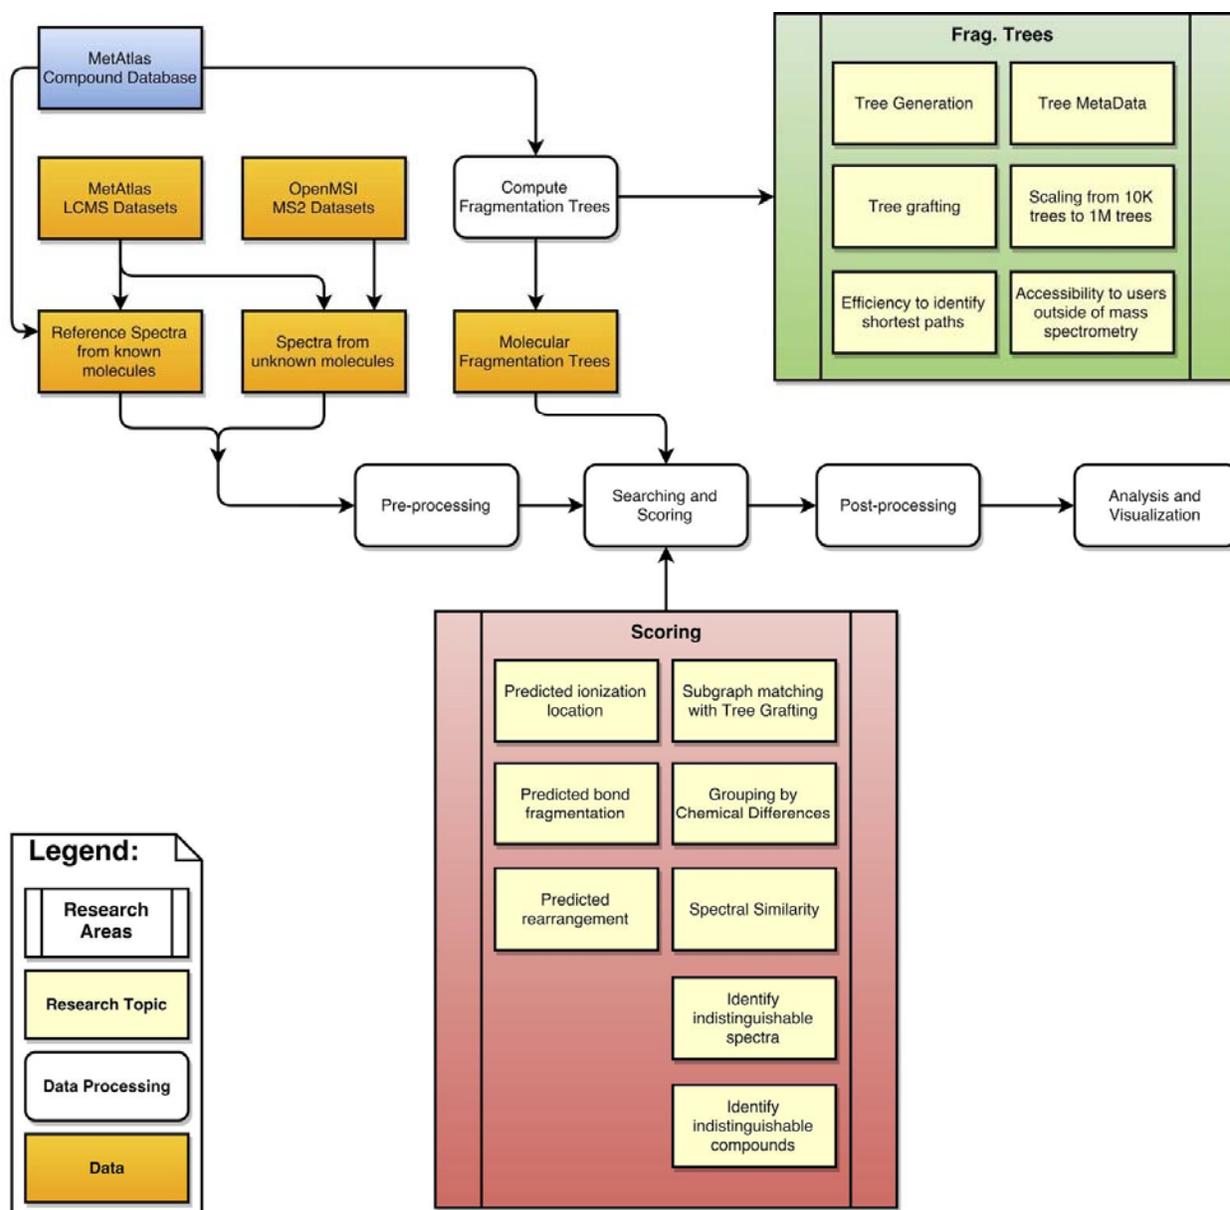


Figure 1. Roadmap for network-based chemical identification. Structures are obtained from a compound database; spectra are obtained from a data analysis and management system; and trees are obtained from a database of fragmentation patterns. Algorithms search and score each spectrum. Last, these are presented to the user through visualizations.

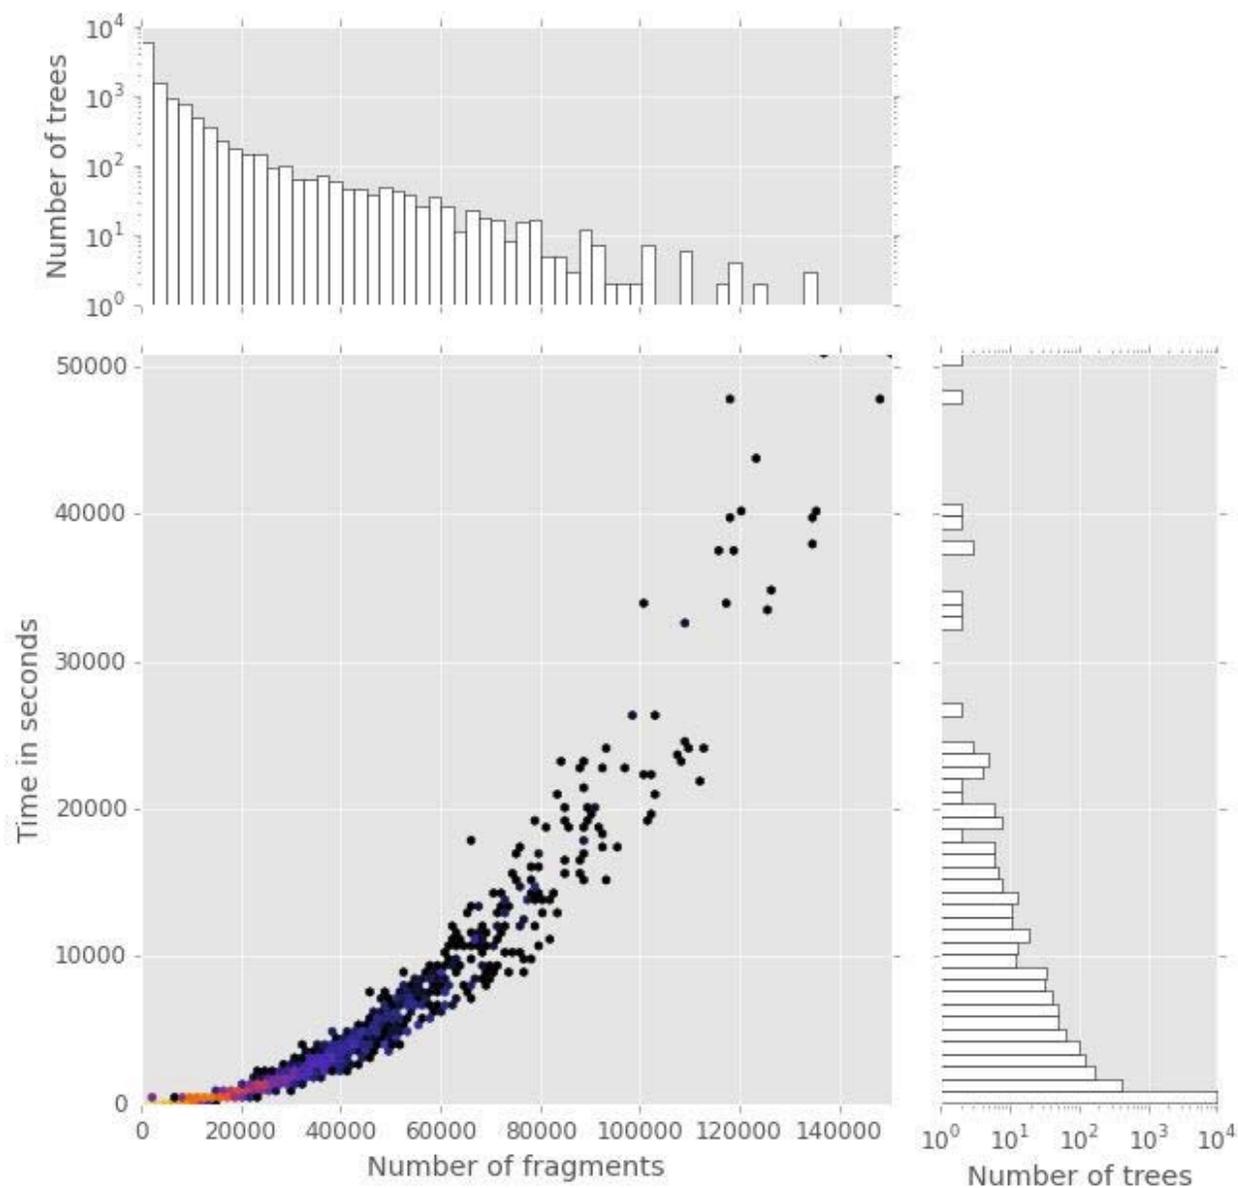


Figure 2. Density plot showing the compute times relative to the number of fragments for generating fragmentation trees for 11,539 molecules consisting of 3 to 302 atoms. For the calculation of complete fragmentation trees, the compute complexity is expected to grow roughly at a rate of n^m , where n is the number of bonds and m is the depth of the tree.

Case Study Title: Scalable algorithm to compute whole genome similarity in metagenomics

Lead Author(s): Chirag Jain and Srinivas Aluru, Georgia Institute of Technology

1. Description of Research

1.1 Overview and Context

Analysis of microbial genomes is essential to study the prokaryotic diversity in an environment. The ANI similarity technique is a well-accepted norm used to compare two or more microbial genomes based on the average nucleotide identity of the conserved genes among them. Previous research has demonstrated the robustness of ANI similarity to measure the genetic and evolutionary distance between genomes (Goris et al. 2007). However, this technique uses blast alignments to compute the distance, thus limiting its scalability to compute the pairwise distances between thousands or tens of thousands of genomes. With growing interest in metagenomics and advancement in the high throughput sequencing technologies, there is now a need to come up with an efficient algorithm and heuristics that improve the scalability of this method without compromising its accuracy. A recent attempt was made to replace this technique with an order-of-magnitude faster alignment-free method based on sketching; however, this only yields accurate results when the two genomes being compared have 98% or more nucleotide identity (Ondov et al. 2015). We have experimentally verified the loss in accuracy for evolutionarily distant genomes. In this project, we are working on bridging this gap—maintaining the accuracy and sensitivity of ANI similarity measurement while being almost as fast as the already proposed alignment-free technique.

1.2 Research Objectives for the Next Decade

The broad goal of this research is to design high-performance combinatorial algorithms that can speed up the analysis of microbial diversity in large-volume sequencing experiments within metagenomics. Going forward, we would like to maintain a database of all prokaryotic genomes along with metadata about their geographical origin and importance. With our fast sequence similarity method, we wish to support a Web service that allows users to upload and search an unknown genome by comparing it against all the prokaryotic genomes like a Web search engine.

2. Computational and Data Strategies

2.1 Approach: Originally, the sketching technique was proposed in the late 1990s for clustering Web documents (Broder 1997). In our method, we treat each gene sequence of a genome as a document with multiple short substrings (k-mers) and follow a similar approach. We compute Jaccard similarity indices for all pairs of genes across the two genomes to identify the set of core-

conserved gene pairs. The Jaccard similarity indices of the potential pairs are then used to estimate the genetic similarity of the two given genomes.

2.2 Codes and Algorithms: We have implemented a prototype version in C++, version 11. To make the implementation process easy and efficient, we make extensive use of C++, version 11, STL data structures and algorithms.

3. Current and Future HPC Needs

3.1 Computational Hours: We are still in the early stages of implementation. Right now, comparison of two genomes takes about a second with our method. However, as we go forward, we aim to improve further and achieve sufficient scalability to use this method to do pairwise comparison of 2,000 genomes. We expect the runtime to be multiple hours on a single core. By 2025, computation hours will increase as more microbial genomes are sequenced around the world. The factor of growth could be as large as 25.

3.2 Parallelism: The aim of this project is to support similarity queries in a fraction of a second. We plan to use both distributed and shared memory parallelism to reduce runtime on large clusters.

3.3 Memory: The proposed algorithm is expected to have low memory requirements relative to input size because it has the advantage of representing the sequences with small sets of sketches.

3.4 Scratch Data and I/O: For our implementation, the data is read at the beginning of the program execution. With the support of good I/O bandwidth (over 1 GB/sec), the overall compute time is usually expected to outweigh the file I/O time. In addition, we have the option to pre-compress all the genomes into compressed formats for faster I/O in future computations.

3.5 Long-term and Shared Online Data: Because the algorithm works on the compressed form of complete genomic sequences, we expect that the long-term online storage will not grow beyond 100 GB for comparing thousands of genomes.

3.6 Archival Data Storage: There are 66,000 prokaryotic genomes hosted on NCBI servers (see <http://www.ncbi.nlm.nih.gov/genome/browse>); they require about a terabyte of memory storage in the uncompressed fastq format. To test and deploy our algorithm on large volumes of data, we will need few terabytes of archival data storage. In the future, we plan to maintain a repository of all the microbial genomes sampled in metagenomics experiments globally to enable a search through a Web interface for biologists.

3.7 Workflows: Our workflow includes the following steps: (1) download the new genomes to off-chip memory; (2) pre-compute the compressed sketches and save them in fast accessible off-

chip memory; (3) compute the pairwise similarity using the sketches; and finally (4) output the similarity results to the console.

In the future, it will be useful to have services that can accept software and offer storage, memory, and disk space along with the Web interfaces so biologists can directly run the software on their data using the underlying hardware. This will be analogous to the NCBI's BLAST initiative, except it will allow the researchers to supply novel software programs associated with peer-reviewed publications. Similar to Blast, MG-RAST portal at Argonne National Laboratory is designed for researchers to upload and analyze the metagenomics raw sequence datasets. As of 2013, MG-RAST had already completed the analysis of over 25 Tera base pairs of DNA data contributed by thousands of researchers worldwide.

4 Requirements Summary Worksheet

| Code: _____ | Column 1: Current Usage | Future Usage: 2020 (As a factor of column 1)**** | Future Usage: 2025 (As a factor of column 1)**** |
|--|--------------------------------|---|---|
| Computational core hours (Conventional)* | 10 | 25X | 50X |
| Computational node hours (Homogeneous many-core)** | 1 | 15X | 20X |
| Computational node hours (w/GPU or accelerator)*** | No accelerator usage supported | | |
| Memory per node | 128 GB | 4X | 8X |
| Aggregate memory | 128 GB | 4X | 8X |
| Data read and written per run | 16 GB | 4X | 8X |
| Maximum I/O bandwidth needed | 1 GB/sec | 4X | 8X |
| Percent of runtime for I/O | 20 | 20 | 20 |
| Scratch file system space needed | 16 GB | 4X | 8X |
| Permanent online data storage | 16 GB | 4X | 8X |
| Archival data storage needed | 1 TB | 25X | 50X |

*Please use "core hours" for "conventional" processors. (i.e., node-hours * cores_per_node). Intel "Ivy Bridge" is an example conventional processor.

**Please use "node hours" for homogenous many-core architectures. A self-hosted Intel Xeon Phi "Knights Landing" is an example.

***Please use "node hours" for "GPU or accelerator" usage.

**** E.g., 32x column 1.

Extreme-Scale Genome and Metagenome Assembly and Analysis

Nikos Kyrpides, Dan Rokhsar, Aydin Buluç (abuluc@lbl.gov), Rob Egan, Evangelos Georganas, Leonid Oliker, Steven Hofmeyr, and Katherine Yelick (kayelick@lbl.gov), LBNL

Relevant offices: BER (Biology and Environment), ASCR, and EERE (Bio Manufacturing)

Overview: Description and Impact

Microbiome research has far-reaching applications in bioenergy, environmental remediation, food production, and nutrition and human health, as highlighted in a recent report from OSTP [1,2]. Microbiome research will also allow scientists to build better models of carbon and nutrient cycling and better predict how ecosystems will respond to stresses like climate change or sea level rise. Widespread screening of microbiome samples is now possible because the cost of sequencing has dropped by over 350,000 times in the past 15 years, with the cost per megabase now under 2 cents. Uncultured microorganisms represent the vast majority of the microbiome's diversity, containing an untapped wealth of information including new life forms with novel functions and enzymatic activities, and biosynthetic clusters that could be used to manufacture novel natural products and chemicals needed for energy and environmental challenges [3,4].

Metagenomics is currently the leading technology in studying the uncultured microbial diversity and delineating the microbiome structure and function. While accessing an unprecedented number of environmental samples that consist of thousands of individual microbial genomes is now possible; however, there is beginning to be a computational bottleneck, because the sequencing cost improvements notoriously exceed that of Moore's Law.

Effective and accurate microbiome analysis depends on the efficiency of the assembled sequences. Metagenomes assembled into long contiguous sequences (contigs) are not only critical for the identification of the usually long biosynthetic clusters but also key for enabling the discovery of new lineages of life and viruses. However, for most of these energy-related environmental samples, there is no existing reference genome, so a first step in analysis is de novo assembly, a challenging computational problem due to the high error rates or short read lengths of sequencers. Metagenome assembly is further complicated by sequences that are repeated across genomes, polymorphisms within a species, and variable frequency of the genomes within the sample.

De novo assembly is one of the most demanding computational bioinformatics challenges, and correctness challenges due to errors in the data may require parameter sweeps to filter errors and maximize quality. The computation requires large amounts randomly accessible memory and a pipeline of algorithms that often runs for days on large-memory SMP machines; some problems are too large for current systems. Recent work by the LBNL team has shown that with distributed memory algorithms and data structures, as well as a PGAS programming model on low-latency networks, an individual problem scales well to tens of thousands of cores. However, the computational power of existing systems is not adequate for the assembly of metagenomics datasets available in public databases, such as JGI's Integrated Microbial Genomes (IMG) [5] and NCBI's Sequence Read Archive (SRA). Once metagenomes are assembled, downstream analysis includes all-versus-all comparison of novel genes (on the order of 30 billion in IMG), which will require exascale computing resources.

System Requirements

I. Code and Tools: We will use HipMer [6], a high-performance re-implementation of Meraculous for distributed-memory supercomputers. Meraculous is a hybrid k-mer/read-based whole-genome assembler that avoids explicit error correction steps. It has most recently been used to construct the de novo assembly of the highly repetitive plant genome of bread wheat (*Triticum aestivum*) [7]. Although Meraculous is one of the best assemblers available in terms of accuracy, it is known to be too slow for large genomes. HipMer reduces the assembly time for large genomes such as human and wheat from days to minutes. Programming models, data structures, and runtimes play key roles in the success of HipMer to exascale. HipMer is implemented primarily in UPC, with some MPI-based components. One-sided atomic operations are crucial for the scaling of its key components, including large distributed-memory hash tables. Similarly, software caching, dynamic message aggregation, and efficient overlapping of computation with I/O—as well with communication—eliminate significant bottlenecks.

II. Models and Algorithms: HipMer has a modular design and contains the building blocks common to many graph-based assemblers. It has been used on metagenome data and is currently being extended with the algorithms necessary for state-of-the-art metagenome assembly, which includes multiple passes to extract and assemble genomes that occur in different frequencies in the sample. Changes in sequencing technology, such as the longer read and higher error rates with PacBio technologies may also require adaptations.

III. End-to-End Requirements: Modules of HipMer stress different supercomputer components. The initial modules require high I/O and network bandwidths. We will explore exascale system features such as NVRAM storage, scalable object stores, and general support for complex workflows. The alignment parts of HipMer are computationally intensive and can be offloaded to an accelerator (with existing GPU and KNL implementations of some kernels), while the graph traversal pieces are limited by the network latency, injection bandwidth, and remote atomic speed. This workload will stress the exascale systems in ways that are distinct from most modeling and simulation problems but reflective of data analytics.

Related Research

Computational genomics benefits from advances in many fields. Advanced distributed data structures (e.g., hash tables and Bloom filters) and distributed streaming algorithms are crucial for processing terabytes of data. Genomics researchers also use machine learning algorithms (e.g., principal component analysis) to separate clean and contaminated contigs in metagenomes. These data analytics algorithms open entirely new classes of exascale analytics applications.

10-Year Problem Target

Although microbiome research has been identified as one of the priorities for the economy for the next 10 years, we have barely scratched the surface of the complexity and diversity of the microbial world. The 10-year grand challenge is to perform assembly and analysis of large complex metagenomes, and use this to screen massive databases of metagenomes to identify novel biological processes for synthesis of chemicals and materials. There are currently over 30,000 raw (i.e., unassembled) metagenome datasets from around the globe available at NCBI. A

HipMer analysis (estimated from runs on a 2.7-TB Twitchell Wetland data) would require ~100M core hours to assemble 1000 such large metagenomes. The data sets will grow by ~16 times by 2025, with the compute requirements growing by 256 times, resulting in 25.6B equivalent core hours. In addition to biomanufacturing applications, this type of analysis will also aid in the discovery of new forms of life [8], viruses, and alternative genetic codes [9]. In addition, post-assembly analysis includes clustering the assembled genes that do not match known isolate genomes; there are on the order of 20 billion of these. This is a problem with quadratic complexity that requires exascale resources even without assuming growth in datasets.

References:

- [1] Handelsman, J., and E.R. Stulberg, 2015, “What’s Next for the Microbiome?,” White House, May 21. Available at <https://www.whitehouse.gov/blog/2015/05/21/whats-next-microbiome>.
- [2] Life Sciences Subcommittee of the National Science and Technology Council, 2015, *Report of the Fast-Track Action Committee on Mapping the Microbiome*, November. Available at https://www.whitehouse.gov/sites/default/files/microsites/ostp/NSTC/ftac-mm_report_final_112015_0.pdf.
- [3] Eloë-Fadros, E.A., et al., 2016, “Global Metagenomic Survey Reveals a New Bacterial Candidate Phylum in Geothermal Springs,” *Nature Communications* 7:10476.
- [4] Hadjithomas, M., et al., 2015, “IMG-ABC: A Knowledge Base to Fuel Discovery of Biosynthetic Gene Clusters and Novel Secondary Metabolites,” *MBio*. 6(4):e00932.
- [5] Markowitz, V., et al., 2014, “IMG/M 4 Version of the Integrated Metagenome Comparative Analysis System,” *Nucleic Acids Res.* 42:D568-73.
- [6] Georganas, et al., “HipMer: an Extreme-scale De Novo Genome Assembler.” *SC’15*.
- [7] Chapman, et al., 2015, “A Whole-genome Shotgun Approach for Assembling and Anchoring the Hexaploid Bread Wheat Genome,” *Genome Biology* 16(26).
- [8] Woyke, and Rubin, 2014, “Searching for New Branches on the Tree of Life,” *Science* 346(6210).
- [9] Ivanova, N., et al., 2014, “Stop Codon Reassignments in the Wild,” *Science* 344(6186):909–913.

Appendix: Requirements Summary Worksheet

Current Usage Scenario: Wetlands metagenome assembly (soil samples across several physical sites from the Twitchell Wetland in the San Francisco Bay-Delta). There are 2.7 TB of data, but this has to be read several times during assembly. K-mer analysis and contig generation take 25% of the overall runtime for wheat (single) genome assembly via HipMer at the highest scale. For metagenomics, the pipeline needs to be iterated many times. **We assume 10 iterations.** Numbers in the table assume a **10K core run** on Edison, which is extrapolated to take **an hour to completion for a single iteration** (a single iteration of K-mer analysis and contig generation takes 15 minutes).

Future Usage 2020: Looking at past data (2010–2015), total metagenome base pairs as well as the size of the largest metagenomics dataset increases by 4 times in 5 years. The majority of the computation is linear in the size of the dataset, and few pieces have quadratic complexity. In the short term, we expect a $N * \log(N) * \log(P)$ scaling **for node/core hours** (divide by P for the time it takes) with data size of N and node count of P. An amount of data that is 4 times larger would require 4 times more cores (assuming memory per node staying constant). Both $\log(N)$

and $\log(P)$ will each contribute a factor of 2, increasing the core hours by a factor of $4 \times 2 \times 2 = 16$. Aggregate memory requirements scale similarly.

Future Usage 2025: Similar extrapolation.

| Code: HipMer, C/C++/UPC | Column 1: Current Usage | Future Usage: 2020 (As a factor of column 1)**** | Future Usage: 2025 (As a factor of column 1)**** |
|---|-------------------------------|---|---|
| Computational core hours (Conventional)* | 100,000 | 16 | 256 |
| Computational node hours (Homogeneous many-core)** | 100,000/cores_ per_node | 16 | 256 |
| Computational node hours (w/GPU or accelerator)*** | | | |
| Memory per node | 64 GB | 1 | 1 |
| Aggregate memory | 13.8 TB | 16 | 256 |
| Data read and written per run | 100 TB | 4 | 16 |
| Maximum I/O bandwidth needed | 100 GB/sec | 1 | 1 |
| Percent of runtime for I/O | 30% | 1 | 1 |
| Scratch file system space needed | 10 TB | 4 | 16 |
| Permanent online data storage | 3 TB | 4 | 16 |
| Archival data storage needed | 3 TB | 4 | 16 |

* Please use “core hours” for “conventional” processors (i.e., node-hours * cores_per_node).
Intel “Ivy Bridge” is an example conventional processor.

**Please use “node hours” for homogenous many-core architectures. A self-hosted Intel Xeon
Phi “Knights Landing” is an example.

***Please use “node hours” for “GPU or accelerator” usage.

****E.g., $32 \times$ column 1.

Case Study Title: Computational Challenges with Multi-omic Association Analysis**Lead Authors: Bernard Ng and Sara Mostafavi****Center for Molecular Medicine and Therapeutics****1. Description of Research****1.1 Overview and Context**

Our research focuses on identifying the interactions between different genomic processes in relation to complex traits and diseases. The raw and processed data together are in the terabyte (TB) range. To analyze the different data types jointly often requires a few tens, sometimes up to hundreds, of gigabytes (GB) of random access memory (RAM), and processing the data requires parallel computing to keep computational time practical. High-speed networking is needed to transfer and retrieve raw and processed data between computational and storage servers. We also use patient genotype data, most of which has to be stored on local servers and cannot be stored/handled on cloud-based servers.

1.2 Research Objectives for the Next Decade

The scientific goals are to identify disease biomarkers and use them to aid diagnosis and design new treatments. To keep computation tractable, many analyses involve a certain degree of approximation. If higher computational power can be obtained for a low cost, permutation and resampling strategies can be employed to generate more reliable results.

2. Computational and Data Strategies**2.1 Approach**

Our research projects typically encompass two serial analysis components: (1) visualization, exploratory data analysis, and prototyping; (2) RAM-intensive computation on large data matrices. Our current strategy is to perform the first step on local servers (e.g., powerful desktops), and once prototyping is done, we perform step 2 on our remote computer cluster. For step 2, there are two bottlenecks in our computational pipelines: RAM and I/O speed. Current datasets that we work with are large (potentially >10 GB) and very slow (if possible) to load into memory at once. Once loaded, RAM is the main bottleneck in performing our analysis. Our current strategy is to split the data and the analysis into parts, perform each part with a separate server, and combine the results. All additional steps might increase the chance for errors. Further, certain parallel processing tools require creating copies of the data for each worker, which limits the number of workers that can be simultaneously employed on a single server. We foresee that with more RAM on servers and graphic processing units (GPUs) in the future, as well as more flexible and user-friendly languages for GPU programming, most of the aforementioned problems would be alleviated.

2.2 Codes and Algorithms

A major component of our work involves data visualization and exploratory analysis. After preprocessing the raw data with standard packages, we run variants of regression models with MATLAB and R. These variants involve solving moderate-sized optimization problems.

3. Current and Future HPC Needs

| Code: Assuming a single run of our typical set of regression analyses. | Column 1: Current Usage | Future Usage: 2020 (As a factor of column 1)**** | Future Usage: 2025 2 (As a factor of column 1)**** |
|--|---------------------------------------|--|--|
| Computational core hours (Conventional)* | 8064 | 10x | 11x |
| Computational node hours (Homogeneous many-core)** | 336 | 10x | 11x |
| Computational node hours (w/GPU or accelerator)*** | 48 (for a small part of the analyses) | | |
| Memory per node | 32 GB | 512 GB | 512 GB |
| Aggregate memory | 5 TB | 10 TB | 10 TB |
| Data read and written per run | TB | TB | TB |
| Maximum I/O bandwidth needed | GB/sec | GB/sec | GB/sec |
| Percent of runtime for I/O | 5% | | |
| Scratch file system space needed | 500 GB | 1 TB | 1 TB |
| Permanent online data storage | 5 TB | 50 TB | 50 TB |
| Archival data storage needed | TB | TB | TB |

*Please use “core hours” for “conventional” processors. (i.e., node-hours * cores_per_node). Intel “Ivy Bridge” is an example conventional processor.

**Please use “node hours” for homogenous many-core architectures. A self-hosted Intel Xeon Phi “Knights Landing” is an example.

***Please use “node hours” for “GPU or accelerator” usage.

****E.g., 32x column 1.

Comprehensive Characterization of Microbial Organisms in Soil Communities Using Big Omics Data

Chongle Pan, Oak Ridge National Laboratory

1. Description of Research

1.1 Overview and Context

Metabolic activities of microbial communities can be inferred from metagenomics and metaproteomics (proteogenomics) analyses. The nutrient flows in a community can be traced using proteomic stable isotope probing analyses. We employ high-performance computing (HPC) to process the large amounts of sequencing data and mass spectral data generated in these omics analyses. Results from these omics analyses are used to study phosphorus cycling in tropical rainforest surface soils, carbon cycling in grassland subsurface soils, microbial degradation of algae blooms in coastal waters, and symbiotic association between *Populus* microbiomes and their hosts.

1.2 Research Objectives for the Next Decade

Soil communities are the least understood type of microbial communities in natural environments, because most microorganisms in soil cannot be cultured in the laboratory, and the extremely high complexity of soil communities defies our current technologies for comprehensive proteogenomics analyses. In the next decade, obtaining a better understanding of soil communities by proteogenomics is an important scientific goal in terrestrial ecology.

Because a typical soil community is composed of thousands of species with relative abundances of 0.01% to 1%, adequate proteogenomics analysis of a single soil sample may require a trillion base pairs (Tbp) of metagenomics sequencing and a million mass spectra of metaproteomics measurement. One Tbp of metagenomics sequencing would provide ~20-fold coverage of the large number of microbial species at 0.01% relative abundance. Because of the high spatial and temporal variability, a terrestrial ecosystem requires such deep proteogenomics analyses across a large number of soil samples collected across space and time. The data processing goal is to scale up the capability of computing for the increased size of individual proteogenomics data sets and scale up the capacity of computing for the increased number of proteogenomics data sets.

2. Computational and Data Strategies

2.1 Approach

Deep metagenomic sequencing requires HPC for assembly and functional annotation of a large amount of genomic sequences of constituent organisms in a community. Processing mass spectral data with HPC enables identification of not only the sequences of proteins but also their post-translational modifications and isotopic enrichment levels from metaproteomics measurements. We currently use MPI to distribute computation across nodes and OpenMP or OpenACC to parallelize computation within nodes. In next decade, the increased size of proteogenomics data sets will require better management of data movement and residence in the memory hierarchy and network storage.

2.2 Codes and Algorithms

Our current proteogenomics work flow includes two MPI/OpenMI algorithms. The Omega algorithm is an overlap graph assembler for metagenomics. Omega first constructs an overlap graph by aligning all reads against each other and then simplifies the overlap graph to extract genomic sequences. Omega uses MPI-3 remote memory access (RMA) to store all reads in a distributed array, and the read alignment computation is parallelized by MPI/OpenMP. The overlap graph obtained was partitioned to contiguous subgraphs, stored in a distributed adjacency list based on MPI-3 RMA, and simplified by parallel MPI/OpenMP threads. The Sipros algorithm is a database-searching algorithm for metaproteomics. Sipros identifies peptides by matching their predicted fragments with observed mass spectra. Mass spectra are distributed across compute nodes and compared to indexed candidate peptides by MPI/OpenMP threads. Load balancing was performed dynamically at both the OpenMP level and the MPI level.

We also use two data-parallel frameworks that distribute computation by single-node algorithms across many nodes. The Sigma framework parallelizes the alignment of reads onto reference genomes by dividing the read data set across compute nodes. Sigma uses CUSHAW2-GPU for read alignment on compute nodes with GPU and uses Bowtie2 on compute nodes with only CPU. The UniFam framework parallelizes the searches of millions of protein sequences against hidden Markov models (HMM) of protein families using HMMER 3. Both frameworks provide parallel post-processing to filter the large amounts of output and aggregate to useful results.

3. Current and Future HPC Needs

3.1 Computational Hours

See the table.

3.2 Parallelism

Our codes can scale up to 64 CPU cores per node and up to 1,000 nodes. We will use more remote memory access and increase data redundancy on local memory to increase the coarse-grained parallelism .

4. Requirements Summary Worksheet

| | Column 1: Current Usage | Future Usage: 2020 (As a factor of column 1) | Future Usage: 2025 (As a factor of column 1) |
|---|-------------------------------|---|---|
| Computational core hours (Conventional) | 15 million | 10X | 40X |
| Computational node hours (Homogeneous many-core) | | | |
| Computational node hours (w/GPU or accelerator) | 1 million | 10X | 40X |
| Memory per node | 32 GB | 500 GB | 2,000 GB |
| Aggregate memory | 32 TB | 100 TB | 400 TB |
| Data read and written per run | 20 TB | 200 TB | 800 TB |
| Maximum I/O bandwidth needed | 500 GB/sec | 1,000 GB/sec | 4,000 GB/sec |
| Percent of runtime for I/O | 10% | 10% | 10% |
| Scratch file system space needed | 20 TB | 100 TB | 400 TB |
| Permanent online data storage | 50 TB | 500 TB | 2,000 TB |
| Archival data storage needed | 100 TB | 1,000 TB | 10,000 TB |

*--Please use “core hours” for “conventional” processors. (i.e., node--□hours * cores_per_node). Intel “Ivy Bridge” is an example conventional processor.

**--Please use “node hours” for homogenous many-core architectures. A self-hosted Intel Xeon Phi “Knights Landing” is an example.

***--Please use “node hours” for “GPU or accelerator” usage.

Case Study Title: Data Fusion

Lead Author(s): H. Steven Wiley and Ronald Taylor (PNNL)

1.1 Overview and Context: Data fusion is the process whereby multiple types of high-throughput molecular data (e.g., genomic, transcriptomics, proteomics) are combined to provide information on the processes that link them. Most previous work on data fusion has used statistical approaches to link changes in the level or amounts of different molecular species with changes in system properties. As an alternate approach, we have designed a way to link multiple data types by mapping all data onto a genomics-based framework, specifying their relationships in terms of mechanistic models. This approach is compatible with large-scale NoSQL data frameworks as well as most high-throughput, genomics-based analytical data. Because of the continuing exponential growth of these types of data across a broad range of biological studies, this data fusion framework provides an attractive platform for extracting new biological insights from existing data.

As a test case, we explored multiple approaches to calculate protein abundance from peptide data, as well as relationships between RNA transcripts and protein levels across multiple samples and conditions. These data were then used to estimate translational efficiency as a function of codon usage; ribosome binding sites; operons by transcript abundance; and normalization of protein expression using transcript abundance.

1.2 Research Objectives for the Next Decade: Future studies will extend this work to available proteomics and transcriptomics data in public repositories, which will require extensive data renormalization prior to data fusion and the creation of appropriate models that specify the relationship between different data types. Thus, we propose to load and reprocess all currently available data using the abundance of highly conserved proteins and transcripts as internal standards. All BAM files from sequencing studies (DNA- and RNA-Seq) could be loaded to look for microdiversity (SNPs) and its impact on gene and protein expression. Linking protein and transcript variance to metabolite expression can build more accurate metabolic models of both individual organisms and communities.

2. Computational and Data Strategies

2.1 Approach: Multiple, linked key values tables in HBase are created at different levels of abstraction, with position in reference genomes being the lowest. Previous work used Hadoop/MapReduce, but we are transitioning to Spark, planning for 200 GB of RAM per node for the Spark component of each node, and 56 GB for the OS and Hadoop using 32 nodes, which should be enough for 1 TB of data. Future work should require perhaps a hundredfold greater capacity.

2.2 Codes and Algorithms: Transcriptomics and proteomics files from both microbial and eukaryotic studies were parsed and uploaded into tables in an HBase data warehouse on a PNNL Linux cluster. Datasets were placed in a small number of flexible HBase tables using locus tag, peptide ID, parent protein, and counts from RNA-Seq corresponding to each gene (rnaSeqCount), taxonomy, and genome tables containing different fields in their records. The protein ID was set to match the locus tag that was the primary key for the records in the rnaSeqCount table that stored the RNA-Seq count data on a per-gene basis. All calculations on protein rollups, peptide predictions, and translational efficiency were performed within the data warehouse system and stored in linked tables.

3. Current and Future HPC Needs

Quantitative information requested in Section 3 must be put in a table following the template below.

HPC needs for “Big Data” searching, processing, and integration in biology are trivial. The problem is the complexity of the data. All the genomics-based biological data in the world generated by DOE will fit into a relatively small commodity cluster. In 10 years, it will fit into a large cluster. The problem is development of the codes, data frameworks, and algorithms needed to process and analyze the data. Gene homology mapping and genomics assembly might require more horsepower, but this is a different type of problem than what I am describing here.

3.1 Computational Hours: How many hours on conventional¹ cores (no accelerators) do the codes involved in the study use now? If you use GPUs or other accelerators, use the next row for that answer. How do you expect your computational requirements to increase through 2020 and 2025? Include all hours you will need to reach the scientific goals you listed in 1.2 above.

3.2 Parallelism: Please describe the scale at which your codes use coarse-grained (multi node) and fine-grained (on-node or accelerators) parallelism today. Please describe current plans to increase either level of parallelism.

3.3 Memory: Describe your current and future memory requirements in terms of the minimum shared memory pool (node) and aggregate memory required for you to run. Note that future systems may have much less memory as a function of peak performance than systems have today. Note also that the memory hierarchy will be potentially complex (on-chip fast memory and significantly slower to off-chip memory).

3.4 Scratch Data and I/O: How much online scratch storage space do you need for your runs (current/future), including checkpoint/restart data? Please estimate your I/O bandwidth requirement (bandwidth = data read or written / time to read or write). What percentage of your total runtime are you willing to devote to I/O?

3.5 Long-term and Shared Online Data: How much active, online long-term storage do you need today and in 2020 and 2025? Please describe any requirements for sharing or accessing the data.

¹ Conventional = current multicore CPUs like Intel “Ivy Bridge.”

3.6 Archival Data Storage: Archival data is accessible online, but may involve a delay in accessing it (e.g. data stored on HPSS tapes). How much data do you have stored in a data archive currently? How much will you need in 2020 and 2025?

3.7 Workflows: Please briefly describe your current workflows and requirements for 2020 and 2025.

3.8 Many---Core and/or GPU Readiness: Future systems will contain “lightweight” cores and/or hardware accelerators (e.g., GPUs) with deepening memory hierarchies. Are your codes ready for this? If yes, please explain your strategy for exploiting these technologies. If not, what are your plans for dealing with such systems and what do you need to help you successfully transition to them?

Answers to the following questions are not required as these issues will be dealt with more globally in the white papers. However, if you have a unique situation that you feel will need to be addressed, please do so. Keep it short!

3.9 Software Applications, Libraries, and Tools: If you think your needs for HPC software (applications / libraries / tools / compilers / languages / etc) will change between now and 2020, 2025, please describe them here? Be sure to consider workflows, analytics, and I/O software.

3.10 HPC Services: If you anticipate needing additional HPC services not provided today in 2020---2025, what are they? Possibilities might include advanced training, data analytics and visualization assistance, support servers, collaboration tools, web interfaces, federated authentication services, gateways, etc.

3.11 Additional Needs: If there is anything else important to the success of your project that you have not mentioned, please do so here.

4. Requirements Summary Worksheet

Please fill out the following table to the best of your ability prior to the meeting. If you are not able to make an estimate, leave the entry blank and we will discuss it at the meeting.

| Code: _____ | Column 1: Current Usage | Future Usage: 2020 (As a factor of column 1)**** | Future Usage: 2025 (As a factor of column 1)**** |
|---|----------------------------|--|--|
| Computational core hours (Conventional)* | | | |
| Computational node hours (Homogeneous many---core)** | | | |
| Computational node hours (w/GPU or accelerator)*** | | | |
| Memory per node | GB | GB | GB |
| Aggregate memory | TB | TB | TB |
| Data read and written per run | TB | TB | TB |
| Maximum I/O bandwidth needed | GB/sec | GB/sec | GB/sec |
| Percent of runtime for I/O | | | |
| Scratch file system space needed | TB | TB | TB |
| Permanent online data storage | TB | TB | TB |
| Archival data storage needed | TB | TB | TB |

*---Please use “core hours” for “conventional” processors. (i.e., node---hours * cores_per_node). Intel “Ivy Bridge” is an example conventional processor.

**---Please use “node hours” for homogenous many---core architectures. A self---hosted Intel Xeon Phi “Knights Landing” is an example.

***---Please use “node hours” for “GPU or accelerator” usage.

****--- E.g., 32× column 1.

D.2 Case Studies Addressing Environmental Research

Software Engineering for Portability (corresponds to Section 3.2.7.3)

Exascale Computing for Subsurface Science: A Case Study Based on the PFLOTRAN Simulator

Glenn Hammond (gehammo@sandia.gov), Jeff Johnson (jnjohnson@lbl.gov)

Science Challenges

Modeling at high-fidelity earth-system processes:

- Fundamental subsurface biogeochemical processes
- Impact of climate change on terrestrial ecosystem processes
- Fate and transport of contaminants at DOE facilities, legacy waste sites, and nuclear waste repositories
- Terrestrial storage of supercritical CO₂

The following are areas where major advances are expected in the next 5–10 years:

- Full 3D representation of geologic formations at 10 to 100 times the level of detail
- Tighter coupling of multiscale and multiphysics processes
- Maturation of data assimilation methodologies (that leverage subsurface simulators as engines to run ensemble calculations for coupled process models)

Current and Future Computational and Data Strategies

The current approach to modeling subsurface processes within subsurface simulator is to couple these processes in either a fully coupled or sequentially coupled approach, or a hybrid of the two. For instance, the code PFLOTRAN (Hammond et al., 2014) solves for multiphase flow and reactive multicomponent biogeochemical transport as separate nonlinear systems of partial differential equations (PDEs) that are coupled sequentially. However, multiphase flow fully couples the PDEs for two fluid phases and energy, solving the resulting nonlinear system of equations using a Newton-Krylov solver. Reactive transport also solves for its PDEs simultaneously. It is not anticipated that this implicit solution approach will change significantly in the future as it is required to stably accommodate the disparate timescales involved.

Aspects of Computing that Accelerate or Impede Progress

Accelerators

1. *Scalable Newton-Krylov solvers.* Iterative solvers are the only option for solving large systems of linear equations, and are an elementary building block for any simulation approach that integrates disparate time scales. These solvers are currently our largest bottleneck for scalability because they require global communication (Hammond et al., 2012). “Multi-level” methods such as multigrid can significantly improve scalability, but they still need some form of global communication, and their application to system PDEs is an open area of research. Further, these solvers need preconditioners, and nonlinear preconditioning is currently very problem specific (“dark art/sorcery”) and needs either more systematic analysis or to be separated from the critical path.
2. *Workforce development.* Teams comprising researchers and software developers that trust each other and work well together would greatly accelerate the development of next-generation models. The cohesion of the project/team would require long-term funding for stability, which can be counterbalanced by accountability for tangible results. Without cohesive teams of interested and appropriately selected personnel with time for focused work, enthusiasm wanes and progress is slow. Without long-term funding, staff members become overcommitted and multi-institutional efforts encourage the pursuit of pet projects. Without accountability, mediocrity prevails.

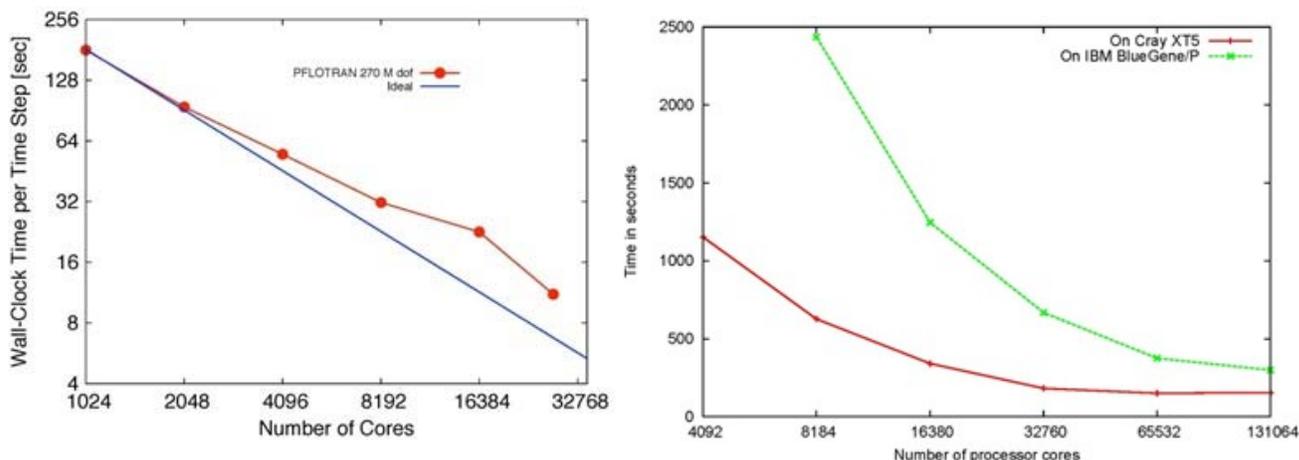
3. *Improved solution methodologies.* As posed, many multiphase models have discontinuities which hinder the performance of scalable solvers. Focusing efforts on alternatives with nicer continuity properties would help solvers do a better job.

Impediments

1. *Current trends in hardware and programming model complexity.* Even in current-generation software, the complexity of models makes high-performance code bases difficult to maintain. In an exascale landscape where researchers are forced to consider details such as hierarchical memory, the population of qualified HPC developers is becoming increasingly rarified and specialized. Interdisciplinary projects suffer the most from this issue.
2. *Reliance upon global communication.* At larger process counts, the cost of global communication becomes the dominant factor in performance. Communication minimizing algorithms must be developed and deployed in applications.
3. *Lack of parallel input (the “I” in I/O).* As models become larger and more complex with large datasets being assigned to the model domain, the reading of these datasets (of increasing size) will become prohibitive.

Software Applications, Libraries, and Tools

Although PFLOTRAN makes use of PETSc’s scalable solvers, the HDF5 parallel I/O library, and MPI for message passing, the largest barriers to exascale subsurface simulators are still the lack of scalable solver algorithms and I/O. This was demonstrated through a SciDAC II groundwater project where PFLOTRAN was run on the Jaguar XT5 supercomputer at ORNL utilizing up to 262,144 cores. The performance of the code was assessed on Jaguar XT4 on up to 27,580 cores and Jaguar XT5/ANL’s IBM BGP on up to 131,064 cores (see figures below). The major bottleneck in the latter was attributed to the global reduction within the iterative Krylov solver based on profiling results (Hammond et al., 2012).



It is likely that OpenMP can be used to accelerate certain portions of the code in an expedient fashion. Through ORNL’s Center for Accelerated Application Readiness in 2011, the 0D chemical reactions within PFLOTRAN were refactored by Cray and nVidia developers to use accelerators. The results of the one-year effort can be summarized as follows: (1) for an exaggerated biogeochemical system (e.g., composed of ~30 chemical species), the maximum speedup experienced was ~14x. For a more realistic biogeochemical system (e.g., 10–15 species), the maximum speedup was < 4.

Summary

Based on our experience applying HPC to BER research problems, the major barriers to exascale performance are the same algorithmic barriers that curtail petascale performance, with the additional constraints introduced by accelerators and hierarchical memory. To realize scalable performance at the exascale, we believe that greater emphasis should be placed on software development, both in the study

and development of scalable algorithms, and in the practices of staffing, training, and motivating teams of engineers and scientists.

References

- Hammond, G.E., P.C. Lichtner, C. Lu and R.T. Mills (2012) Chapter: PFLOTRAN: Reactive flow and transport code for use on laptops to leadership-class supercomputers, *Ebook: Groundwater Reactive Transport Models*, Editors: Zhang F., G.T. Yeh, and J.C. Parker, Bentham Science Publishers, p141-159, eISBN: 978-1-60805-306-3, doi:10.2174/978160805306311201010141.
- Hammond, G.E., P.C. Lichtner and R.T. Mills (2014) Evaluating the Performance of Parallel Subsurface Simulators: An Illustrative Example with PFLOTRAN, *Water Resources Research*, 50, doi:10.1002/2012WR013483.

This page is intentionally left blank.

