The NCAR-Wyoming Supercomputing Center Science Justification

Proposed to

The National Science Foundation

by

The National Center for Atmospheric Research

and

The University Corporation for Atmospheric Research

in partnership with

The University and State of Wyoming

4 September 2009

NCAR
NATIONAL CENTER FOR ATMOSPHERIC RESEARCH

UNIVERSITY OF WYOMING

The National Center for Atmospheric Research is sponsored by the National Science Foundation. Any opinions, findings and conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Table of Contents

Project Summary

The need for dedicated petascale[1] resources for the geosciences was argued persuasively in a 2005 NSF-sponsored report (Bryan et al., 2005). More recently, due in part to the alarming long-term and global threats to society revealed in the recent IPCC Fourth Assessment Report on climate change (Parry et al., 2007), calls for sharply expanding the resources devoted to the Earth System sciences[2] have become more frequent and insistent (Stern, 2006, Lyons et al., 2008, Shukla et al., 2008, Shukla et al., 2009, Ramanathan et al., 2009). The goals of such an initiative are summarized in the NRC's 2009 report, *Restructuring Federal Climate Research to Meet the Challenge of Climate Change* (Ramanathan, 2009):

> *To plan for the effects of climate change, the next generation of global climate models will have to provide numerical simulations on a spatial scale of a few kilometers, with enhanced vertical resolution and better representation of the upper atmosphere.*

The computing limitations are clearly delineated there as well:

> *...most if not all the priorities listed above are dependent on, and have to date been limited by, inadequate computational resources....*

It is clear that an effective response to these calls will require dedicated petascale and eventually exascale[3] high performance cyberinfrastructure[4] (HPCI).

To address these limitations and enable the next series of advancements in the Earth System sciences, NCAR/UCAR and partners the State and University of Wyoming propose to build and operate the NCAR-Wyoming Supercomputer Center (NWSC) in Cheyenne, Wyoming – less than 100 miles from the NCAR/UCAR campus in Boulder, Colorado and 35 miles from the UW campus in Laramie, Wyoming. The proposed NWSC will have the power, space, and cooling to support a 4 MW supercomputer, sufficient to provision a 1.0 to 1.5 petaflops peak system with an anticipated data production rate of 23 to 35 petabytes/year in 2012. In alignment with the core values of our community, the design of the NWSC will set a new industry standard in data center energy efficiency. The specific facility construction plan we propose has been developed in close consultation with the community over a number of years and is described in detail in the companion document, the *NCAR-Wyoming Supercomputing Center Project Execution Plan* (PEP).

Upon completion of the NWSC, the partnership will deploy a petascale system in it that will provide a 15 to 20-fold increase over resources currently available at NCAR, and approaching the aggregate computing resources currently available on the TeraGrid. Further, we propose that the NWSC be well integrated with the CI of the NSF's TeraGrid eXtreme

---

[1] Petascale means computing at a rate of about 10^{15} arithmetic operations per second.

[2] The term "Earth System science" describes a multidisciplinary approach to the study of the planet Earth and its environs that stresses the investigation of the interactions among the Earth's components (including the oceans, atmosphere, cryosphere, and biosphere), solar physics, as well as the relevant mathematical, computational, and social sciences. Earth System science describes the collection of disciplines involved in this multidisciplinary effort,and updates, in a less atmosphere-centric way, the often-used term "atmospheric and related sciences."

[3] Exascale means computing at a rate of about 10^{18} arithmetic operations per second.

[4] Cyberinfrastructure means the overall the technological solution to the problem of efficiently connecting data, computers, and people with the goal of advancing science.

Digital (XD), the Track-1 “Blue Waters” system, and the DataNet program. It should also be well connected to computing facilities operated by other Federal agencies such as the DOE, NOAA, and NASA, as well as high performance systems located at colleges and universities.

Finally, the NCAR-Wyoming Partnership proposes to leverage the world-class capabilities of the NWSC to expand and bolster the impact of the partnership’s education, outreach, and training and diversity programs in the Earth System and computational sciences, including UCAR’s Significant Opportunities in Atmospheric Research and Science (SOARS) program, NCAR’s Advanced Study Program (ASP), and its Summer Internships in Parallel Computational Science (SIParCS). At the University of Wyoming, the NWSC will strengthen its STEM education, workforce development, and outreach activities to underrepresented groups.

Intellectual merit. The vision of seamlessly modeling the Earth System across scales actually represents an emerging research paradigm shift that requires an increasingly interdisciplinary and integrative approach to realize. At present, we can only glimpse the scientific rewards of fully applying this new paradigm with adequate computational support. Quoting from the NRC’s 2007 *Decadal Survey of Earth Science and Applications from Space: National Imperatives for the Next Decade and Beyond* (NRC, 2007):

> *Understanding the complex, changing planet on which we live, how it supports life, and how human activities affect its ability to support life in the future is one of the greatest intellectual challenges facing humanity. It is also one of the most important for society as it seeks to achieve prosperity and sustainability.*

When set in the context of the urgency of the climate change issue, the emerging Earth System science research paradigm requires a bold plan for aggressive deployment of dedicated computing capability on a scale sufficient to dramatically accelerate the rate of scientific advancement. This has caused us to radically rethink the appropriate scale of the HPCI resources required and the facilities that house them.

The massively parallel[5] supercomputers and related CI housed in the NWSC will enable dramatic improvements in resolution, better process representations, longer simulation length, and better statistics for a broad spectrum of important Earth System science applications. As explained in the Community Science Objectives section below, with this system, mesoscale ocean general circulation models will become routine in climate studies, and scientists will be able to use this system to address the greatest uncertainty remaining in the climate system – the representation of clouds. For the first time, regional climate simulations with nested grids at meteorological resolutions will become feasible. This will help scientists investigate the connection between climate and hurricane frequency and strength, study the localized effects of regional climate change on agriculture and water supplies, and numerous other computationally demanding Earth System processes. Solar researchers will be able to use vastly enhanced computer models of the Sun that enable the first-ever simulations of the full life cycles of sunspot assemblages, unlock the mechanisms behind the 11-year solar cycle, model coronal mass ejections, and understand the role of solar variability in climate.

Team qualifications. The NWSC partnership between NCAR and the State and University of Wyoming brings together the financial resources needed to construct the facility and the skill

[5] Petascale systems are expected to have between 50,000 and 100,000 CPUs or “cores.”

to operate dedicated HPCI for Earth System modeling. It also unites NCAR's scientific expertise in the atmospheric and related sciences with Wyoming's research strengths in the areas of subsurface flow, energy, and computational science, thus providing the partnership with a complementary portfolio of energy and climate research capabilities.

NCAR/UCAR. NCAR and UCAR are uniquely qualified to lead this effort by their long history of leadership in and service to the atmospheric sciences community, their experience and familiarity with Grid computing, and their successful conduct of large, complex projects. In particular, NCAR's experience with Grid technology, developed through working on the DOE Earth System Grid (ESG) and membership in the TeraGrid, makes it technically feasible to construct distributed HPCI and extend it into partner institutions. Construction of the NWSC will be overseen by the same NCAR project management personnel responsible for delivering the HIAPER MREFC program, an $81.5 million project that built and outfitted a high-performance mid-size jet for atmospheric research. NCAR's cultivation of proven large-project management practices over the course of the program are the credentials it brings to the NWSC design, construction, and commissioning effort to benefit the Earth System sciences community. NCAR and UCAR's decades of experience working with the 73 member universities in the UCAR consortium qualify them to construct and work within effective governance models. NCAR has worked in partnership with the broader atmospheric and related scientific community to develop world-class community modeling infrastructure for execution on NCAR systems and other major computational resources available in the U.S. The latter include NSF TeraGrid resources and U.S. Department of Energy (DOE) supercomputing systems. Further, NCAR has unmatched experience in providing dedicated, discipline-specific supercomputing derived from decades of providing supercomputing resources and data services to the atmospheric and related sciences (DSS, 2007, Semtner et al., 2009, Leovoy et al., 2002, Dunning et al., 2001, Semtner et al., 2006, Good et al., 2006).

The State of Wyoming. An important enabling partner in the construction and operation of the NWSC, the State of Wyoming has committed substantial financial resources to the project. As an EPSCoR[6] state, Wyoming has a particularly important role to play in facilitating the NCAR-Wyoming partnership's interaction with other EPSCoR jurisdictions.

The University of Wyoming. An academically valuable partner on the team, UW will also help translate the technological and educational impacts of the NWSC into benefits for Wyoming, the region, and especially other EPSCoR states. Specifically, over the past 10 years UW has built a strong group specializing in computational geosciences and particularly in multi-phase fluid flow in porous media – or the movement of water, natural gas, petroleum, and carbon dioxide within sandstone and carbonaceous reservoirs at depth. Now UW has one of the nation's strongest groups in this area. By partnering with the University of Wyoming, NCAR and the NWSC will bring together energy researchers, climate researchers, and experts in modeling carbon sequestration as a climate change mitigation strategy under the same research umbrella.

Supporting broader participation. A particularly important aspect of the proposed NCAR-UW partnership is its interaction with EPSCoR jurisdictions. The partnership will leverage Wyoming's EPSCoR experience to help extend the benefits of the NWSC facility to other EPSCoR

[6] The Experimental Program to Stimulate Competitive Research (EPSCoR) is a federal program to strengthen research and education in science and engineering throughout the United States. The term "EPSCoR states" refers to those states that have been determined to have historically received insufficient research and education funding.

state universities, where such capabilities are likely to be just emerging or perhaps absent. Through EPSCoR, the NWSC partnership model will promote scientific and infrastructural development in these states and will provide a vehicle for EPSCoR and non-EPSCoR universities to work together to advance both computational science and the geosciences.

Broader impacts. There is a rising demand by decision makers at all levels of society for improved understanding and prediction of climate and weather phenomena. The decisions to be based on this knowledge have staggering societal implications, and this provides a compelling justification for the NWSC project. Urgent questions about climate change include the impacts of regional change on water supplies, agriculture, and the spread of disease; shifts in the likelihood of extreme weather events such as heat waves, droughts, and floods; and the advisability of specific adaptation and mitigation strategies – the latter including ideas such as removing anthropogenic carbon dioxide from the atmosphere by sequestering it in deep geologic formations. These new issues add to a long list of short-term threats such as severe storms for which society has long been demanding improved forecasts. The NWSC is a direct response to these needs, enabling the deployment of high performance CI that researchers require to perform high-resolution simulations of weather phenomena, global and regional climate, coastal oceans, sunspots, subsurface flow, and more. It will also increase – by over an order of magnitude – the resources available for experimental weather forecasts and support for meteorological field campaigns, thereby accelerating technology transfer between researchers and operational forecasters.

Earth System research and education will be transformed by the creation of the NWSC. The next generation of Earth science researchers and computational scientists will be attracted by the importance of the problem and the scale of the facilities available to them: their talent will be developed through practical research experience gained on petascale equipment from existing and successful education, outreach, and training programs such as UCAR's Significant Opportunities in Atmospheric Research and Science (SOARS) and NCAR's Advanced Studies Program (ASP) and Summer Internships in Parallel Computational Science program (SIParCS). Further, integration of the NWSC facility with other NSF HPCI will provide important cross linkages between these and similar programs at other resource providers. Integration will also broaden the number and types of researchers across the nation making effective use of the facility – particularly for multidisciplinary research – and thereby support the NSF's vision of a transformative national petascale cyberinfrastructure for science and engineering.

A portion of the NWSC system will be used by University of Wyoming scientists and their collaborators. This computational resource's impact on the University of Wyoming's academic strategic plans, particularly its growing capabilities in computational science, STEM education, and workforce development will be transformational, attracting faculty candidates, collaborators, and talented students to the state. The NWSC will thus accelerate growth of the University of Wyoming's research capabilities in the related disciplines of atmospheric science, climate science, and ecology.

The NWSC will also have a profound long-term impact on the economy of the State of Wyoming and the region. In particular, the NWSC will provide an anchor to attract industries with complementary computing facility needs. As a catalyst for CI development, IT economic sector growth in Wyoming has the potential to translate into economic benefits to the state such as well-paying jobs. In particular, through partnership, these beneficial effects can be extended to other EPSCoR states.

Finally, the overhead from operating U.S. servers requires almost as much energy as the servers use themselves. By deploying an innovative, energy-efficient design, the NWSC represents an engineering advancement that, if replicated throughout the country, could improve data center efficiency and increase the competitiveness of U.S. industry.

Note on the NWSC Proposal Document Set

This *Science Justification* document is the second of three components that together constitute the NCAR-Wyoming Supercomputing Center proposal. The first, called the *Project Summary*, is a four-page preface bundled with this *Science Justification* and designed to give the reviewer to the overview proposed activity, as well as to highlight its intellectual merit and broader impacts. The third and final part of the proposal is a separate document named the *NCAR-Wyoming Supercomputing Center Project Execution Plan*. It provides a detailed discussion of project requirements; infrastructure; organization; description; scope; schedule; budget; configuration management; quality management; communications; risk management; procurement management; project management control system; environment, health and safety planning; facility commissioning, acceptance, and warranty service; technology acquisition and transition to operations; and closing.

This document has three primary purposes: 1) Provide details of the scientific community's research objectives that the proposed NWSC facility is designed to address, 2) Describe the research infrastructure and technical requirements needed to meet the identified objectives, and 3) Discuss the broader impacts of the project. It is organized as follows:

1. **Introduction** provides the background information, history, and context needed to understand the sections that follow.

2. **Community Science Objectives** surveys a representative subset of the important Earth System science objectives driving this proposal to construct the NWSC.

3. **Technical and Research Infrastructure Requirements** translates the community science objectives into technical, operational, and support requirements for the cyberinfrastructure housed by the facility.

4. **Facility Requirements** translates technical and research infrastructure requirements into specific facility requirements.

5. **Broader Impacts** presents the societal impacts of the proposed research, the expected impacts of the NWSC and the infrastructure it houses on research and education, and impacts on the State and University of Wyoming and the region.

6. **Acknowledgements**

Appendix A presents background information about NCAR, including its computational facilities and capabilities.

Appendix B presents background information about the State and University of Wyoming.

Appendix C provides references for the research cited in this *Science Justification*.

1. Introduction to the NWSC Science Justification

In recent years, numerous reports have called for dramatically increased investment in, and better coordination of, high performance computing (HPC) resources supporting the U.S. science and engineering research enterprise (Atkins et al., 2003, Colella et al., 2003, Graham et al., 2004, and Benioff et al., 2005). These calls to action were in part motivated by recognition of the crucial role the computational sciences had assumed in the overall scientific enterprise. Specifically, in 2003, the National Blue-Ribbon Advisory Panel on Cyberinfrastructure (Atkins et al., 2003) stated:

> *The classic two approaches to scientific research, theoretical/analytical and experimental/observational, have been extended to in silico simulation to explore a larger number of possibilities at new levels of temporal and spatial fidelity.* p. 4.

The report then called on the NSF to invest $1 billion per year in cyberinfrastructure for scientific and engineering research.

At the Department of Energy, the case for large-scale computer simulation was also being made in the 2003 *A Science-Based Case for Large-Scale Simulation* (Colella et al., 2003) report, which concluded that many scientific and engineering applications were on the threshold of new knowledge, including those in climate and weather. This "SCaLeS" report recommended that:

> *The Government must provide long-term funding for national high-end computing centers at levels sufficient to ensure the regularly scheduled deployment and operation of the fastest and most capable high-end computing systems that address the most demanding computational problems.*

In 2005, the Presidential Information Technology Advisory Committee (PITAC) issued a report entitled *Computational Science: Ensuring America's Competitiveness* (Benioff et al., 2005). The report reiterated many of the points made by Atkins and SCaLeS, eloquently articulating the role of computers as scientific instruments:

> *Computational science provides a unique window through which researchers can investigate problems that are otherwise impractical or impossible to address….*

Regarding the importance of advanced computing to the nation, the PITAC report went on to emphasize the critical role that computing plays in national security, scientific leadership, and economic competitiveness. Echoing a similar call for cyber-leadership from the NSF in Atkins et al., (2003), the 2005 PITAC report also recommended

> *… that the Federal government provide an infrastructure that includes and interconnects computational science software sustainability centers, data and software repositories, and high-end computing leadership centers with each other and with researchers.*

These calls have led to a variety of responses by the U.S. government, including creation of the Office of Cyberinfrastructure at the NSF and the development of a new CI vision by NSF (Crawford et al., 2007). NSF's strategy has focused on deploying general-purpose HPC CI, serving all of science and engineering, through the TeraGrid program (Catlett et al., 2007). Currently the TeraGrid operates roughly 1.6 petaflops[7] of peak computing power distributed across 14 computing systems at 11 sites.

---

[7] 1 petaflops = 10^{15} floating point (arithmetic) operations per second.

However, along with these general calls for enhanced cyberinfrastructure for science and engineering, persuasive arguments were also advanced for the adequate provisioning of discipline-specific cyberinfrastructure, for example by Atkins et al., (2003), which refers to this issue by stating:

> *These [recommended cyberinfrastructure] programs would also create any discipline-specific cyberinfrastructure required to support these applications, often based on extensions to the more generic cyberinfrastructure. The justification for this is the belief that disciplinary experts, in close partnership with computer scientists, are best able to judge the merits, impact, and importance of applications and specialized cyberinfrastructure focused on their field.* p. 64.

In the intervening years, the need for enhanced discipline-specific HPC CI dedicated to developing a comprehensive understanding and predictive capability for the Earth System, has been articulated in a number of reports (Bryan et al., 2005, Shukla et al., 2008, Ramanathan et al., 2009). In addition, a variety of reports have concluded that U.S. leadership in the computational Earth System sciences is an essential component of a balanced U.S. geosciences research program (Bryan et al., 2005, Bleck et al., 2002, Kinter et al., 2004, Cohen et al., 2005). In addition, there have also been repeated and independent calls outside the U.S. for the creation of greatly augmented discipline-specific HPC infrastructure devoted to tackling problems in the Earth System sciences and managed by that community (Hewitt et al., 2007).

Most recently, the World Climate Research Programme (WCRP) World Modeling Summit for Climate Prediction Workshop report (Shukla et al., 2009) asserted unequivocally that:

> *The weather and climate modeling community does not have sufficient computing power to build and develop the next generation of cloud system resolving models…. It is essential that computing power be increased substantially (by a factor of 1000), and scientific and technical capacity be increased (by at least a factor of 10) to produce weather and climate information of sufficient skill to facilitate regional adaptations to climate variability and change.* p. 7.

Further, the National Research Council's 2009 report, *Restructuring Federal Climate Research to Meet the Challenges of Climate Change* (Ramanathan et al., 2009), stated that

> *... most if not all the priorities listed above are dependent on, and have to date been limited by, inadequate computational resources and the associated technical personnel.* p. 221.

> *New computing configurations will be needed to deal with the computational and data storage demands arising from decadal simulations at high resolution with high output frequency.* p. 6.

It is clear from these findings that the HPCI needs of the Earth System sciences community are not being met. While it is tempting to ascribe this condition to an inadequate government response to the initial calls for enhanced cyberinfrastructure, it is also fair to note that deploying integrated, national-scale infrastructure of the kind envisioned by these reports takes longer and is more difficult to implement and sustain than originally envisioned. It is also the case that the scale of computing needed by the Earth System sciences doesn't fit well with the "niche" model envisioned for discipline-specific cyberinfrastructure that emerged from the Atkins report. When examined closely, what emerges from these reports is that the Earth System sciences really are a multiplicity of closely related complex science problems, with a rather unique set of simulation requirements, needing enormous computational resources, and having a set of societal impacts so immediate and so profound that the answers really cannot wait.

Origins of the NCAR-Wyoming Supercomputing Center project. The National Center for Atmospheric Research (NCAR) and the University Corporation for Atmospheric Research (NCAR) recognized as early as 2004 that 1) the projected computing requirements of climate system models were growing faster than computer performance, and 2) computing technology trends, particularly in the areas of heat density and power consumption, would eventually cause the systems to outstrip the power and cooling capabilities of existing facilities. As the NSF's Federally Funded Research and Development Center (FFRDC) supporting the atmospheric and related sciences, NCAR and UCAR, with NSF guidance, immediately began to work with the research community to better define the science goals, application characteristics, and computing requirements of the Earth System sciences. Between 2005 and 2008, a series of three NSF-sponsored workshops and two NSF-sponsored panels collected input from the scientific community (Loft et al., 2005, Laursen et al., 2006, Kinter and Seidel, 2008, Snavely et al., 2008). At the same time, NCAR and UCAR met with HPCI vendors and data center and other experts to examine various alternatives for effectively responding to the challenge.

Further refinements of the project direction. Two important questions that defined the response were: how should the needed resources be provisioned, and where should they be housed? Kinter and Seidel, 2008 articulated the need for dedicated HPCI operated by a discipline-specific organization. Summarizing their arguments, this model of provisioning resources has the following points in its favor. 1) It has proven its utility over many years for the atmospheric and related science community. In particular, NCAR's scientific computing enterprise possesses a long track record of service to the atmospheric and related sciences. The strongest evidence of this success is the fact that despite the availability of larger computing resources elsewhere, NCAR's systems remain heavily subscribed (see Appendix A for details). 2) The technical requirements of important scientific applications such as climate system models are unique: the discipline-specific approach to supercomputing allows services and system design to be tailored to meet these requirements. 3) It ensures that model development and related research in Earth System processes can occur in a controlled yet responsive environment in which researchers can prepare and test complex models. 4) It provides the organizational focus, capabilities, and skill-sets required to support important field and computational campaigns – including those driven by unfolding natural disasters – with on-demand resources.

Finally, the panel examined the question of obtaining these computing resources from the TeraGrid and concluded that the needs of the research community would be met best if a dedicated facility were created, but the panel also recommended that it be well-integrated with other NSF HPCI such as the resources of the TeraGrid and the Track-1 system.

Regarding location, NCAR and UCAR conducted a careful review of a variety of options including refurbishment or expansion of existing UCAR facilities, purchase or lease of a facility, or collocation of equipment in other facilities operated by the Federal government or private industry. Eventually, NCAR and UCAR concluded the best, most impactful, and most cost-effective response to the challenge was to propose the construction a new supercomputing facility along the Front Range[8]. Locating it near UCAR's Boulder campus allows NCAR to more effectively manage the facility and to bring its computational and domain specific expertise to bear. After a regional competition, UCAR/NCAR chose to partner with the State and University of Wyoming to construct the NWSC.

[8] The Front Range is an area where the eastern edge of the U.S. Rocky Mountains meets the western edge of the Great Plains. The Front Range runs roughly north-south from southern Wyoming to south-central Colorado.

2. Community Science Objectives

We have compiled in this section a representative sampling of the science objectives of Earth System science modeling community, focusing on the ways these objectives will benefit from access to dedicated petascale CI. The examples were chosen to represent both the breadth of sub-disciplines involved and the multiscale nature of the problems being studied. Objectives were obtained from the atmospheric sciences (including climate science), oceanography, solar physics, space weather, Earth science, and computational science. For each of these disciplines, we have solicited information from sub-disciplines conducting research at various scales. Further, to ensure a representative cross section, we sought researchers using or developing multi-component modeling systems for use in fields such as climate and space weather.

In soliciting this input, scientists were asked to focus not only on their current capabilities and needs but also to envision how having access to petascale resources would influence their research objectives over the next five years. Apart from introductory material in sections 2.1 through 2.3, the description of these objectives and plans is in the scientists' own words.

This survey of science objectives builds upon, updates, and refocuses the 2005 assessment of the community documented in a pair of NSF-sponsored reports by the Ad Hoc Committee and Technical Working Group for a Petascale Collaboratory for the Geosciences, entitled *Establishing a Petascale Collaboratory for the Geosciences: Science Frontiers* and its companion *Technical and Budgetary Prospectus* (Bryan et al., 2005, Loft et al., 2005). While the scope of these reports was a bit broader than the mission contemplated for the NWSC, they still present a valuable and comprehensive analysis of many of the same areas of scientific study. In particular, the reports summarize many objectives that, as will be seen, still emerge from discussions with researchers today. Quoting from Bryan et al., 2005, petascale computing is a critical computing scale that

> *…will permit the computational geosciences community to dramatically expand the range of processes and spatial and temporal scales represented in simulations, making it possible to understand how the many Earth system elements interact to produce the complex phenomena that shape Earth's past, present, and future.*

These findings are echoed in another report focused on atmospheric and climate science: the 2008 NRC National Academies report entitled *The Potential Impact of High-End Capability Computing on Four Illustrative Fields of Science and Engineering.* This report documents and prioritizes the grand challenges in atmospheric science and the role of computing in attacking them (Lyons, J.W., et al., 2008). The committee classified the following eight grand challenges in atmospheric science as critically dependent on advancement of high-end computing capabilities, denoted with [1], or somewhat dependent, denoted with [2]. These challenges are listed below and matched to the community science objectives described in this proposal (the section number is indicated in parentheses, where applicable):

- [1] Extend the range, accuracy, and utility of weather prediction (2.4.3)
- [1] Improve understanding and timely prediction of severe weather (2.4.3), pollution, and climate events (2.4.1)
- [1] Improve understanding and prediction of seasonal-, decadal-, and century-scale climate variation on global, regional, and local scales (2.4.2)
- [1] Create the ability to accurately predict global climate and carbon-cycle response to forcing scenarios over the next 100 years (2.4.1)

- [1] Understand the atmospheric forcing and feedbacks associated with moisture and chemical exchange at Earth's surface (2.4.3)
- [2] Understand the physics and dynamics of clouds, aerosols, and precipitation (2.4.4)
- [2] Develop a theoretical understanding of nonlinearities and tipping points in weather and climate systems
- [2] Model and understand the physics of the Ice Ages, including embedded abrupt climate change events such as the Younger Dryas, Heinrich, and Dansgaard-Oeschger events (2.4.1)

Finally, a 2008 DOE-sponsored workshop report, entitled *Scientific Grand Challenges: Challenges in Climate Change Science and the Role of Computing at the Extreme Scale* focused even more tightly on climate science (Washington et al., 2008). This report assessed the research needs and opportunities for climate change science over the next two decades. It reiterates the importance of petascale and ultimately exascale computing to many of the climate research priorities of:

- How will sea level, sea-ice coverage, and ocean circulation change as the climate changes?
- How will the distribution and cycling of water, ice, and clouds change with global warming?
- How will extreme weather and climate change on the local and regional scales?
- How do the carbon, methane, and nitrogen cycles interact with climate change?

A principal finding of this report, consistent with the other community input discussed above, is that "An increase in the spatial and temporal resolution and the need for accelerated throughput will require expanded high-end computing resources at the extreme scale."

The Community Science Objectives sections that follow begin with background information on the societal impacts of Earth System phenomena in Section 2.1. Section 2.2 presents the research tradeoff dimensions of Earth System modeling, then Section 2.3 describes the role and scope of community models and modeling infrastructure in Earth System science research. Sections 2.4–2.7 present the science objectives from four major sub-disciplines: atmospheric science (2.4), space weather (2.5), oceanography (2.6), and Earth science (2.7). Finally, Section 2.8 discusses the role of computational science[9] as a crosscutting enabler of scientific advancement. The perspective that computational science research in support of geoscience modeling is itself a legitimate research objective for petascale systems is derived by community input (Snavely, et al., 2008, Washington et al., 2008), which strongly emphasized the importance of the participation of computational scientists in the teams that develop geoscience application software for massively parallel petascale systems.

[9] Computational science is defined as the application of computer simulation and other forms of computing to problems in various scientific disciplines. Computational science involves contributions from applied mathematics, computer science, and software engineering.

2.1 Societal Impacts

U.S. leadership in geosciences disciplines is vital to the nation's well being. U.S. economic competitiveness and national security hinge on making correct, scientifically informed decisions about national policies regarding the Earth System. Discoveries in this field will enable scientists to better inform decision-makers who formulate energy, environmental, agricultural, and natural resource management policies.

The sense of urgency conveyed by Earth System scientists through this rising chorus of reports is motivated by a spectrum of societal concerns ranging from the devastating, immediate, and localized impacts of natural disasters to the alarming long-term and global threats from a shifting climate revealed in the most recent IPCC Fourth Assessment report (Parry et al., 2007).

Earth System science research is central to developing tools used for forecasting, evaluating the risk of, and mitigating the effects of natural disasters. The impact and importance of enhancing these efforts are underscored by the human tolls and economic impact of recent natural disasters:

- The European heat wave in August 2003 caused over $12 billion in uninsured agricultural losses, $1.5 billion in forest fire damage, and an estimated 22,000–35,000 premature deaths Europe-wide.

- The August 2005 Hurricane Katrina caused $40.6 billion in insured losses and 1,353 direct fatalities. Total losses to the economy of the United States may have exceeded $100 billion (Johnson et al., 2006).

- From August to November 2004, nine hurricanes raked the Caribbean. At least 2,000 people were killed and hundreds of thousands were left homeless (Walter et al., 2005).

- A single natural disaster, the December 2004 Indian Ocean tsunami that claimed at least 225,000 lives, underscores the vulnerability of people living in coastal regions (Inderfurth et al., 2005).

- Space weather incidents are blamed for causing substantial financial losses. For example, the March 13, 1989 collapse of the Hydro-Quebec power grid was precipitated by geomagnetically induced currents. In the outage, more than 5 million people lost electrical power for 12 hours (Odenwald, 2009a). According to a study by Oak Ridge National Laboratory, a similar event in the United States could result in economic losses upward of $6 billion (Odenwald, 2009b).

The connection between extreme weather and climate change. A changing climate system will cause changes in the frequency and intensity of extreme weather events (defined as those events in the top or bottom 10% of all events) (Ramanathan et al., 2009, p. 25). Integral elements of society such as building codes (e.g., the definition of the 100-year flood plain) and the premiums charged by the insurance industry rely on an accurate understanding of the probabilities of such events. There is evidence that the probabilities of extreme events are changing. For example, according to Chapter 2 of Ramanathan et al., (2009) the following events have been documented over the past several decades:

- The number of heat waves and warm nights has increased, while cold days and nights and frost have decreased. The latter is a contributing factor to the spread of the pine bark beetle in the forests of the western United States.

- The incidence of drought has increased, especially in Africa, Australia, the southwestern United States, and Southern Asia.

- A statistically significant increase in the number of heavy and very heavy rainfall events in the continental United States, and the associated increased chance of flooding events.

- Intense extratropical cyclones appear to be increasing in number and strength, and their tracks have been shifting northward.

The slow-motion disaster. Beyond severe weather events that receive a lot of public attention, many impacts of climate change are predicted to be more gradual but no less costly and devastating. These effects, summarized here and described in greater detail in Ramanathan et al., (2009) and CCSP (2008) for example, include:

- *Water.* Significant degradation is predicted in the availability of fresh water resources, even in developed countries.

- *Food.* Alterations in regional climate will affect agricultural production and will disproportionally impact subsistence farmers in developing regions, particularly in the tropics.

- *Ecology.* Steep declines in biodiversity and in ecosystem resilience are foreseen. The world's coral reefs – and the fisheries they support – are particularly vulnerable to CO_2-induced changes to the ocean's pH. In the tropics, feedbacks between drought, forest mortality, fire, and land-clearing practices may lead to a shift to drier climates there, and to the collapse of biodiversity in rainforests.

- *Coasts.* Sea level rise will be gradual, but will continue for many centuries into the future: its impacts and cost to society will be amplified by the large proportion of the world's population living in coastal regions.

- *Health.* Direct effects of climate change on health include increased mortality from heat waves, floods, and severe storms. Indirect effects include changes in the prevalence of some "tropical" disease vectors in temperate regions: for example, the risk of malaria returning to the United States is increased, and the incidence of Dengue fever and Lyme disease may increase, following changes in the distribution of their hosts.

Better climate simulations at higher resolutions are expected to provide the basis for improving our understanding of these issues. Such fundamental knowledge is essential for effective national and international greenhouse mitigation and adaptation efforts to be implemented. Furthermore, the benefits of improved predictions could amount to billions of dollars. The 2006 Stern Review in the UK (Stern, 2006) supports this assertion:

The accuracy of climate predictions is limited by computing power. … It is important to continue the active research and development of more powerful climate models to reduce the remaining uncertainties in climate projections.

And the *Hadley Centre Review Supercomputing Final Report* (Hewitt et al., 2007) states:

The costs to the UK of adapting its basic infrastructure ... to climate change could run to many hundreds of billions. Without better information on future climate and climatic extremes, the UK is likely either to lose or to waste many billions per year through under – or over – investment in adaptation.

The magnitude of these changes, and the remaining uncertainties around these predictions, are leading to rising demands from resource managers and policy makers for more precise scientific answers to a variety of specific questions ranging from the regional impacts of climate change and the advisability of specific climate adaptation and mitigation strategies. Responding to this call for scientific knowledge to support informed decisions will require an enhanced and sustained investment across a broad front of Earth System science research topics – including enhanced observational facilities and computational capabilities – in the coming decades (Washington et al., 2008, Shukla et al., 2008, Ramanathan et al., 2009).

As will be shown, bridging this gap between scientific demands for, and the availability of, such resources will not only enable higher-resolution, more-accurate predictions with lower levels of uncertainty, but it will also open new scientific frontiers by enabling increasingly multidisciplinary simulations with the inclusion of more realistic physical, chemical, and biological phenomena. By addressing the existing gap between modeling needs and current computational capabilities, the geosciences community will be able to accelerate the progress of research and open new scientific frontiers in all areas of Earth System simulation.

2.2 The Tradeoff Dimensions of Earth System Modeling

Advances in numerical simulations in the geosciences are generally sought in six areas, or dimensions, that are discussed in this section: increased spatial resolution (Resolution), the inclusion of new processes and phenomena plus the inclusion of more diverse processes (New Science), higher-fidelity physics (Better Science), increased capabilities for data assimilation (Data Assimilation), increased ensemble sizes for better statistics (Ensembles), and execution over longer timescales (Timescales). Figure 2.1 (after Buja and Palmer) illustrates these tradeoffs graphically. Each concentric hexagon in the diagram represents a level of sustained computational power, from the terascale (recent past), through the petascale (by 2010), to the exascale (predicted for 2018). For a particular computational level – and assuming the time to solution is fixed by practical or real-time limitations – advances in any one of these six dimensions (each a vertex of the hexagons) must come at the expense of the others. Depending on their specific research interests, different modelers will make different choices in how to expend an increment of available resources along the different axes. The ability to open new scientific frontiers by advancing all areas simultaneously would require both radical advances in current computational capabilities and the development of novel enabling algorithms. The latter include nested and adaptive mesh techniques that focus computational expense only when and where needed.

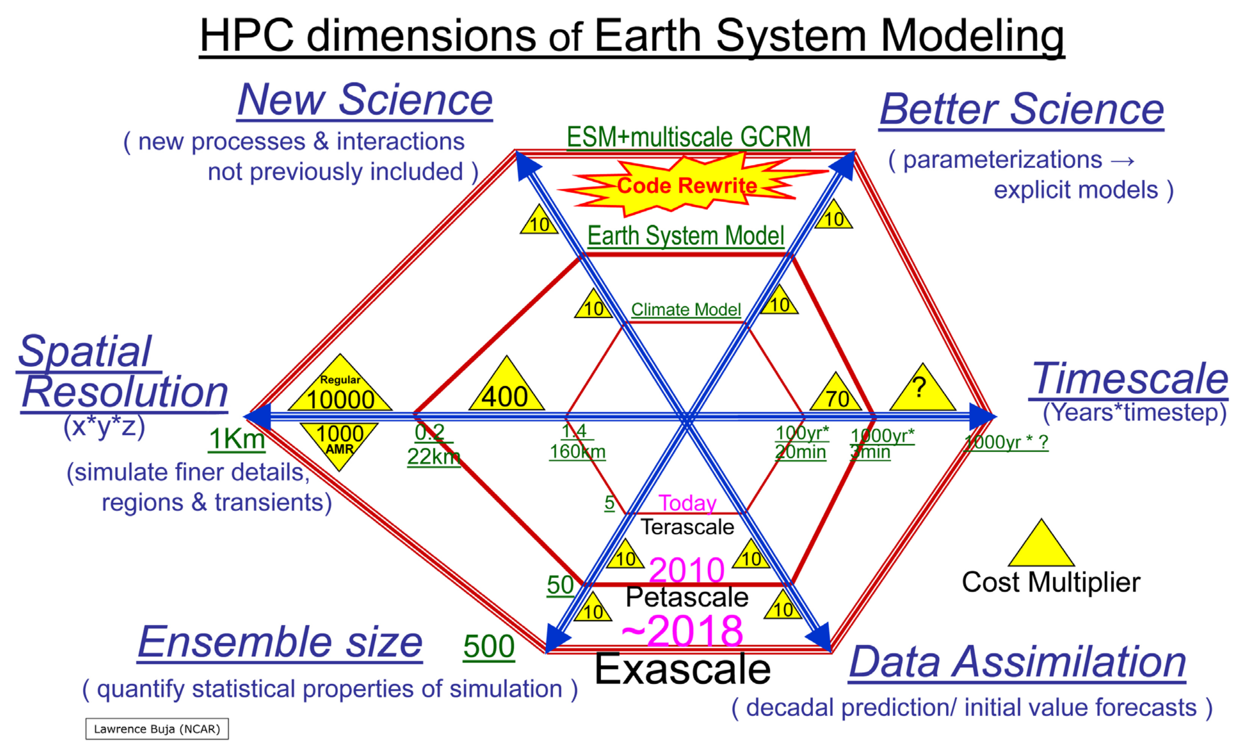

Figure 2.1. The six tradeoff dimensions of Earth System science (with the cost estimates and objectives taken from climate modeling) for each dimension moving from terascale to petascale to exascale computing. For increasing levels of computing, the capabilities will be increased along each of these dimensions according to the research priorities of a specific community.

Resolution. The Earth System exhibits a wide range of spatial and temporal scales that interact in complex nonlinear ways: these in turn drive the need for higher "resolution" of these scales in virtually all geoscience simulations. Increased model resolution produces not only more accurate simulations of known physical processes, it also has the compounding effect of resolving more accurate and correct physics that must otherwise be modeled or parameterized with great uncertainty (Pope, 2000, Sagaut, 2002, Randall, 2003). Achieving increases in model spatial resolution generally requires rapid growth in the required computational power, which scales as the fourth power of the increase in spatial resolution[10]. For problems with wide ranges of spatial scales, algorithmic techniques such as adaptive mesh refinement can be used to extend the range of resolution for a given level of computational power (Skamarock, 1989).

Timescale. The Earth System evolves on drastically disparate timescales – from seconds to millenia. For example, ocean turnover, biogeochemical systems such as the carbon cycle, and continental ice sheets all evolve over millennia: such models coupled to a rapidly evolving atmospheric model may require millions of timesteps to complete a single year of simulation.

New Science. Because most geoscience problems involve a wide range of disparate but interacting physical phenomena, most simulations include only a subset of the known important

[10] This power law assumes the problem is fully three-dimensional. The fourth power comes from the need to shorten the numerical timestep as the grid spacing is increased, a mathematical requirement derived from the so-called Courant-Friedrichs-Levy (CFL) condition. Thin layer problems, such as the atmosphere and ocean, are quasi two-dimensional, and scale as the third-power with resolution, once the third dimension has been adequately resolved.

physical phenomena due to computational constraints. The availability of substantially more capable computational hardware will enable the inclusion of more diverse and accurate physical processes in models. For example, climate simulations will enable "New Science" by including robust physical models for atmospheric chemistry, detailed ocean salinity, a full stratosphere, and biogeochemical feedback effects. The large computational requirements inherent in the inclusion of additional physics is well illustrated in the area of atmospheric chemistry, where the inclusion of hundreds of species and reaction rates is seen as necessary for accurate modeling of aerosol mixing (Rodriguez and Dabdub, 2004). These phenomena may also introduce additional scales, such as fast-reaction timescales in atmospheric chemistry, which in turn may require more sophisticated numerical algorithms (e.g., time-implicit methods) to remain computationally tractable (Butcher, 2003).

Better Science. An example of such "Better Science" is large eddy simulation (LES) of turbulence, which is predicated on the existence of a range of small "universal" eddies that are insensitive to the macroscopic problem parameters (Kolmogorov, 1941). LES methods require resolving all turbulence eddies down to the small universal range, which can then be modeled using a sub-grid-scale model. LES methods can thus result in very demanding resolution requirements but entail a much lower degree of uncertainty in the modeling of unresolved physics due to the universal nature of the unresolved eddies (Smagorinsky, 1963).

Data assimilation. Data assimilation techniques are now well established as an integral tool for enhancing simulation accuracy and reliability (Kalnay, 2003), but the computational expense of these methods has limited their use in current high-end simulation efforts. The amount of available remote sensing data is expected to grow by a factor of 10^6 over the next decade (Bryan et al., 2005), and the effective use of this large volume of data will require significant increases in computational capabilities. Clearly, evolving and applying data assimilation techniques to more complex simulation problems, while making effective use of the rapidly growing volume of observational data, represents a formidable computational task. Additionally, algorithmic advances will be required to advance the state-of-the-practice in data assimilation problems, for example, through the development of more efficient and effective optimization techniques rooted in applied mathematics and control theory (Nocedal and Wright, 1999).

Ensemble size. Earth System simulations are probabilistic in nature, partly due to the chaotic nature of geo-turbulence and the uncertainties in the experimental data and approximations used in physical modeling. Quantification of simulation uncertainty is arguably at least as important as the simulation outcome itself (Kalnay, 2003, Molteni et al., 1996). "Ensemble" runs are often used to quantify these uncertainties, using large numbers of repeated simulation runs with slightly modified inputs to assess variations in the predicted outputs. Increasing the number of ensemble runs constitutes a straightforward approach to obtaining more accurate estimates of the simulation uncertainty. While increasing the number of ensemble runs is an inherently parallel procedure that scales linearly with the number of runs, the ensemble sizes grow geometrically with the number of relevant input parameters. The number of parameters increases rapidly for coupled interdisciplinary simulations. The need to drastically increase the number of ensemble runs while at the same time undertaking more complex simulation problems in itself exacerbates the gap between the required and currently available computing power. More advanced techniques rooted in stochastic methods for calculating simulation uncertainties are also appearing, but these promise to be even more computationally intensive (Palmer, 2001, Majda and Khouider, 2002, Heinz, 2003).

2.3 The Importance of Community Modeling and Modeling Frameworks

Community models play a pivotal role in Earth System modeling. As new components and physical processes have been added to Earth System models, they have correspondingly become increasingly complex software systems. The interdisciplinary nature of this complexity is such that many scientists have found it convenient and productive to pool their resources and work together to deploy community models and modeling frameworks. This strategy leverages the work of others, reduces duplication of effort, and ultimately the turnaround time between first asking a scientific question and obtaining an answer. Self-describing data and model components promise to reduce or eliminate errors – such as using different units for the same quantity – that might otherwise go unnoticed or prove difficult to debug. Community models and modeling frameworks are thus valuable tools that enable scientists to explore more complex and ambitious research problems than they could accomplish individually, and therefore increase the collective impact of the community's research effort. For these reasons, adapting community models and modeling frameworks to efficiently use petascale resources is a critical enabling strategy for Earth System modeling research.

There are many examples of community models and modeling frameworks in use in the Earth System science community that are already being employed on large-scale systems. Two examples of community models and one modeling framework are described here.

The Community Climate System Model (CCSM). A fully coupled global climate model, CCSM provides state-of-the-art computer simulations of the Earth's past, present, and future climate states. The CCSM has about 400 users and is sponsored by the National Science Foundation (NSF) and the U.S. Department of Energy (DOE), and is maintained and administered by the Climate and Global Dynamics Division (CGD) at NCAR. High-resolution configurations of CCSM have been run in production on 7,700 processors on the Atlas system at Lawrence Livermore National Laboratory, and on nearly 6,000 processors on the Kraken system at the National Institute for Computational Sciences (NICS).

The Weather Research and Forecast Model (WRF). A next-generation mesoscale numerical weather prediction system, WRF is designed to serve both operational forecasting and atmospheric research needs. It features multiple dynamical cores, a 3-dimensional variational (3DVAR) data assimilation system, and a software architecture allowing for computational parallelism and system extensibility. WRF is suitable for a broad spectrum of meteorological applications across scales ranging from meters to thousands of kilometers. WRF allows researchers the ability to conduct simulations reflecting either real data or idealized configurations. WRF provides operational forecasters with a model that is computationally flexible and efficient, while offering the advances in physics, numerics, and data assimilation contributed by the research community. A WRF "nature run" was performed on a 65,536-processor Blue Gene/L system located at IBM's Watson Research Center (Michalakes et al., 2007).

WRF has a rapidly growing user base. Of the 10,432 registered users in June 2009, roughly 65% were foreign. Annual WRF workshops and tutorials are held at NCAR. WRF is currently in operational use at the National Centers for Environmental Prediction (NCEP). The effort to develop WRF is a collaborative partnership, principally between NCAR, the National Oceanic and Atmospheric Administration (NOAA), NCEP, the Forecast Systems Laboratory (FSL), the Air Force Weather Agency (AFWA), the Naval Research Laboratory, the University of Oklahoma, and the Federal Aviation Administration (FAA).

The Earth System Modeling Framework (ESMF). ESMF is a high-performance, flexible software infrastructure to increase ease of use, performance portability, interoperability, and reuse in climate, numerical weather prediction, data assimilation, and other Earth science applications. ESMF provides an architecture for building complex, coupled modeling systems, including data structures and utilities for developing individual models. ESMF focuses on the benefits of defining modular components. Components are a unit of software composition that has a coherent function and a standard calling interface and behavior. Components can be assembled to create multiple applications, and different implementations of a component may be available. In ESMF, a component may be a physical domain, or a function such as a coupler or I/O system. ESMF also includes toolkits for building components and applications, such as regridding software, calendar management, logging, error handling, and parallel communications. The scalability of ESMF bundle redistributions has been tested on up to 2,048 processors[11].

ESMF has over 30 Earth System applications using its coupling superstructure and over 4,000 registered downloads of its software as of January 2009. ESMF is the result of interagency collaboration between NASA, NOAA, NSF, and the Department of Defense.

2.4 Atmospheric Science

Dramatic improvements in computing power – along with rapid advances in disk capacity, tape densities, and network bandwidths – are transforming how society generates and consumes information. Since the early days of computing, the atmospheric sciences have led in the adoption of computing technology for scientific research and continue to do so today[12]. The explosion of digital observational data about the atmosphere, coupled with once-unimaginable computer modeling capabilities, are allowing atmospheric scientists to tackle a broad range of complex, interdisciplinary, grand-challenge problems in new and more realistic ways. For example, climate scientists who were once confined to performing low-resolution simulations using atmospheric general circulation models with prescribed sea surface temperatures now work with dynamically coupled models of sea ice, ocean, and land surface processes (Collins et al., 2006, Meehl et al., 2006). Additional processes such as ocean ecosystems, atmospheric chemistry, dynamic vegetation, land ice, and the carbon-nitrogen cycles are being included in climate simulations, and model resolutions are being dramatically increased in exploratory simulations of regional climate. Meteorologists are planning to improve resolution of convective storms and hurricane cores to levels of tens of meters to improve the understanding and predictive skill of these severe weather events, and cloud physics researchers are seeking to unify large-eddy and direct numerical simulation to provide better descriptions of the fundamental microscale processes in clouds that can improve their representation in climate and weather applications.

[11] See http://www.esmf.ucar.edu/metrics/performance/timing_0707_bundle_redist.pdf for details.
[12] For example, the first successful numerical weather prediction experiment was performed by Jule Charney and John von Neumann and others, in 1950, on the ENIAC digital computer.

2.4.1 Climate Simulation and Prediction

Climate prediction is a major science driver requiring enhanced computing, and the specific example for which Figure 2.1 was prepared. With the wide public acceptance of the findings from the IPCC AR4 report, the climate science community is moving beyond simply running the basic SRES[13] scenarios proving that global warming is occurring; they are now focusing on new areas of climate change decision support. This will involve putting climate change in the context of the evolving natural and anthropogenic environmental systems (i.e. "new science" dimensions) to assess the adaptation and mitigation implications of local, national, and international policies on emissions, energy use, and resource management. Providing climate knowledge-support for society's decision makers will require developing a climate prediction capability to understand and accurately simulate our complex energy-chemistry-climate system. Such a capability still needs to be developed and demonstrated, but it will surely push both the "data assimilation" and "ensemble" dimensions of Figure 2.1.

At the same time, researchers remain concerned that climate forcings are only on the order of 1 W/m^2, while known biases in present-day climate models are order 10 W/m^2. Because of this basic climate problem, climate projections rely on a number of assumptions such as: when integrated over long times, small forcing signals will emerge; tuning for global balance is sufficient despite large regional biases; and the parameterizations and compensating errors that produce the observed records over the past century will hold over to the next. Such assumptions must be transformed into knowledge, and a necessary ingredient is increased computing power so that the "better science" of Figure 2.1 can be obtained by having all the important climate processes either be explicitly "resolved" or faithfully parameterized (e.g., see Hack et al., 2006, for improvements in CCSM fidelity realized by increasing resolution). Through this effort a greater understanding of the Earth's climate system will be derived from both new observations and better modeling. The next-generation climate models will be true Earth System models bringing in "New Science" biogeochemical, ecological, and economic model components. As scientists seek to simulate the full Earth System, the complexity and cost of climate system models is expected increase in the other climate prediction dimensions too (see Figure 2.1). Spatial resolution will increase even further, allowing the formation of fine-scale features and phenomena that were previously unresolved. Simulation complexity, in terms of both completely new processes and better representation of existing processes, will have a substantial impact in the improving model fidelity. Of lesser impact will be the improvements in the time scales, ensemble sizes, and data assimilation cost dimensions.

Climate modeling will benefit from the availability of dedicated high-capability computing resources in two distinct ways. First, increased simulation resolution and more detailed physical parameterizations will be made possible by greater amounts of raw computational power. Second, the development of advanced algorithms will be enabled through research performed using more capable hardware. For example, an effort to increase simulation resolution, made possible by greater amounts of raw computational power, will require both more timesteps and more detailed physical parameterizations. This must also be accompanied by the development of advanced algorithms that will be enabled through computational science research performed using more capable hardware. Global Earth System Models (ESM), with dynamic global carbon and nitrogen cycles, terrestrial and oceanic

[13] SRES = Special Report on Emission Scenarios describes the carbon emission scenarios used in the IPCC Assessments.

biogeochemistry, and land cover/use, will be needed by 2012 to study scenarios of bioenergy and food production in a changing climate. By 2018, with even more powerful exascale computers, it should be possible to conduct multi-year simulation of ESMs operating at cloud-resolving scales globally, with well-characterized precipitation and hydrology. However, as the architectures of exascale computers become more clearly defined, the design of many of the "legacy" climate codes in use today will need to be radically modified.

To be of maximum utility to decision makers, climate models must provide regional detail and be able to resolve and predict the frequency of significant events such as hurricanes and typhoons (Drake et al., 2005). Additionally, these models must have the capability to predict the likelihood of extreme climate fluctuations and abrupt climate change, the response of climate and ecosystems to changes in climate forcing within the coming decades, the availability of natural water resources (particularly the likelihood, intensity, and duration of droughts), the influence of various land-use scenarios on regional climate, the distribution and intensity of atmospheric dust, the impact of emission reduction strategies on regional climate scales, and the viability of different carbon sequestration scenarios (Kinter et al., 2004). To meet these requirements, climate models must incorporate greatly increased spatial resolution as well as improved treatments of sub-grid-scale physical, chemical, and biological phenomena in each component model. Studying extreme (rare) events will require better statistics, which will require increases in both the number of ensemble members and the lengths of simulations.

A current example of climate science pushing in the ensemble direction is an NSF PetaApps program-funded project to determine the impact of unresolved scales on climate predictions via the use of interactive ensembles, which has been allocated 35M CPU-hours on a Teragrid XT5 system, Kraken, at NICS. The objective of this project to carry out a variety of high-resolution interactive ensemble experiments using both multiple atmosphere and multiple sea ice instantiations. The project has also recently been selected as an application for the NSF Track-1 Blue Waters system, to carry out experiments using multiple atmosphere and ocean instantiations at ultra-high resolution.

Perhaps the strongest drivers in the integration length dimension of climate science are paleoclimate applications, with glacial transitions a present research focus. A DOE Innovative and Novel Computational Impact on Theory and Experiment (INCITE) project to simulate the last glacial transition from 25,000 to 1,000 years ago, was allocated 840K CPU-hours on Phoenix and 13.7M processor-hours on Jaguar. A similar set of numerical experiments (and hence similar allocation) for the previous glacial-interglacial cycle is also of great scientific interest, because the previous interglacial was the last time the Arctic experienced summer temperatures significantly warmer than those of the 20th century and when sea level was 4-6 meters higher than present.

An example of an experiment pushing the state of the art in the resolution direction is a CCSM simulation with an eddy-resolving 0.1° ocean coupled to a 0.25° atmosphere that was recently performed at Lawrence Livermore National Laboratory (LLNL) on the Atlas Linux Cluster. A single simulated year in this configuration required about 230,000 CPU-hours, about 500 times that for a standard IPCC configuration with a nominal 1° ocean and 1° atmosphere configuration.

Beyond these current efforts, the computational cost of improving certain science components is reasonably well known. For example, extending the atmosphere model into the stratosphere increases the overall cost of a coupled simulation by about a factor of four. Modeling the carbon cycle is very costly (a high-priority challenge identified by Lyons, et al., 2008). The expense of the carbon cycle comes from the ocean ecosystem model, which alone increases ocean model computational costs by a factor of three. Further, the ecosystem requires about 3,000 years to equilibrate, increasing the timescale dimension as well.

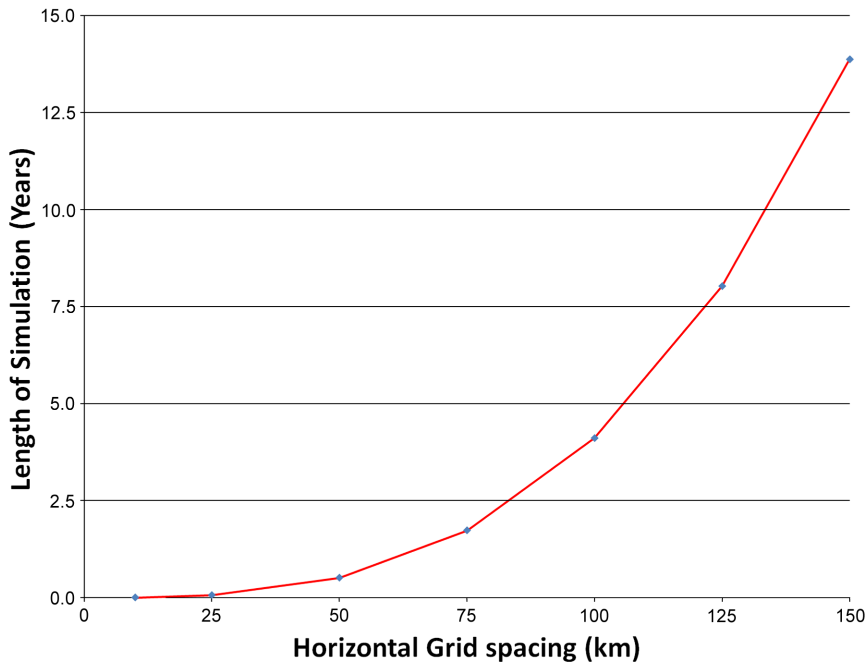

Figure 2.2. The tradeoff between horizontal resolution vs. simulation length. A low-resolution model (high horizontal grid spacing) can simulate significantly more model years than a high-resolution model (low horizontal grid spacing) in a given period of wall-clock time (after Michalakes et al., 1995).

Advances in decadal predictions will involve ocean initialization and integration and will push simulation costs in both the ensemble and data assimilation dimensions. Preliminary indications are that the former increases the ocean computational cost by an order of magnitude, while the latter adds about another factor of two.

Looking ahead a decade or more, the target spatial resolutions for the cloud-resolving atmospheric and eddy-permitting oceanic simulation components of coupled climate system models have been estimated to be 1 km and 10 km, respectively. By comparison, current models utilize spatial resolutions on the order of 100 km. To carry out long simulations at these high resolutions, increases in computational power by at least three orders of magnitude (i.e., 10 petaflops sustained) over capabilities provided by the NWSC's 2012 petascale system will be needed (see Figure 2.2) to maintain adequate integration rates[14] (Wehner et al., 2008). In addition the inclusion of more complete physics, by adding important interactive physical, chemical, and biological processes to each component model, will be required to gain a better understanding of the ecological implications of climate change. Added phenomenology of this kind has the potential to add several more orders of magnitudes to the computational complexity of these models. As a better understanding of underlying physical phenomena is obtained (often as the result of observational programs), the fidelity of coupled climate

[14] The climate modeling community expresses simulation throughput in simulated years per wall clock day. As outlined by the Ad Hoc Committee and Technical Working Group for a Petascale Collaboratory for the Geosciences, climate researchers have expressed a goal of achieving overnight turnaround of a standard 20-year atmospheric run (Bryan et al., 2005).

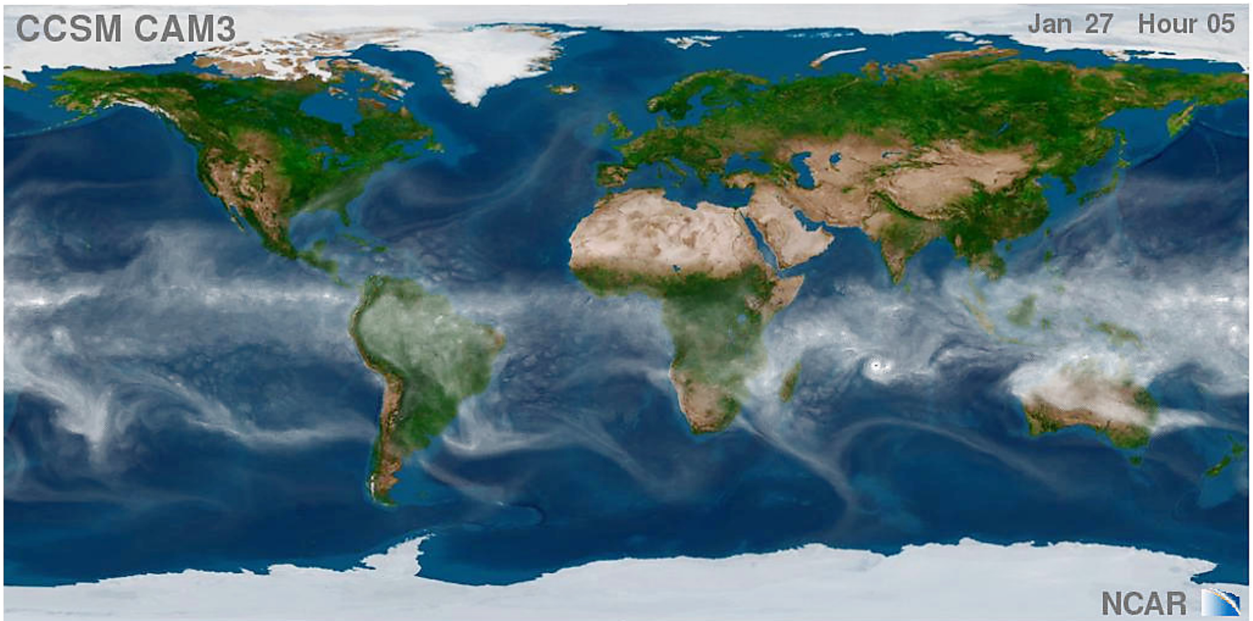

Figure 2.3. A high-resolution 40-km CAM run at NCAR produced a detailed representation of a cyclone in the southern Indian Ocean. Accurate storm structure – indicated by the eye-like feature – and faithful representation of storm frequency and intensity are key goals of climate modeling.

models will need to be increased by the replacement of existing parameterizations of sub-grid-scale physical processes with more realistic and accurate treatments. This will require more sophisticated algorithms and even greater amounts of computational resources. Similarly, assimilation into climate system models of the ever-expanding observational datasets for the atmosphere and oceans provided by remote sensing systems rely on researchers having access to robust new algorithms and even larger high-end computing resources (Lyons et al., 2008, p. 59). Until computer and numerical methods catch up to these extraordinary demands, local refinement techniques, such as those discussed in Section 2.4.2 on Nested Regional Climate Modeling, as well as other modeling compromises, must be used as interim solutions.

2.4.2 Nested Regional Climate Modeling

Nested Regional Climate Modeling (NRCM) provides the opportunity to advance further and in more dimensions of Figure 2.1 regionally, than are possible globally. At NCAR, the NRCM unites the strengths of two leading tools of atmospheric science: the worldwide reach of a premier global climate model and the detail of the world's leading county- and city-scale weather model.

The sophistication of these models places great demands on even the world's most powerful computers. Global climate models typically track the atmosphere at points separated by 160 or more kilometers, so they cannot easily depict such features as hurricanes or thunderstorm complexes. Weather models can narrow the resolution to 1.5 kilometers or less, but they cannot span the globe with this capacity or cover periods longer than a few days.

As we adapt to ongoing climate change and work to reduce our climate-warming emissions, we need a clearer sense of how the atmosphere might respond in vulnerable locations, from the storm-plagued Gulf and Atlantic coasts to the water-scarce West. As concluded by the IPCC and the U.S. Climate Change Science Program, some risks are already clear:

- Extreme weather events such as heavy rainfall and heat waves will become more frequent.
- More precipitation will fall as rain instead of snow.
- Snow will melt sooner, increasing runoff and the risk of flooding in early spring.
- In summer, the risk of drought and wildfire will go up as the warmer air dries the land.
- Ocean temperatures in the tropical Atlantic will likely continue to warm, providing more fuel for hurricanes and making it easier for them to become intense when other conditions are favorable.

All of these changes have enormous implications for public safety, agriculture, water resources, and urban planning. But current climate models cannot portray the local processes or the complex topography required to predict such high-impact events on a local or regional scale.

The Nested Regional Climate Model can provide the next level of regional guidance for policymakers who must deal with tomorrow's climate. One of the strengths of NRCM is its ability to provide detail where and when it is most needed. For example, by using nested domains (as illustrated by focus clarity in Figure 2.4), NRCM scientists can downscale (zero in) on regions of particular interest, such as hurricane or drought-prone areas or mountainous regions, without the much higher cost of simulating the entire globe with such detail.

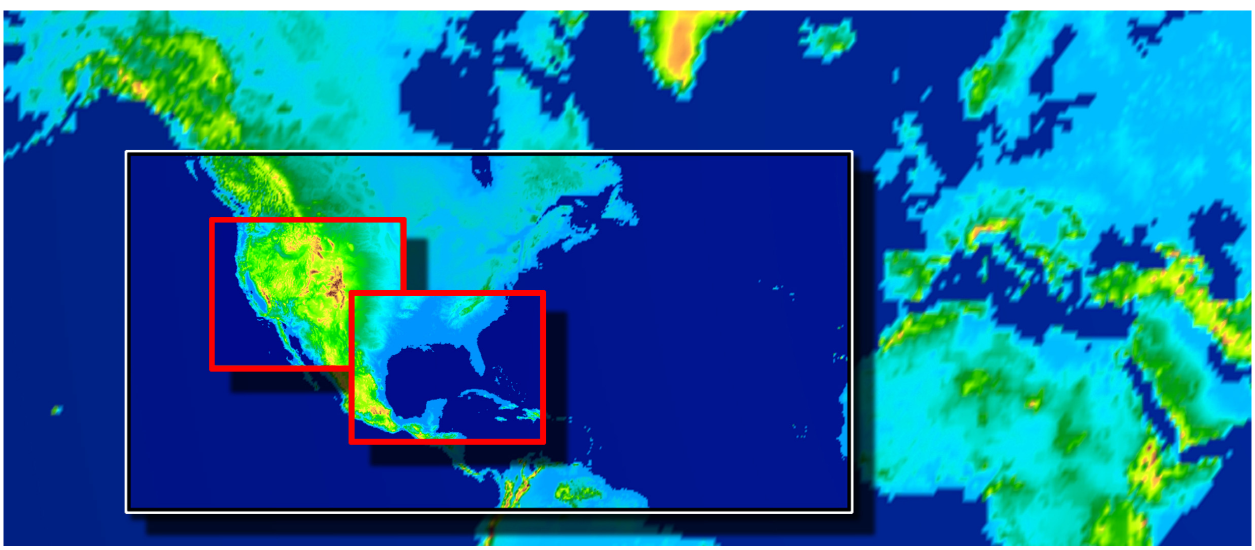

Figure 2.4. In one NRCM experiment, the atmosphere is simulated in three dimensions at resolutions ranging from about 30 kilometers between data points across a large part of the Northern Hemisphere to as fine as 4 kilometers in targeted areas of North America (red boxes). This strategy enables scientists to forecast future climate in detail for specific regions. This provides an adaptable tool for improving regional climate predictions without overtaxing supercomputing resources. (Contrast between coarse- and fine-resolution boxes has been increased for illustrative purposes.)

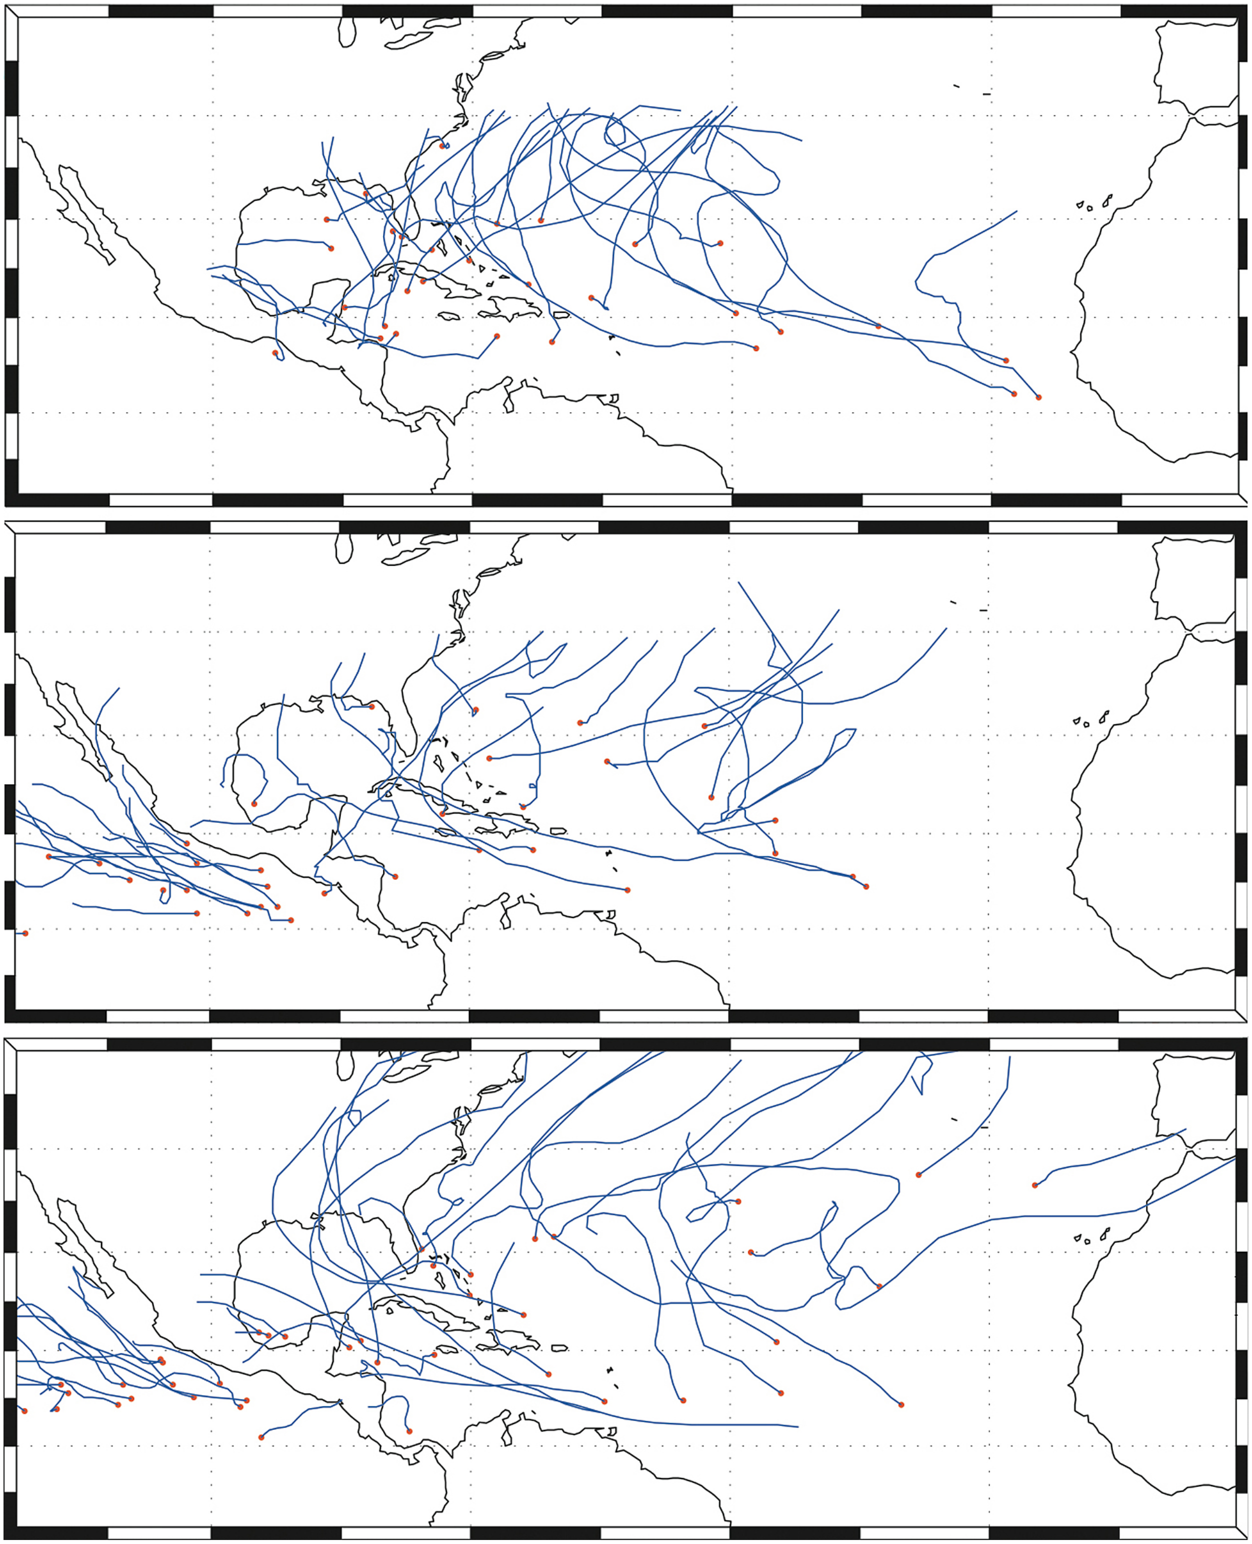

Figure 2.5. The NRCM has already demonstrated the increased skill possible through enhanced resolution in a reproduction of the 2005 Atlantic hurricane season, the busiest on record. When run with 12-kilometer spacing between grid points (top), the NRCM does a much better job than an identical simulation at grid resolution of about 35 kilometers (middle) in simulating the patterns of activity observed in the 2005 season (bottom), including the areas where the hurricanes form and make landfall.

The NRCM approach can also focus on especially important regions, such as the tropical Pacific Ocean, where small changes can lead to large upscaling effects across North America. This innovative approach bypasses the traditional horizontal-resolution/simulation-length tradeoff shown in Figure 2.2, resolving fine-scale, regional areas much sooner than will be possible through brute-force increases in the resolution of global climate models. Still, depending on the area refined and the degree of refinement, NRCM can potentially be orders of magnitude more computationally expensive than current climate models without regional nesting. However, these investments may be worth it, in terms of improved model performance, not only locally but also globally. For example, increasing the resolution of the coupled ocean off the west coast of North America has widespread impacts on tropical precipitation, and hence global atmospheric circulation. The improved global circulation patterns justify doubling the computational expense to capture regional details. Similar effects are expected from increasing ocean resolution off Chile/Peru and off southwest Africa. Numerical studies are a powerful tool in determining the underlying physical mechanisms. The signals are large (5 mm/day = 1.8 m/year), even in these 10-year seasonal averages.

With petascale computing resources, the NRCM will be able to simulate a variety of 21st-century climates, from which statistical portraits of the future atmosphere can be produced. The NRCM cannot tell us exactly where a drought, a hurricane, or a winter storm will strike 30 years from now. However, it can give us a better idea of how and where weather patterns are likely to shift from decade to decade and how specific high-impact weather events, such as hurricanes, may change in frequency, intensity, size, and rainfall. This is the guidance that stakeholders need to help manage resources, protect lives, and deal with changes in society and our environment.

2.4.3 Research on Extreme Weather

The notion of what constitutes a weather extreme has no unique definition, but in general the idea encompasses events that are relatively rare, at least from a local perspective, and typically produce a large adverse influence on human infrastructure and society. Modeling of extreme weather centers on two types of activities: 1) predicting individual events, either in a deterministic or probabilistic sense, and 2) simulating the statistics of extreme events over a time scale long enough that many events occur. Because many extreme events are often linked strongly with precipitation (or lack thereof in the case of droughts), accurate numerical representation of the dynamics of precipitation systems is an important aspect of simulating many extreme events. Other extreme events involve much more the physics of the planetary boundary layer (PBL), whether the issue is air quality, transport and dispersion of hazardous material, or intermittent turbulence in stable situations. Realistic simulation of all these events requires fine resolution, and the simulation of sufficient statistics requires long-time simulations as well. Following are three examples of key challenges.

Hurricanes. Over the last several decades, hurricane track forecasting has improved significantly whereas very little progress has been made in hurricane structure, wind intensity, and rainfall forecasts along the track. The shortcomings of predicting the structure of the inner core result in part from insufficient model resolution. It is currently not known precisely what resolution is required for a physically realistic simulation of inner-core structures. Recent results (Chen, 2006; Davis et al., 2008) have suggested that a grid spacing of 1-2 km in the horizontal and 50 m in the vertical might be necessary. However, the variability of inner-core structures

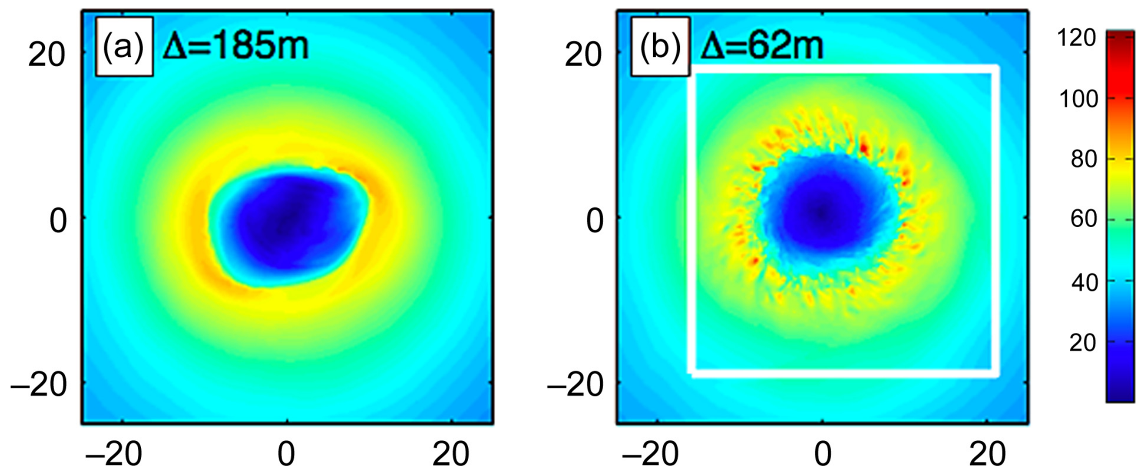

Figure 2.6. Snapshots of the distribution of wind speed in meters per second at the 10-meter level at t=9.75 days with grid intervals (a) 185 m and (b) 62 m, domain indicated by white box. Note the development of turbulent motions in the eye wall in (b). Axis labels indicate distances in kilometers. Color key for wind speed (in mph) appears at right.

appears sensitive to resolution. For example, recent idealized results with the WRF model run at an ultra-high 62 m horizontal grid spacing (see Figure 2.6) indicate that fully developed turbulence emerges at a grid spacing of about 100 meters (Rotunno et al., 2009). This simulation is roughly equivalent to an attempt to use LES technique to compute vertical and horizontal turbulence effects (i.e., non-PBL, no artificial filter) in the hurricane core. Even these simulations have not reached the point of convergence, such that theories about hurricane intensity can be tested with confidence. Such high-resolution simulations are needed to examine the intrinsic predictability of hurricane structures to help us better understand what structural aspects of hurricanes might be simulated deterministically versus stochastically. This bears directly on resources and methods devoted to improvement of data assimilation and initialization techniques, for instance. It also directly bears on the proper representation of hurricanes in coarser-resolution models that might be used for ensemble prediction, or even coarser simulations of future climate scenarios.

In addition to inadequacies of resolution, our lack of skill in forecasting hurricane intensity and attendant rainfall may be attributed to a lack of full coupling to the ocean and representation of ocean waves and detailed processes at the air-water interface. Extremely high winds around the eye, intense rainfall, large ocean surface waves, and copious sea spray push the surface-exchange parameters for temperature, water vapor, and momentum into regimes unreachable at current levels of model resolution and coupling. Within the hurricane core, there are competing processes associated with 1) the heat and moisture fluxes that fuel the storm, and 2) the dissipation of kinetic energy associated with wind stress on the ocean surface. Air-sea interaction is especially important in the region between the center of the eye and eyewall where there are extremely large gradients in the wind, temperature, and pressure fields. Proper treatment of the air-sea exchange processes may require additional modules specific to air-gas mixtures that characterize the air-sea transition, and it is currently an active research area to study the best way to model these processes and communicate the effects back to the atmospheric model.

In a similar way, there has been little direct modeling of the details in wind resulting from the land-atmosphere interaction in landfalling hurricanes. At present, turbulence is treated primarily with traditional formulations only suited for grid spacings of a few kilometers. However, statistical prediction of wind variability, including gustiness, and incorporation of these results into damage models is an emerging area of study. But realistic simulations of these effects will also require extremely high resolution, down to perhaps tens of meters.

Recent attempts to predict hurricane intensity in real cases have begun utilizing advanced data assimilation methods that are computationally intensive. Methods currently being investigated are variational and ensemble techniques, as well as hybrid approaches. To perform data assimilation at the same resolution as a single forward integration requires more than an order of magnitude increase in computational power. Ensemble data assimilation currently requires on order 100 members to adequately depict error covariance structures needed for the spatial influence of observations to be realistic.

The partly stochastic nature of hurricane intensity change also requires a probabilistic modeling approach. It is currently impractical to integrate 100 members to provide probabilistic intensity information using a 1-km mesh around the hurricane core for integrations of 5 days or more. This problem is important for quantifying predictability of hurricane intensity and understanding the interrelationship between hurricane track and intensity. Large ensembles can also be used to understand the sensitivity of the predicted state to initial perturbations.

Severe thunderstorms and tornadoes. Recent work has shown benefits from performing simulations using grid increments such that cumulus parameterization can be removed (e.g., 1-4 km, Kain et al., 2008; Weisman et al., 2008). It has also been demonstrated that more realistic simulation of moist convection requires grid increments around 100 m (Bryan et al., 2003). Further refinement is needed, perhaps as fine as 20 m, to simulate tornadoes (Fiedler, 2009). The need to simulate multiple storms within a convective cluster that may span hundreds of kilometers, and to simulate the severe weather associated with these storms places computing requirements into the petascale range.

Many of the issues with initialization and predictability discussed for hurricanes also apply to convective storms. Ensemble and variational data assimilation methods have been developed for convective storms, but these require roughly 100 times the computing power of a single forecast. However, these methods are critical to realistically initializing storm structure such that accurate short-range (1-3 hour) prediction is possible. The need for estimating probability distributions of severe convection requires ensembles of 30-100 members integrated for 1-2 days. Even with modest resolution (e.g., 1 km horizontal and 300 m vertical), this rapidly becomes a petascale computing problem.

Boundary layer modeling. As the resolution of mesoscale models increases into the large-eddy simulation (LES) regime, LES modeling is approaching the mesoscale (Wyngaard, 2004). Mesoscale LES has important practical applications in the areas of transport and dispersion of pollutants and other aerosols. Emerging areas such as wind energy require detailed simulation of the vertical profile of turbulence within the lowest 200 m of the atmosphere. This requires resolving the large vertical shears and intermittent turbulence characteristic of nighttime conditions. These simulations require grid increments of 10-20 m, perhaps finer in the vertical. To extend such simulations to mesoscale domains of 100 km or more places this computational problem in the petascale range.

2.4.4 Cloud System Research

A key factor that limits the reliability of anthropogenic climate change modeling and retards the prediction of climate is our inability to accurately represent the complex effects of clouds on climate. The types of clouds and precipitation that develop as a result of multiscale interactions are critical, e.g., low, thick cumulus clouds have a cooling effect on climate while high, thin cirrus clouds trap long-wavelength radiation and warm the Earth. In Figure 2.7, the organized cumulonimbus at the top of the image have a significantly greater impact on the atmosphere than the cumulus clouds at the bottom. Despite many efforts, the clouds-in-climate problem has resisted solution. One reason is that, due to the limitations of computing power, global atmospheric models presently used for climate research have a coarse grid spacing (~100 km). Another reason is the difficult task of using parameterizations to represent the statistical interactions of populations of clouds with large-scale motions.

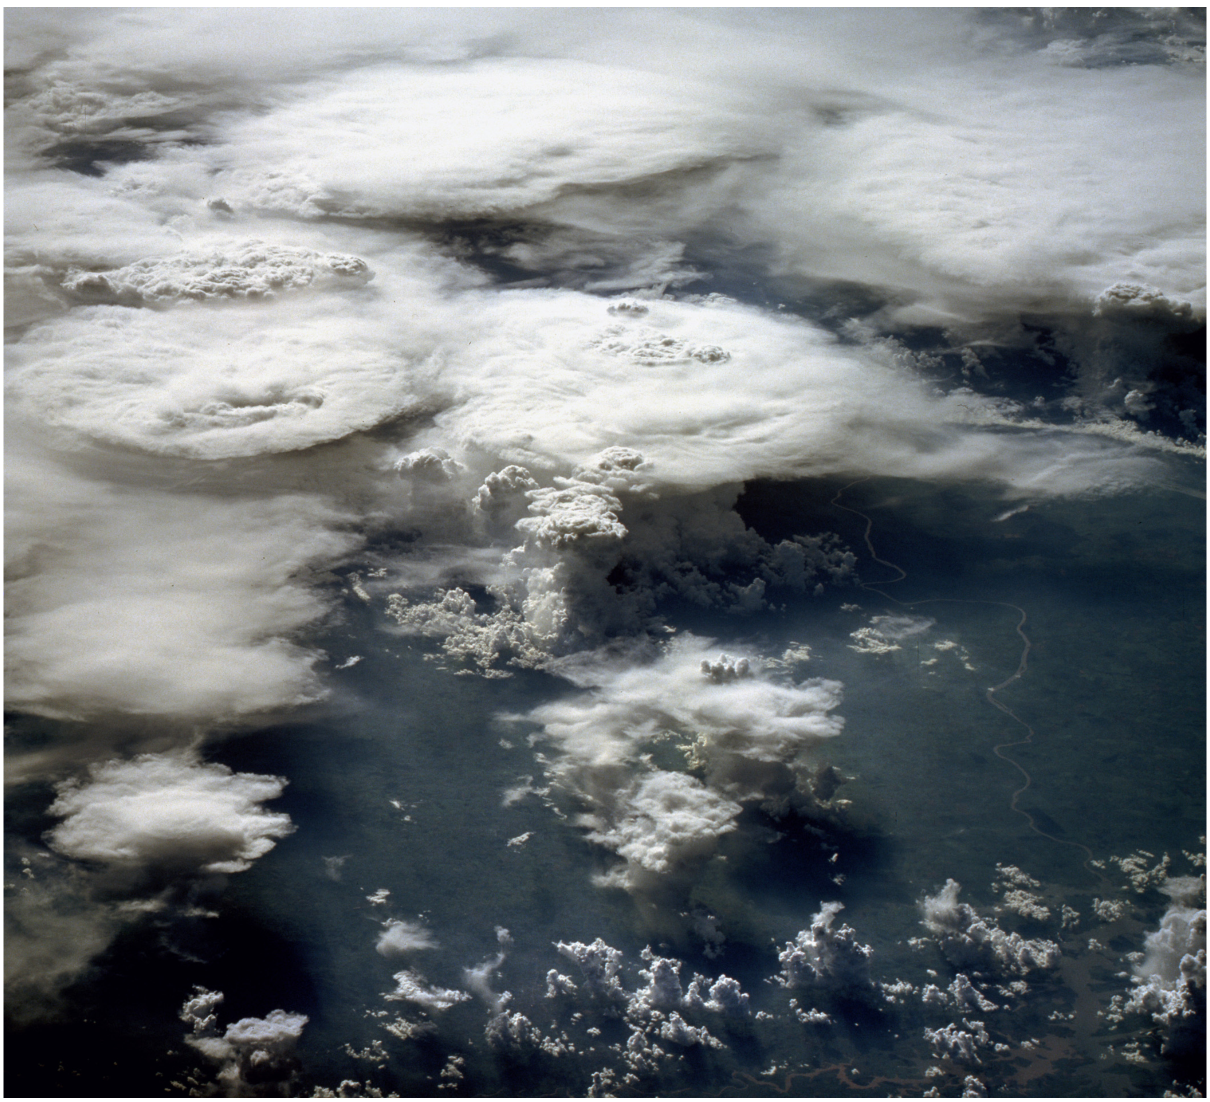

Figure 2.7. This space-shuttle view of cumulus clouds rapidly developing into cumulonimbus formations indicates the widespread presence of small-scale cloud phenomena, which are known to be important in affecting the larger-scale phenomena.

The availability of petascale computing resources to climate and weather researchers will enable increases in model spatial resolution that are necessary for resolving and studying the formation and interaction of cloud systems. By refining the horizontal mesh size down to 1 km, not only will the accuracy of currently captured processes be increased dramatically, but the simulation of new and important sets of processes will be enabled. The value of such a capability in a global cloud-resolving model has already been glimpsed by researchers at the 40-teraflops Earth Simulator in Japan (Miura et al., 2007).

Increasing resolution alone will not address the shortcomings of current parameterizations of convective processes; new techniques and sub-grid-scale representations of these multiscale convective processes must be developed. The key point is that these exploratory research activities can be dramatically accelerated and advanced through routine access to large-scale computing resources. The atmospheric modeling capability that will emerge will revolutionize atmospheric science by making it possible for researchers to numerically simulate the myriad of dynamical, microphysical, and radiative processes that make up the interactions of clouds with weather and climate systems (Randall et al., 2003). This will have immediate application to studies of weather prediction, climate change, water resources, atmospheric chemistry, and even ecosystem dynamics.

Explicit representation of clouds and convection in atmospheric models will likely accelerate progress in superparameterization and cloud-system resolving (or convection-permitting) models. In superparameterization, atmospheric convection is represented by cloud-system resolving models (grid spacing ~1 km), which mitigates several problems with traditional convective parameterization. For example, the mesoscale organization of convection is represented and the diurnal cycle of convection is improved substantially. Superparameterization has been tested in the NCAR Community Climate System Model (CCSM; Khairoutdinov et al., 2005). The global cloud-system-resolving model approach (Satoh et al., 2008) represents key elements of large-scale convective organization such as the Madden-Julian Oscillation. The explicit approach gives optimism that the longstanding problem of representing mesoscale organization in global models identified decades ago (Houze and Betts, 1981) can be solved.

Explicit approaches provide useful information on how moist convection acts as an integral part of the Earth System. Organization of convection in the mesoscale part of the motion spectrum (1 km–100 km) is vitally important since it directly affects the amount and distribution of precipitation on both local and global scales. Our comprehension of multiscale convective organization and how it interacts with the global circulation is moving forward (Moncrieff, 2004, Moncrieff et al., 2007), and has been given impetus in a new internationally coordinated project: Year of Tropical Convection (http://www.ucar.edu/yotc/). Without access to petascale computing, it will not be possible to move forward and remain competitive at the international level. In the latter respect, it is imperative that NCAR and the U.S. University Sector secure a petascale computing capability at the earliest opportunity.

Climate models are reasonably consistent regarding the warming due to carbon dioxide, but differ widely in their prediction of the hydrological cycle, e.g., some models suggest that average precipitation will increase sharply, while others indicate that it will not change much. A major challenge is prediction of the regional distribution and amount of precipitation. Quantifying this aspect of the climate system requires advanced knowledge of scale-interactions among convection, radiative transfer, the planetary boundary layer, and surface exchange processes across an enormous range of scales. Cloud-system-resolving models are the approach of choice because they couple the physical process at realistic scales.

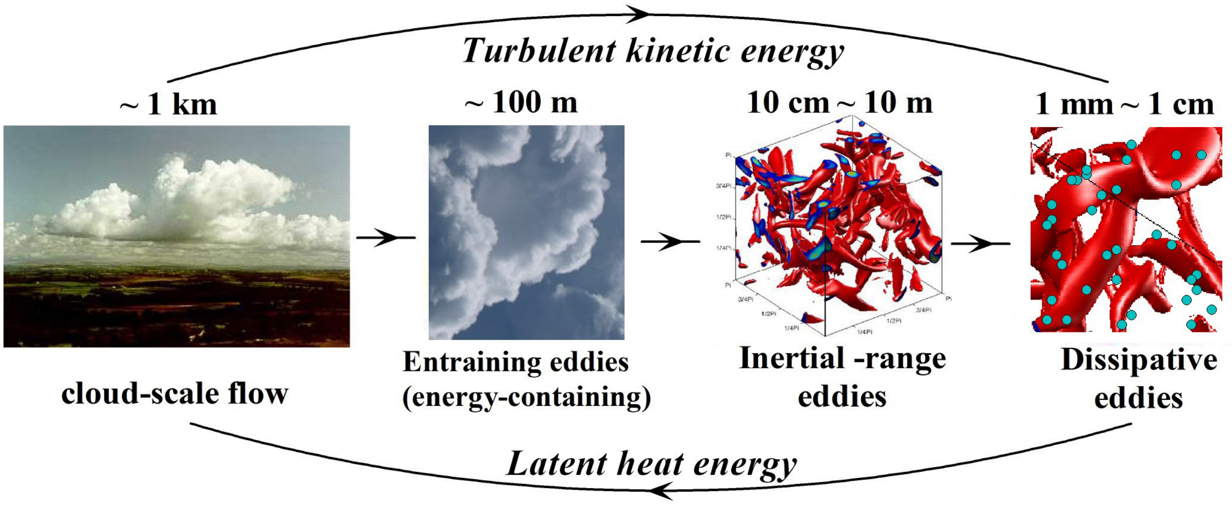

Figure 2.8. Multiscale interactions in atmospheric clouds. The turbulent kinetic energy flows from cloud-scale motion to dissipative eddies. Latent heat energy flows from individual droplets to cloud-scale motion.

Reliable weather and climate prediction at both local and global scales depends on our understanding of microphysical processes, small-scale dynamics, and interactions between clouds and the planetary boundary layer. The scale-dependent interactions are illustrated in Figure 2.8. While cloud systems may extend for hundreds of kilometers, the individual water droplets within them are typically 5–20 μm in radius. There are also several important intermediate scales. Cloud droplets are entrained by turbulent eddies of scale ~100 meters. Raindrops form through collisions arising from differential sedimentation and interaction with small-scale eddies about 1 millimeter to 1 centimeter in size. Multiphase characteristics are important. A droplet changes size through condensation and evaporation. The latent heating during phase changes at the droplet scale makes moist air buoyant. Turbulent kinetic energy moves downscale and latent heat energy moves upscale (Figure 2.8).

There may be up to 10^{17} droplets in a 1-km^3 cloud, so following all the droplets within a cloud is not computationally feasible. Alternative descriptions of the process are required, but this is a challenging task due to the inherent nonlinearities, inhomogeneities, intermittency, and disparate coupling scales. Cloud-resolving large-eddy simulation (LES) with bin-based microphysics modeling (Morrison and Grabowski, 2009) is being pursued. The droplets in a grid volume are represented by a number-density function and/or its moments. Such simulations typically have a spatial resolution of a few tens of meters; therefore droplet-scale microphysics has to be modeled.

Little is known about the necessary sub-grid-scale parameterization for moist turbulence, limiting the accuracy of cloud-resolving LES in modeling cloud dynamics. In recent years, direct numerical simulation (DNS) of small-scale air turbulence has been applied to study cloud microphysics (Andrejczuk et al., 2004, Wang et al., 2005). In DNS, turbulent air motion at dissipation-range scales (mm–cm) and a limited range of inertial-subrange scales (up to ~10 cm) are resolved, but larger-scale motion is represented by a forcing scheme. The current limitation of the use of DNS to low Reynolds numbers makes it impossible to directly address the effect of larger-scale turbulent eddies on droplet-scale processes, e.g., mixing, condensation, and coalescence.

Petascale computational resources would allow investigators to further develop and optimize LES of cloud dynamics and DNS of cloud microphysics to unprecedented resolutions, making it possible to perform DNS in a computational domain of O(1 m^3) and LES at grid spacing of O(1 m), well into the inertial range of cloud turbulence. Closing the DNS-LES scale gap will enable the development of sub-grid-scale models for inertial-range LES modeling offering a unified framework.

2.4.5 Regional Climate Modeling in Complex Terrain

Climate warming on the order of 2-4° C is predicted for mid-21st Century by the IPCC Fourth Assessment Report (Perry et al., 2007). Water availability is a more important climate change issue for states such as Wyoming, Colorado, New Mexico, and Utah. Predictions of changes in surface precipitation evaporation and their interannual variability are more uncertain (Hoerling and Eischeid, 2007). Climate trends and anomalies of the past two decades have already made large impacts on regional ecosystems and economies. The potential impact of these changes is amplified by the situation of the headwaters of several major North American river systems in the region; climate changes here will affect water availability downstream.

Scientists at the University of Wyoming are interested in examining how climate warming affects mountain ecosystems and their connected ecohydrological processes, as well as how biological and physical changes at the land surface influence precipitation processes and snowpack dynamics. Their research seeks to address the following integrated interdisciplinary questions:

- How do land-atmosphere exchanges of energy and water vary between contrasting ecosystems in a water-limited complex terrain environment?
- How are cloudiness and precipitation controlled by the land surface (including terrain, vegetation, and soil moisture)? How do seasonal-to-interannual ecosystem changes affect the energy, moisture, and aerosol fluxes that feed clouds and precipitation?
- How does ecosystem structure and function affect local-to-regional-scale climate variability – particularly snow accumulation and snowmelt dynamics – in water-limited complex terrain?

Meeting the scientific challenges of Earth System sciences requires unprecedented engagement among diverse disciplines, including the ecological, atmospheric, hydrological, and computational sciences (see Figure 2.9). These studies will be pursued via integrated and overlapping research-infrastructure initiatives (see Figure 2.10):

- The Forest-Steppe and Tundra Ecotones Research (FoSTER) Observatory will provide an integrated ground-based sensor network, monitoring array, and experimental area spanning the critical transitions between steppe and forest, and forest and tundra.

- The SOLPIN (Simulations and Observations of Land-Precipitation Interactions) initiative will examine how the terrestrial biosphere interacts with cloud and precipitation processes. SOLPIN will utilize FoSTER and other land-atmosphere data networks, but it will focus on novel airborne measurement techniques and related numerical experiments.
- The Research Computing Initiative (RCI), a mid-size supercomputer, will provide the computational infrastructure required to manage the data streams from FoSTER and SOLPIN, and develop and test computational models that complement and assimilate the observational data for these projects.

Section 2.4.2 on Nested Regional Climate Modeling describes the importance NCAR and other Earth System scientists see in using regional modeling approaches. The Wyoming EPSCoR work cited above demonstrates clearly how Wyoming and NCAR will collaborate in a major research question to which UW brings considerable strengths in studying ecological and biosphere processes and NCAR contributes expertise in atmosphere and climate modeling.

In particular, the NSF-EPSCoR RCI facility will be used for operational very-high-resolution (< 1 km) Weather Research and Forecasting (WRF) simulations, coupled with land surface and hydrological models, covering the state of Wyoming and the surrounding continental headwaters region. This work will improve weather, wind, and streamflow predictions, and it will improve our understanding of land surface-atmosphere interactions. The simulations will be conducted in collaboration with NCAR scientists and the WRF development team. The model output will be archived and served by the RCI facility to allow assessments on seasonal to interannual timescales.

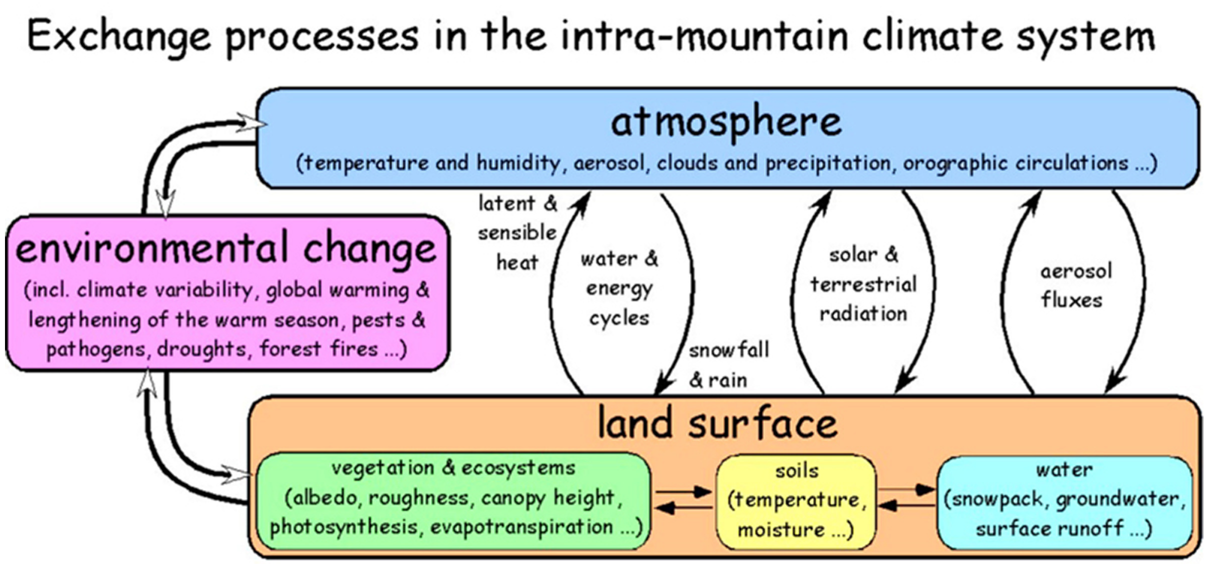

Figure 2.9. Key components and exchange processes in the climate system. Feedbacks between the land surface and the atmosphere, via energy and water cycles, radiation, and aerosol fluxes, lead to dynamic equilibrium between land surface and atmosphere. Environmental change affects the system, altering feedbacks to drive further, often unexpected, environmental changes.

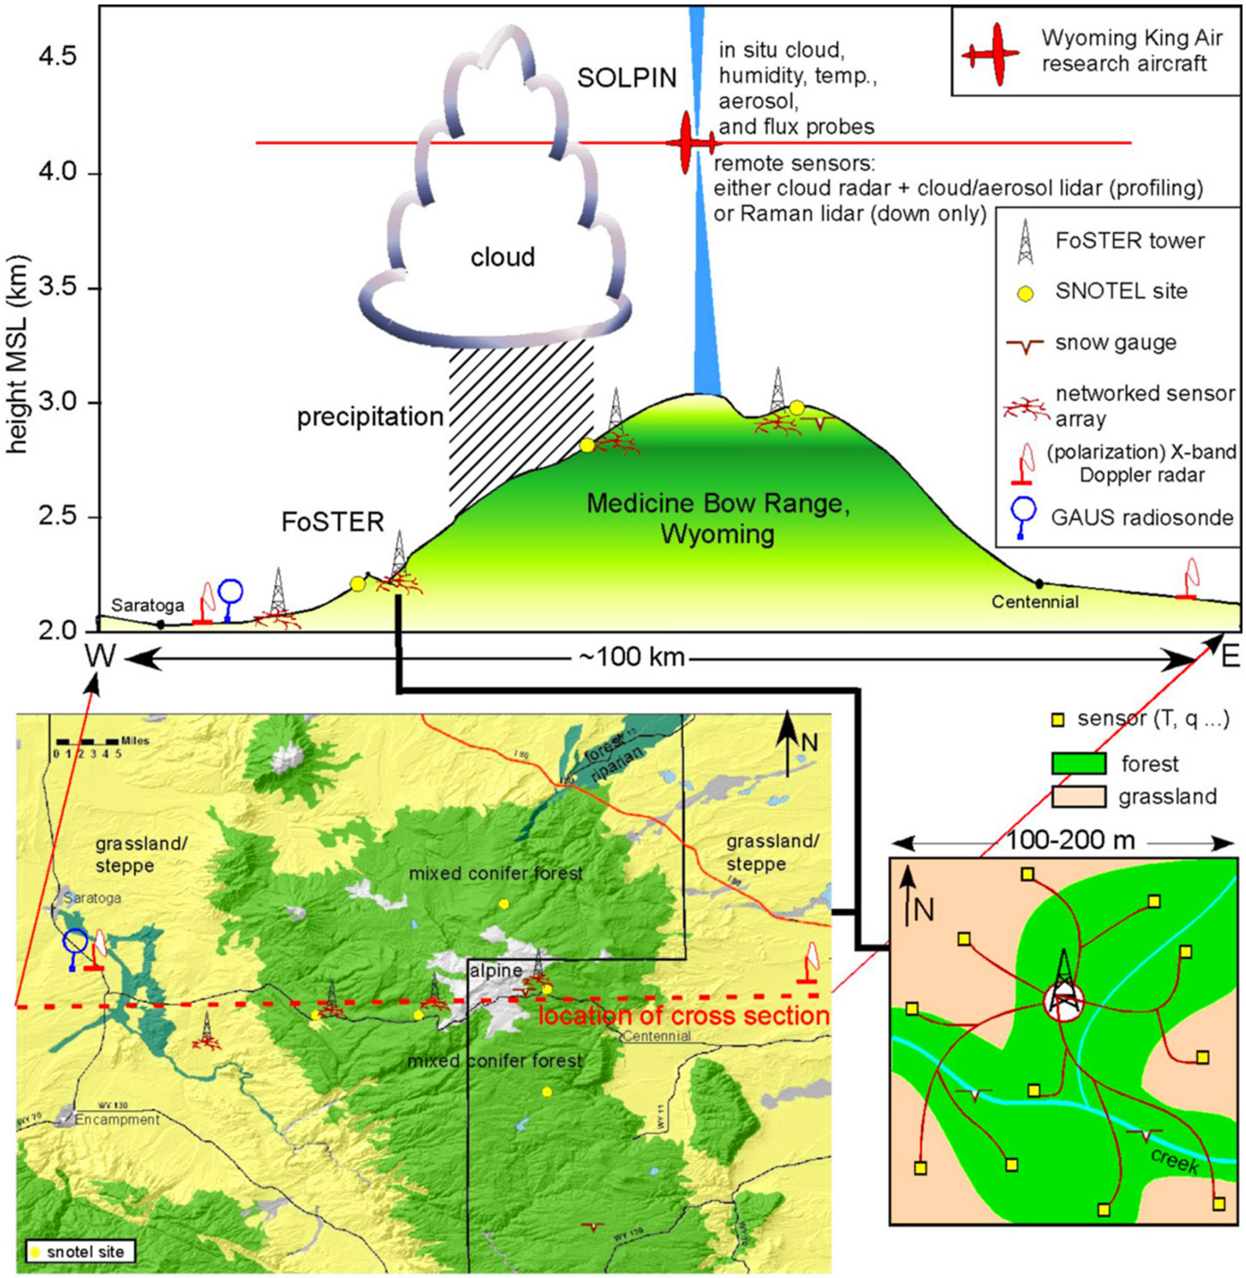

Figure 2.10. Schematic showing the relationships among the component projects of the northern Rocky Mountains regional climate observation and modeling experiment proposed by UW. The FoSTER Observatory will be emplaced along the western slope of the Medicine Bow Range, located 30-50 km west of the UW campus. The SOLPIN initiative will obtain airborne measurements in the Medicine Bow Range as well as other areas in the Rocky Mountain region.

During the transition period from the cold to the warm season, precipitation is least predictable yet most abundant in the region, resulting in much uncertainty in warm-season water availability. Spring also is the time that land surface-atmosphere interactions are believed to become most important. Therefore two field campaigns will be conducted in spring over the Medicine Bow Mountain Range (Figure 2.10). These campaigns will use new instruments enabled by FoSTER and SOLPIN. Their experiment will be designed in close collaboration with NCAR scientists, in particular through NCAR's Colorado Headwaters Research Program, to ensure data collection and structure are most relevant for model improvement.

Again the RCI facility will be central in the field campaign data collection, data serving, and model validation efforts. The RCI will provide an integrative cyberinfrastructure strategy by establishing a campus-level scientific computing entity bridging the gap between small individual clusters and the NSF's cyberinfrastructure assets, such as the TeraGrid and the proposed NWSC facility, following the model advocated by the NSF's own Cyberinfrastructure Council (Crawford et al., 2007). Establishment of such an entity will provide critical hardware for applications that have outgrown individual department capabilities or responsibilities, and will provide a testing ground for software that will eventually run on the TeraGrid and potentially the NWSC hardware. Dedicated staff will provide expertise, training, and experience on large-scale computational and data archiving systems.

The level of computational capability will be matched to the requirements of performing regional scale simulations using established tools such as WRF, which are currently run on large systems such as NCAR's Bluefire supercomputer, and which will eventually be run on the petascale system housed in the NWSC.

2.5 Space Weather

The term space weather refers to the evolution of disturbances that begin on the Sun, propagate through interplanetary space, enter the Earth's magnetosphere, and penetrate down into the ionosphere and thermosphere. These phenomena can have significant impacts on satellites, astronauts, and systems ranging from global positioning systems to power grids. Significant challenges in this field that will benefit from petascale computing include first-principles modeling of solar convection and its contribution to the 22-year solar cycle and, crucially, modeling the emergence of magnetic flux from the solar convection zone along with the conditions that lead to solar flares and coronal mass ejections (CMEs). These processes are only the beginning of the chain of events that lead to impacts in geospace; to bring space weather forecasting into practical use, these regions must all be modeled accurately.

Magnetic fields that give rise to eruptive events such as CMEs are generated in the solar convection zone and in the underlying layer of rotational shear known as the tachocline. Turbulent convection generates fields through magnetohydromagnetic dynamo action and pumps them downward into the tachocline where they are organized into toroidal structures and further amplified by shear. As they gain in strength, these toroidal fields become buoyantly unstable and rise to the surface, creating sunspots, bipolar active regions, and coronal loops. Axisymmetric circulations in the meridional plane transport fields and promote cyclic activity such as the 22-year solar activity cycle. Capturing these diverse processes in a numerical model is a formidable challenge that will greatly benefit from petascale computing resources. These numerical modeling techniques are already beginning to play a key role in making long-term predictions of space weather activity levels. Petascale computing resources can address issues such as the maintenance of rotational shear and meridional circulations by convection, the generation and organization of magnetic fields by convection and shear, the role of meridional circulations in establishing activity cycles, the destabilization and emergence of toroidal field structures, and the subtle coupling between the convection zone and the radiative interior via the tachocline. Importantly, petascale computing will also allow investigation of the coupling between deep convection and the smaller-scale convection that occurs near the solar surface (granulation and supergranulation), as well as the manifestation of interior dynamics and magnetic activity in photospheric observations and helioseismic inversions.

A comprehensive understanding of the process of magnetic flux eruption is the key to a quantitative understanding of solar activity and its consequences on the Earth System as well as to gain predictive capabilities at the source of space weather. Modeling the flux eruption process is one of the most difficult problems in solar physics, since it involves the entire convection zone with a density changing by six orders of magnitude. For these reasons, the phenomena of sunspots and flux eruption are the first topics in our discussion of space weather.

2.5.1 Sunspots and Magnetic Flux Eruption

Detailed modeling of the origin, dynamical evolution, and decay of sunspots is essential for our understanding of the solar magnetic cycle and the impact of solar magnetic activity on the Earth System through solar flares and coronal mass ejections that can buffet Earth's atmosphere. The resulting damage to power grids, satellites, and other sensitive technological systems takes an economic toll on a rising number of industries. The hierarchy of strong magnetic field structures at the surface ranges from large sunspots (up to 10,000 km) to bright faculae (less than 70 km). Radiative energy loss to outer space associated with them gives rise to variations of the solar heating of the Earth, leading to variations of the Earth's climate that are yet to be understood.

Sunspot origin and structure. Sunspots are the most prominent manifestations of magnetic field in the visible layers of the solar atmosphere. Their origin is a dynamo process operating in the solar convection zone. Magnetic fields generated there over solar cycle timescales (~11 years) are transported toward the solar surface through a rapid flux emergence process leading to the formation of active regions, or sunspots. Sunspots are typically somewhat larger than the Earth's diameter, their field strength is ~3,000 Gauss, about 10,000 times stronger than the Earth's magnetic field. Such strong fields modify convective energy transport leading to a central dark region, called the umbra, with a predominantly vertical field and a brightness reduced to about 10-20% of the undisturbed solar surface. Even though the umbra appears dark, the temperature is still 4,000–4,500° K. The umbra is surrounded by a filamentary region, the penumbra, with a brightness of about 75% (see Figure 2.11). The penumbra region has a strongly inclined field and exhibits outflows of several km/s, called the Evershed flow (Evershed, 1909). The umbra of a sunspot shows fine structure on spatial scales similar to those of penumbral filaments. Also, it is not uniformly dark, but shows bright "umbral dots" with sizes up to a few 100 km. These umbral dots, (see Figure 2.11), similar to penumbral filaments, show a central dark lane (Scharmer et al., 2002).

Radiative MHD simulations of sunspot structure. Performing computer simulations of entire sunspots is challenging. Sunspots have a typical size of several 10,000 km and show structure down to the smallest observable scales on the order of a few 100 km, requiring large computing domains to simulate, which has only become feasible recently. First attempts by Schüssler and Vögler (2006) focused on a detailed study of the umbra of a sunspot. Simulations including the transition from umbra toward granulation are more demanding, and were initially performed in 'slab' geometry, i.e., in a narrow rectangular domain cut through the center of a sunspot (Heinemann et al., 2007, Rempel et al., 2009a). Recent advances in supercomputing capabilities allowed scientists at NCAR's High Altitude Observatory (HAO) and the Max Planck Institute for Solar System Research (MPS) to perform the first comprehensive simulation of a pair of sunspots. The simulation domain encompasses a horizontal domain of 98,000 × 49,000 km with a depth of 6,100 km, with a horizontal-vertical grid spacing of 32 x 16 km. The computation including 1.8 billion grid points was performed on NCAR's new IBM Bluefire supercomputer, requiring about 250,000 CPU hours. This allowed scientists to simulate about 2 hours of temporal evolution, which is sufficient to understand the magneto-convective origin of sunspot fine structure (see Figure 2.12).

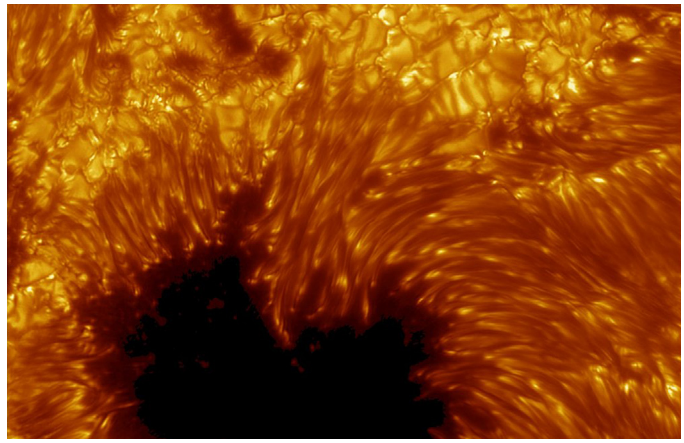

Figure 2.11. Sunspot observed with the Swedish Solar Telescope (SST). This image shows the transition from the dark umbra (bottom) toward solar granulation (top). The penumbra in between shows filaments with central dark lanes. Credit: The SST is operated by the Royal Swedish Academy of Sciences; observation taken 2002 by G. Scharmer.

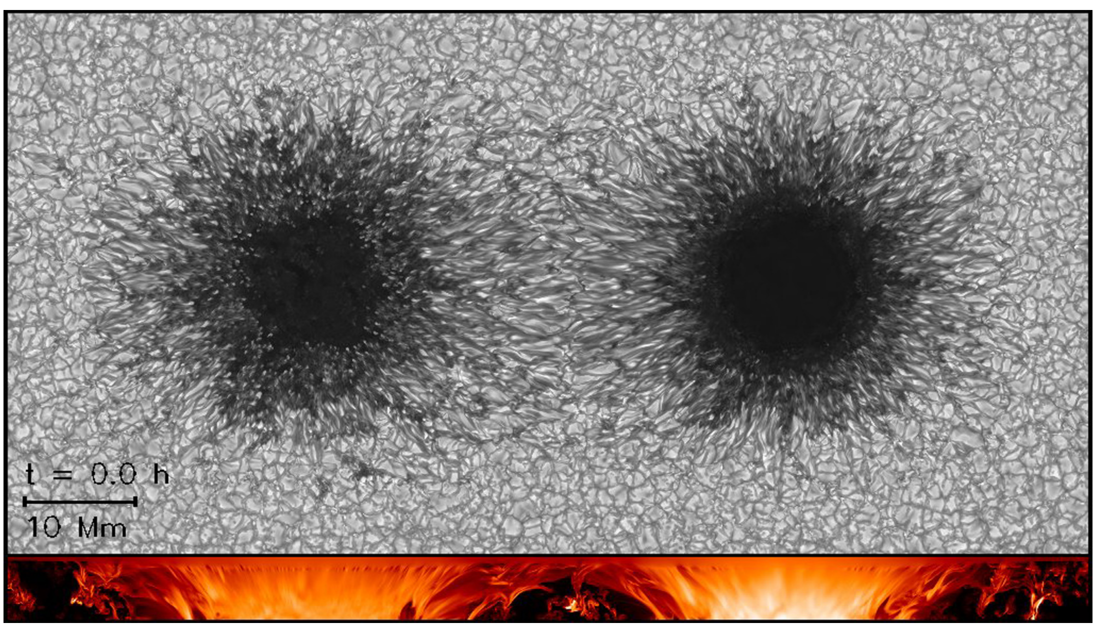

Figure 2.12. In black and white: Bolometric intensity. In color: Subsurface magnetic field strength on a vertical cut through the center of the spot pair, values range from 0 G (black) to 8 kG (white). The simulation encompasses the dark umbra of sunspots with bright umbral dots, the filamentary penumbra and almost undisturbed solar granulation near the horizontal boundaries of the computational domain.

The computer model for these simulations is based on the MURaM solar MHD code, which has been modified to cope with the challenges of simulating strong magnetic field in the solar photosphere. The model solves the MHD equations together with 3-D radiative transfer (rays in 24 directions) and uses the OPAL equation of state to properly represent the partial ionization of the plasma in the upper solar convection zone.

Altogether, numerical simulations indicate a new level of realism in the theoretical modeling of sunspot structure. The basic properties of the simulated umbral dots and penumbral filaments are consistent with a variety of observational results and provide a basis for a physical understanding of umbral and penumbral structure in terms of common magneto-convective processes that are modulated by the varying inclination angle of the magnetic background field. Figure 2.13 displays the connection between magnetic field structure (vertical field strength, top left, and inclination angle, top right) and flows in the photosphere (radial flow, bottom left, and vertical flow, bottom right). Convective motions are present everywhere in the penumbra, and they are the primary driver for the fine structure we observe. Upflows

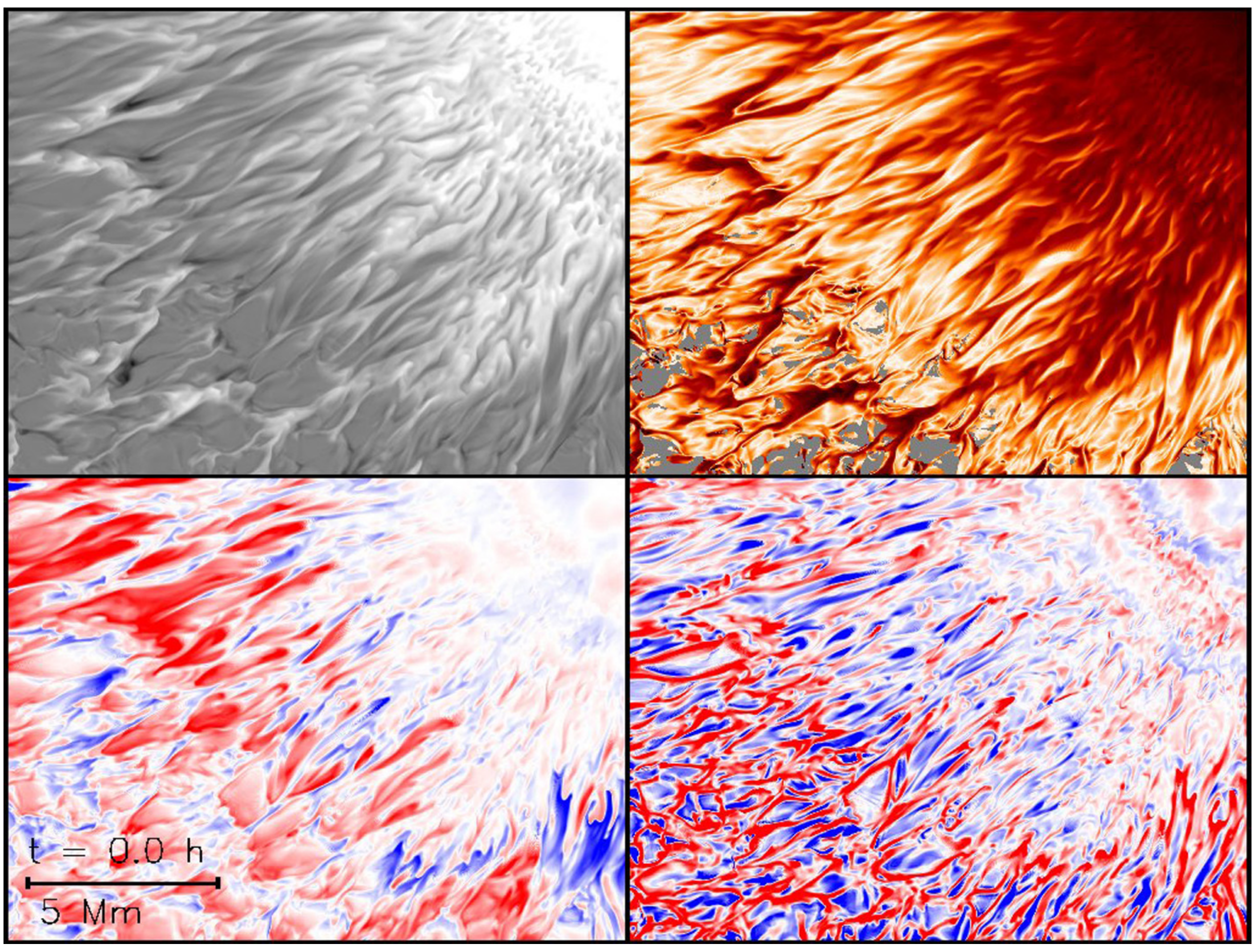

Figure 2.13. Transition from umbra (at right) to penumbra (at left). Quantities shown are in the photosphere. Top left: Vertical magnetic field component (white: upward, black: downward). Top right: Field inclination (darker: vertical, lighter: horizontal). Bottom left: Radial flow velocity (red: outflows, blue: inflows). Bottom right: Vertical flow velocity (red: downflows, blue: upflows).

expand and weaken the magnetic field to a degree that enables overturning motions to take place. The convection affects the vertical field component stronger than the horizontal one, leading to filaments with increased field inclination angle. Horizontal flows are preferentially directed outward. In the inner penumbra, outflows are intermittent with velocities below 1 km/s. The outflows in the outer penumbra largely fill the space and reach several km/s.

Future directions in sunspot simulation. Significant new discoveries can be achieved by expanding the simulation length, size of the computational domain, as well as increasing resolution. The lifespan of sunspots can range from a few hours to months in extreme cases, with a typical lifespan of about a week. The horizontal extent requires an increase of at least a factor of two. The vertical extent of 6,100 km is still rather shallow given the fact that sunspot groups form as part of a flux emergence process that transports magnetic flux from a depth of about 200,000 km. Furthermore, a proper study of the flux emergence process will also require an extension of the computational domain upward to cover the connection to the chromosphere as well as corona. The modeling of flux emergence into the solar atmosphere is currently being carried out with ideal MHD simulations (see "MHD simulations of magnetic flux emergence into the solar corona" below). The latter is relevant for a physical understanding of the triggering mechanisms of flares and coronal mass ejections.

The current resolution of 32 x 16 km does not yet allow for resolving turbulent flows within the elements of sunspot fine structure (umbral dots, penumbral filaments) that have typical diameters of only a few 100 km. A substantial increase in resolution is required to allow for proper modeling of the turbulent properties of these flows as well as detail comparisons with high-resolution observations becoming available in the near future from large-aperture solar telescopes such as the Advanced Technology Solar Telescope (ATST).

Thus, future numerical experiments will consist of simulations covering longer time spans, larger domains, as well as higher resolution:

- An extension of the current experiment to about 8 hours is desirable to study how the sunspot fine structure changes with further thermal relaxation. Output of 3-D data at a higher cadence is required to study in detail the evolution of magnetic field structure over the lifetime of penumbral filaments. This extension will require about 500,000 CPU hours on the Bluefire system at NCAR and will produce about 50 TB of data. The 3-D analysis will be performed with the VAPoR visualization and analysis tool.

- Simulating the formation of an active region remains a demanding numerical experiment. As a compromise between feasibility and the minimum requirements, a simulation in a 200,000 × 100,000 × 24,000 km domain at a resolution of 48 × 24 km covering about 3 days of time evolution will be the target for the next five years. This experiment will require about 1,000,000 iterations in an 8-billion-gridpoint domain for which we estimate about 20 million CPU hours on a computer such as Bluefire. At a moderate output cadence, data on the order of a few 100 TB will be produced. It is likely that a couple simulation runs with different initial states will be required – some of those will be exploratory runs at lower (96 x 32 km) resolution about an order of magnitude less expensive.

- Simulations of sunspot fine structure at higher resolution will require a substantial code development effort to take advantage of adaptive or nested grids. The latter is required to be able to capture the large-scale sunspot structure and details at the same time – simulations of penumbrae in small domains suffer heavily from the influence of boundary conditions. Due to the required code development efforts, an estimate of computing requirements is difficult at this point. Some exploratory studies can be performed on smaller domains with uniform grids (e.g., focus on umbral dots). Simulating the evolution of umbral dots in a 6,000 × 6,000 × 2,000 km domain at 4 km grid resolution would require about 500,000 CPU hours on Bluefire.

- It may be that in the course of pursuing simulations described in the second experiment above, the need for even larger computations will become evident. In particular the overall modification of the solar radiative output associated with the flux eruption of the largest active regions might turn out to be key to understanding solar variability in the context of climate change. Such calculations easily could scale in factors of 100 and more as compared to the case describe in the second experiment.

MHD simulations of magnetic flux emergence into the solar corona. HAO scientists have also been carrying out idealized MHD simulations of the emergence of twisted magnetic fields from the top layer of the solar convection zone into the stably stratified, rarefied solar atmosphere and corona, and of the development of precursor structures for solar eruptions such as flares and CMEs (e.g., Fan, 2009). A 3-D MHD code is used that solves the ideal MHD equations in either a Cartesian or a spherical domain. These simulations have provided insights into how active-region magnetic flux reaching the photosphere can expand dynamically into the atmosphere as a result of the non-linear growth of the magnetic buoyancy instability, and how a flux rope structure containing sigmoid-shaped core field can form in the corona as a result of the shear and rotational motions on the photosphere driven by the Lorentz force of the twisted emerging fields (see Figure 2.14).

These simulations suggest that the observed sunspot rotations in solar active regions is a result of propagation of nonlinear torsional Alfvén waves along the emerging flux tube, transporting significant twist from the tube's interior portion toward its expanded coronal portion. This is a basic process whereby twisted flux ropes are developed in the corona with increasing twist and magnetic energy, leading up to solar eruptions.

Currently, ideal MHD simulations of flux emergence into the solar atmosphere are carried out in a domain with a vertical extent from about a few million meters (Mm) below the photosphere to about 20 Mm above the photosphere, and with a horizontal extent of 20 Mm by 20 Mm, which are small compared to the spatial scale of realistic active regions. For these simulations, a non-uniform grid is used that has the highest resolution (30 km) near the photosphere level and for the central region where the magnetic flux tube is emerging. The grid is stretched toward the sides away from the emerging magnetic structure, and also high up in the corona. The grid size for the current simulations is typically 320 × 300 × 392, and each simulation uses about 50,000 CPU-hours on Bluefire. At minimum, a factor of five increase in the computational domain size (in all three directions) is needed to model flux emergence into the corona on realistic scales for active regions. The simulation domain size will have a vertical extent from about 20 Mm below the photosphere to about 100 Mm above the photosphere and a horizontal extent of 100 Mm × 100 Mm. This would require a non-uniform grid of size of about

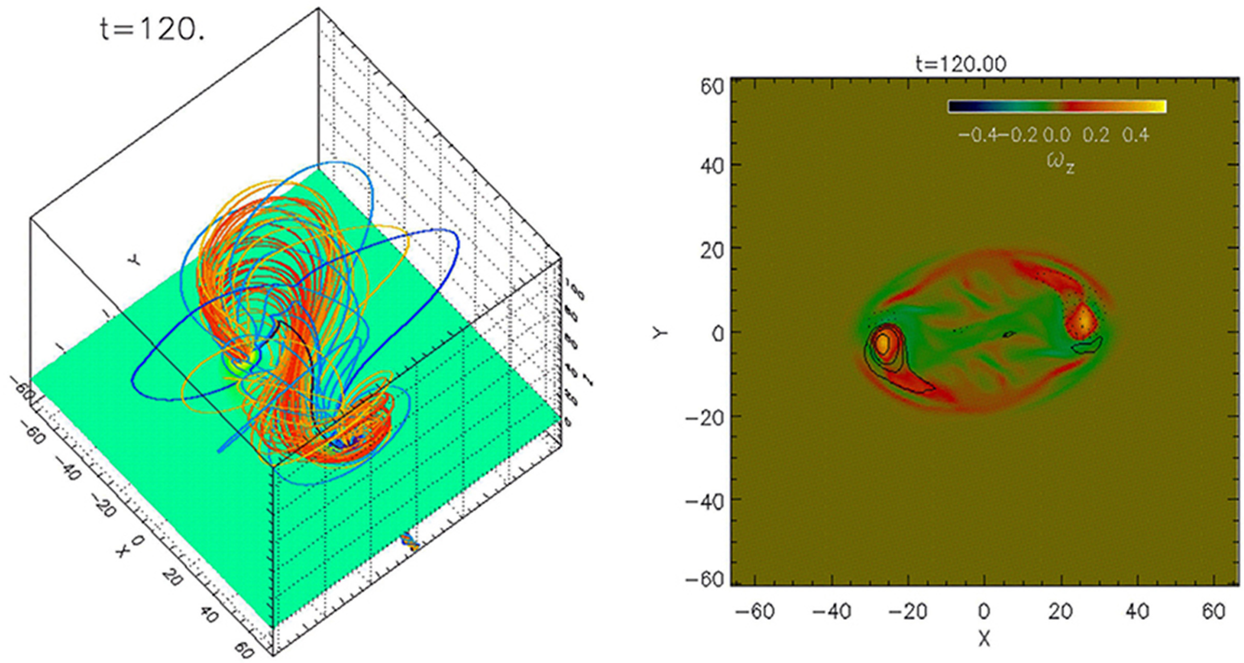

Figure 2.14. The left panel shows the 3-D coronal magnetic field resulting from the emergence of a twisted subsurface flux tube. A coronal flux rope structure with sigmoid-shaped core fields is formed. The right image shows the z-component of the vorticity on the photosphere overlaid with contours of Bz with solid (dotted) contours representing positive (negative) Bz. The image shows counterclockwise vortical motion (i.e., positive z-vorticity) centered on the peaks of the vertical flux concentrations of the two polarities, reminiscent of the observed sunspot rotations in many flare-producing active regions.

1024^3, with the finest resolution remaining at about 30 km. The physical time of flux emergence to be simulated would be about 10 hours. Such a simulation would require 20 million CPU hours on Bluefire.

Coupled models. Currently, many aspects of the flux emergence process are modeled individually (as described in “Radiative MHD simulations” and “MHD Simulations of magnetic flux emergence” above), primarily due to the substantially different physical regimes encountered in the solar convection zone and the solar corona above. With larger computing resources available in the future, it will become feasible to either couple individual models or even combine them into a single model. The latter is an option for the models described in “Radiative MHD simulations” and “MHD simulations of magnetic flux emergence” above since both are based on fully compressible MHD. However, the very different dynamical time scales encountered in the convection zone and corona might make a coupled run (with individual time stepping) more economical.

2.5.2 Earth-Sun Interactions

Modeling the impacts on the Earth of these eruptive events on the Sun also involves significant computational resources. First, the evolution of these structures in the heliosphere as they propagate to geospace must be modeled. Petascale computing combined with AMR techniques will provide the resources needed to refine the grid resolutions to the degree needed to accurately track the interactions between the solar wind plasma and the plasma contained in the CME. Once these structures reach geospace, global-scale magnetohydrodynamic simulations of the magnetosphere must model a region that extends 20 million meters in the Sunward direction and 200 million meters in the anti-Sunward direction with a resolution in the inner magnetosphere on the order of 100 kilometers to accurately represent the reconnection processes that transfer mass, momentum, and energy from the solar wind into the system. Obtaining these resolutions is currently at the edge of the supercomputing resources available today. Furthermore, the ionosphere is tightly coupled to the magnetosphere, requiring the simultaneous execution of an upper atmosphere model that includes modeling numerous chemical interactions occurring throughout this region. The electric and magnetic fields also drive the energization and transport of radiation belt electrons that can be modeled with guiding center particle codes. As petascale computational resources become available, this can be done with large particle simulations. Furthermore, this system is highly driven by the variable solar wind making it an excellent candidate for ensemble forecasting methods like branch prediction.

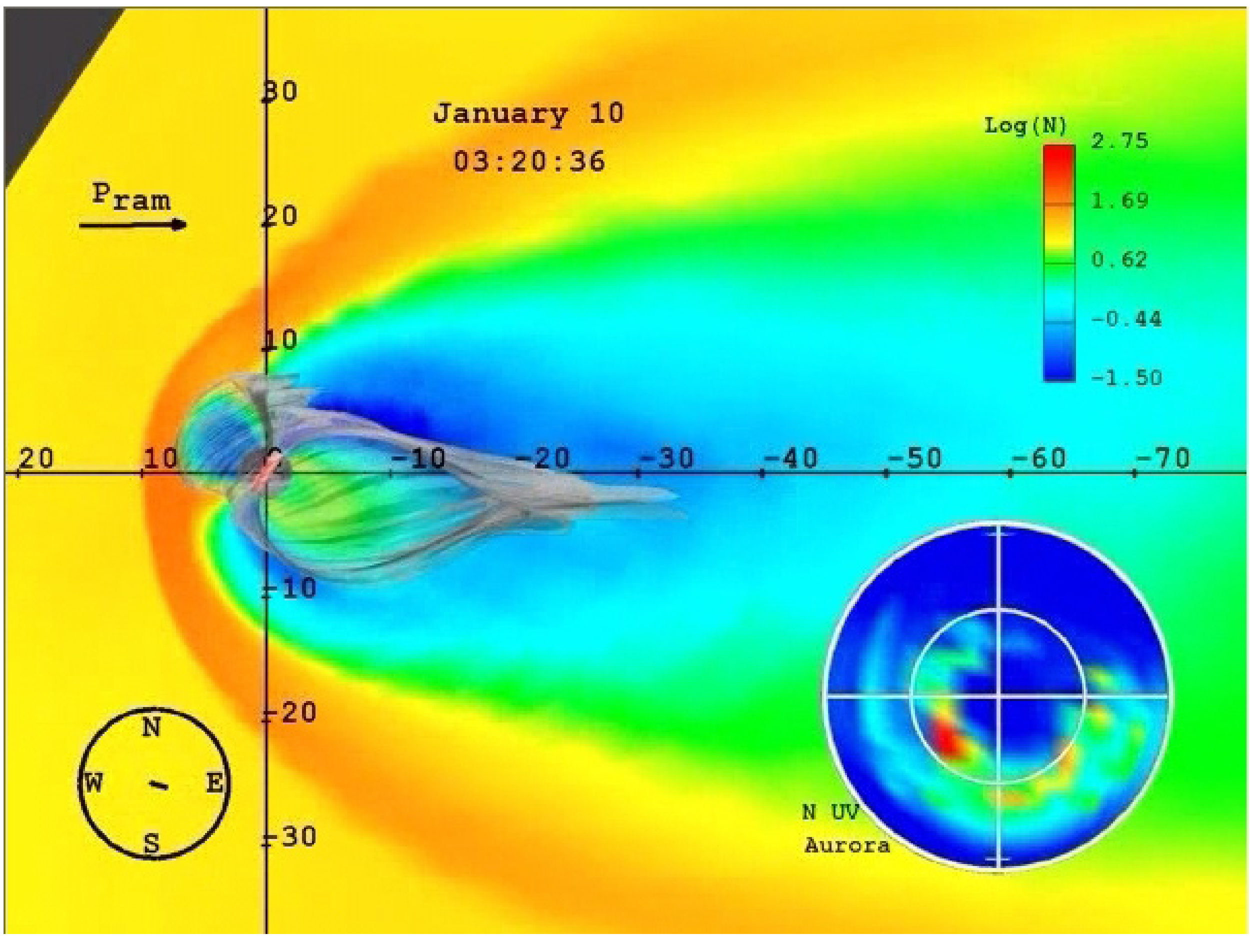

Figure 2.15. Scientific visualization of the impact of coronal mass ejection on the Earth's magnetosphere.

2.6 Oceanography

A comprehensive and predictive understanding of the Earth System must account for the interacting physical and biogeochemical processes governing the storage and transport of energy, freshwater, nutrients, and dissolved gases within the ocean and the coupling of these cycles to the atmosphere, cryosphere, and terrestrial environments. As for any turbulent fluid, these processes span an enormous range of space and time scales, far larger than can be explicitly represented in any numerical model. In addition to the range of scales that arise through intrinsic variability of the fluids system, the complex topography of the boundaries ocean basins impose a similarly broad range of scales. Another feature that distinguishes oceanic flows from those in the atmosphere or freshwater systems is the possibility of density compensation from the dual dependence of density on temperature and salinity. Some ocean temperature or salinity fronts have only a weak signature in density: mixing within their frontal zones is dominated by the unique processes of interleaving and double diffusion, rather than instabilities associated with shear.

Petascale computational resources will lead to more accurate, more precise, and more complete ocean models, and thus increase our confidence in predictions of the Earth System. Accuracy will be improved through replacing ad hoc parameterizations of sub-grid-scale processes by more realistic treatments based on insights gained through very-high-resolution computational process studies, in combination with field observations, laboratory experiments, and theory. Precision will be improved by increased spatial resolution of global- and regional-scale models, allowing more detailed, site-specific predictions. Coupling of additional important linkages among physical, biological, and chemical processes, as well as human interactions, that heretofore had been prescribed or neglected, will lead to more comprehensive predictions. The following three examples provide more specific descriptions of the opportunities for advancing ocean and Earth System science through petascale-class computing. These are by no means comprehensive. Other areas of oceanography that will benefit from enhanced computing include data assimilation for state estimation and initialization of ocean and coupled models used for predictions from days to decades; paleoclimate studies requiring integrations of many millennia to investigate events like ice age initiation and termination; expanded ensemble sizes to provide more robust statistical quantification of forced and natural climate variability; and the ability to investigate a broader range of scenarios in climate change impact and assessment research.

2.6.1 Global Ocean Circulation

The atmospheric components of Earth System models simulate climate by explicitly predicting, then averaging over, the day-to-day weather for periods of years to centuries. The weather systems, arising spontaneously as instabilities of the flow established by the equator to pole gradients in solar heating, have spatial scales of O(1,000 km) and time scales of a few days to a week. The weather systems are more than simply noise on the climate state: their transports of energy, water, and angular momentum are fundamental in establishing the character of the atmospheric circulation and Earth's climate (Peixoto and Oort, 1992). The ocean has weather systems, i.e., mesoscale eddies, dynamically analogous to those in the atmosphere. The majority of the kinetic energy in the ocean is contained in mesoscale eddies and, as in the atmosphere, they are a central part of the dynamical balances of major current systems, and in the transport of energy and material through the ocean. Yet, out of 24 coupled climate models presented in the recent Fourth Assessment Report of the IPCC (Solomon et al., 2007), only one (K-1MD, 2004) attempts to explicitly simulate them. Instead, the flow in nearly every contemporary ocean climate model is represented as essentially laminar, and the effects of ocean mesoscale eddies are entirely parameterized.

Why does this fundamental difference in the approach to modeling atmospheric and oceanic climate dynamics exist? The answer is primarily one of computational economics. Ocean mesoscale eddies have characteristic length scales of O(100 km) or less. With a factor of 10 additional grid points required in each horizontal direction, and a corresponding decrease in the time step, approximately 1,000 times more computational degrees of freedom are required for an eddy-resolving ocean model compared to an atmosphere model to simulate the global domain for same period. This computational cost has been impractically high for climate-length

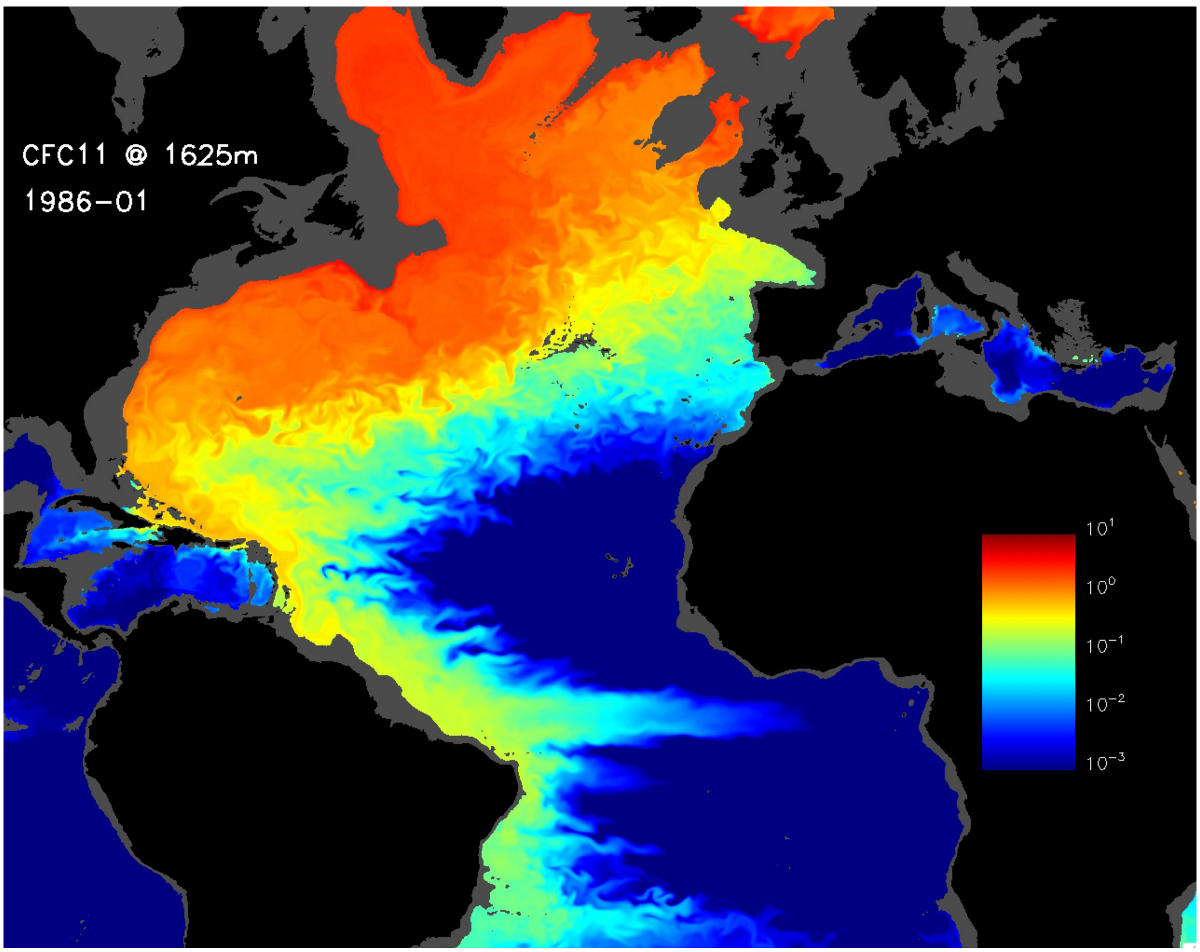

Figure 2.16. The image shows the simulated concentration of Freon-11 (picomoles/kg) in the deep northwest Atlantic from a global eddy-resolving ocean model. Gases such as carbon dioxide and freons enter the deep ocean in a few isolated high-latitude regions of deep convection such as the Labrador Sea, then spread through the ocean interior. The rate of oceanic uptake and spreading of these materials is an important factor in determining the anthropogenic forcing of climate change. Recent observations (Bower et al., 2009) have challenged the conventional picture of this spreading occurring through coherent boundary currents. A very high-resolution simulation of ocean tracers is able to explicitly represent this much more turbulent spreading process.

integrations. With the impending availability of petascale-class supercomputers, we will witness a transition in this state of affairs. Fully coupled climate simulations with ocean component models having resolution of 10 km or better will define the state of the art.

Improvements in the simulated circulation that have been demonstrated (McClean et al., 2006) – when ocean model resolution is increased from O(100 km) typical of current climate models to O(10 km) – have resulted not only from the more accurate representation of mesoscale eddies, but also from improvements in the representation of features of the mean flow (such as western boundary currents and their extensions) that have similarly small scales. One of the most notable systematic biases in coarse-resolution ocean models has been a poor simulation of the path of western boundary currents, especially the position of their separation from the coast. While eastern boundary regions are less energetic, the important processes setting the basic structure of the flow occur on similarly small scales. A dramatic improvement in the fidelity of boundary currents is seen when resolution is increased to around 10 km.

However, there is no indication that the "eddy resolving" solutions obtained to date have converged even for quasi-geostrophic dynamics (Siegel et al., 2001), and the current generation of models remain quite sensitive to any number of choices in model configuration and sub-grid-scale closure. That is, we are still far from "Large-eddy simulation" for global ocean models. It is inevitable that models with increasingly finer resolution will be constructed, though potentially more resource-efficient methods such as nested grids, or adaptive grids may see increased application. Moving into the eddying regime presents new challenges and opportunities in developing parameterizations of processes occurring at yet smaller scales. The expanded range of explicitly resolved scales opens the opportunity for putting parameterizations of processes such as inertial wave excitation and propagation, gravity currents, and mixed-layer restratification on a firmer physical foundation.

2.6.2 Coastal and Estuarine Oceanography

The coastal oceans and their adjacent estuaries affect, and are in turn affected by, human activity of many types. For example, coastal processes influence shipping and navigation, search and rescue, biological productivity, and recreation. In turn, they are subject to the influences of fishing, input of nutrients and pollutants, and coastline and bathymetric modification (e.g., by dredging). Finally, all of these processes and activities are subject to the influences of climate variability and change, in the form of global sea level rise and warming, enhanced hydrologic and atmospheric forcing, and the shifting of marine biogeographic boundaries.

The coastal oceans are also challenging to understand, being both complicated and complex not only physically (Haidvogel et al., 2000), but also chemically and biologically. Figure 2.17 gives one regional example of this interdisciplinary challenge. Despite this complexity, however, end-to-end models of the coastal environment are beginning to demonstrate considerable skill (Haidvogel et al., 2008).

An important goal for the future will be to simulate regional oceanic ecosystems dynamics within the context of a unified Earth system climate model. The primary elements of the modeling system might for instance include a nested hierarchy of physical circulation models for the ocean and the atmosphere; one or more food web models embedded within, and evolving in response to, the physical environment; one or more individual-based models for the relevant higher-trophic-level species; and appropriate mechanisms (possibly utilizing advanced

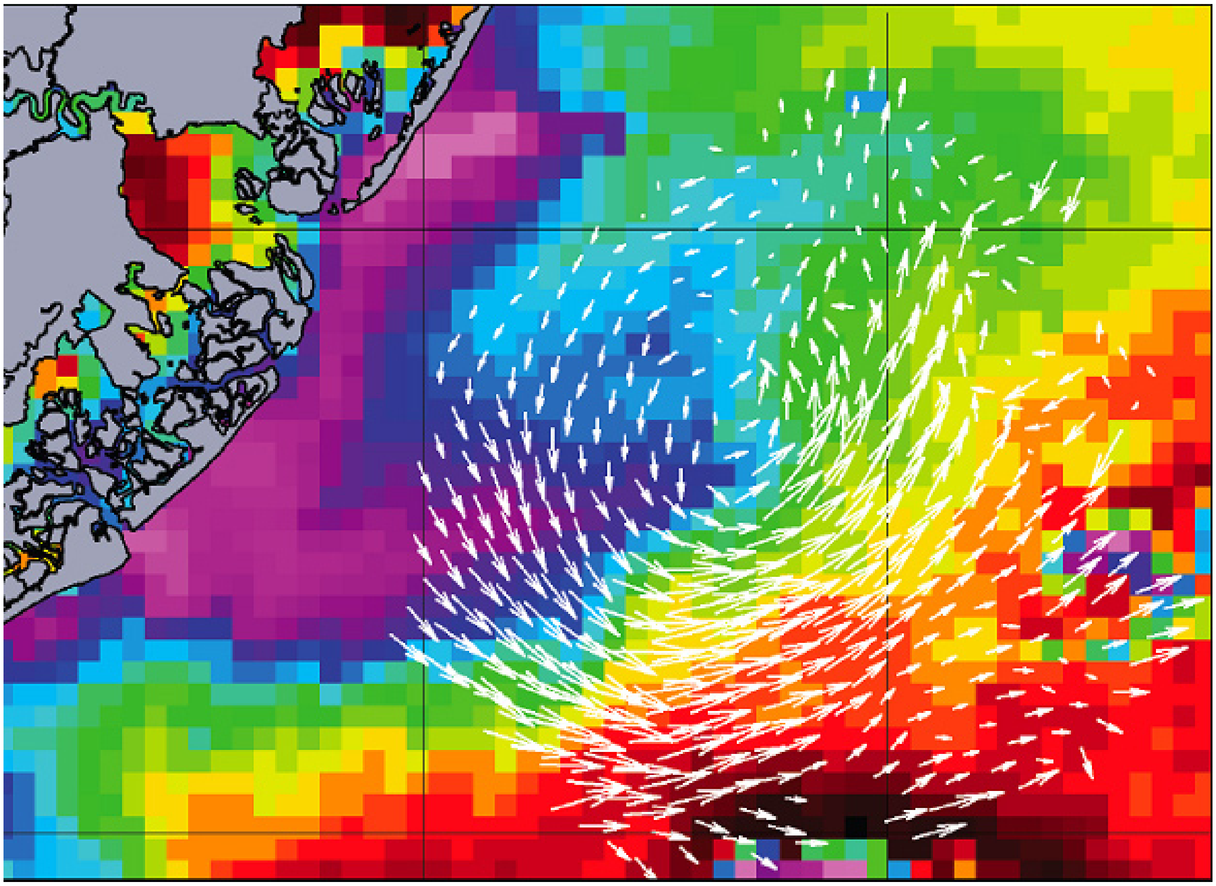

Figure 2.17. Hypoxia conditions off the Southern New Jersey Coast are set up by a combination of topography and upwelling factors. The purple coloration represents cold water from depth, indicative of upwelling. The upwelling brings nutrients to the surface, the enhanced nutrients lead to blooms, the resulting biological material sinks and depletes the deep oxygen, and hypoxia/anoxia results. This process illustrates the intimate connection between biological, chemical, and physical processes in the ocean. (Image courtesy of Dale Haidvogel, Institute of Marine and Coastal Sciences.)

data assimilation) for comparison and/or fusion of these forward models with the available retrospective and contemporary datasets. The challenge of developing and deploying such an integrated system is keen. However, many of the individual pieces are already in place (e.g., within the U.S. Global Ocean Ecosystems Dynamics – GLOBEC – program) (Mantua et al., 2002).

Petascale computing offers the prospect of new and important approaches to studying the coastal environment within the context of unified Earth System Modeling. In particular, a dedicated petascale computer will provide important resources for making progress in multidisciplinary, multiscale Earth System science, such as that emphasized in the U.S. GLOBEC program. This requires the coupling of the principal GLOBEC modeling tools within the framework of a global climate model such as the CCSM. Such coupling efforts are now well advanced, and proof-of-concept demonstrations have been performed showing that regional embedding of coastal regions produces sizeable global changes to the Earth System[15].

15 For an example of global changes produced from local refinement, see Section A.1.2: Nested Regional Climate Modeling.

A second area where future progress will depend on petascale computing and associated technologies is coastal ocean modeling and prediction. The coastal oceanographic community, through the U.S. Integrated Ocean Observing System (IOOS), has put forward a vision of an ocean modeling and prediction system covering the U.S. continental shelves. The envisioned system will use advanced modeling systems and observational networks to meet the needs of environmental assessment, forecasting, management, navigation, and national security. Eleven regional associations within IOOS have been formed to pursue regional themes related to maritime safety, ecological decision support, storm surge response, and other issues. The availability of petascale resources to the coastal modeling community will enable the exploration of pan-regional crosscutting themes (e.g., observational array design), the synthesis of regional observational and model datasets (e.g., the production of reanalysis products), as well as the development and application of advanced forms of model/data fusion and assimilation.

2.6.3 Ocean Process Studies: Saltwater Turbulence

In the summer of 2007, the annual retreat of the Arctic ice cap was the largest in recorded history. Meanwhile, communities across the Arctic suffered flooded basements and other structural problems as permafrost, no longer “perma,” melted. This dramatic warming is due largely to an increased flow of warm water northward from the Atlantic to the Arctic Ocean.

Fortunately, not all of the warm Atlantic water lies adjacent to the surface where it can melt ice and affect climate; much of the warm water is trapped in layers far below the surface. This is possible because the Atlantic water is relatively salty, and can therefore be denser than the surrounding Arctic water despite its higher temperature. Future climate change depends in part on the fate of these submerged, warm layers, which is in turn determined by mixing at the layer boundaries.

As it stands, oceanographers are in a poor position to predict the mixing of salt in seawater, as it is controlled by subtle aspects of seawater chemistry in addition to the well-known complexities of turbulence itself. Predicting Arctic climate is only one of the many areas where oceanographers need to understand how fresh water and salt water mix. From heat storage in the mid-latitude oceans to the biological effects of desalination plants, the mixing of salt is a central and poorly understood piece in a wide range of oceanographic puzzles.

In the ocean, mixing results from turbulence acting at very small spatial scales. Mixing of heat and momentum is accomplished mostly by turbulent eddies around a centimeter in size, which are generated by larger scale motions such as breaking waves. Mixing of salt happens on even smaller scales: a millimeter or less. This makes salt mixing especially difficult to measure and also to resolve in computer models.

Numerical simulations are essential tools for research in all areas of turbulence. A simulation of salt mixing must resolve spatial scales of millimeters and less while also encompassing a large enough volume to permit realistic turbulence to evolve. As a result, the memory requirement is extreme. Recent progress in developing specialized computational tools for the efficient resolution of weakly diffusive solutes like salt (Smyth et al., 2005), and with the computing facilities at NCAR, are making possible direct numerical simulation (DNS) of turbulent mixing in salt water for the first time.

Research is now underway to understand the rich variety of mixing phenomena that this simulation has revealed. Further simulations, made possible by expanded computational capability will be designed to learn how saltwater mixing is controlled by various parameters such as the temperature, salinity, and velocity differences between the water layers. An important outcome will be the development of formulas to represent the effect of salt mixing in the models that oceanographers use to understand global ocean currents, intermediate-scale mixing, and climate evolution.

The Arctic warming described above is far from the only application to benefit from this new knowledge. For example, at moderate latitudes, the ocean surface is warmed by the sun each day. Much of that heat remains at the surface where it can be transferred to the

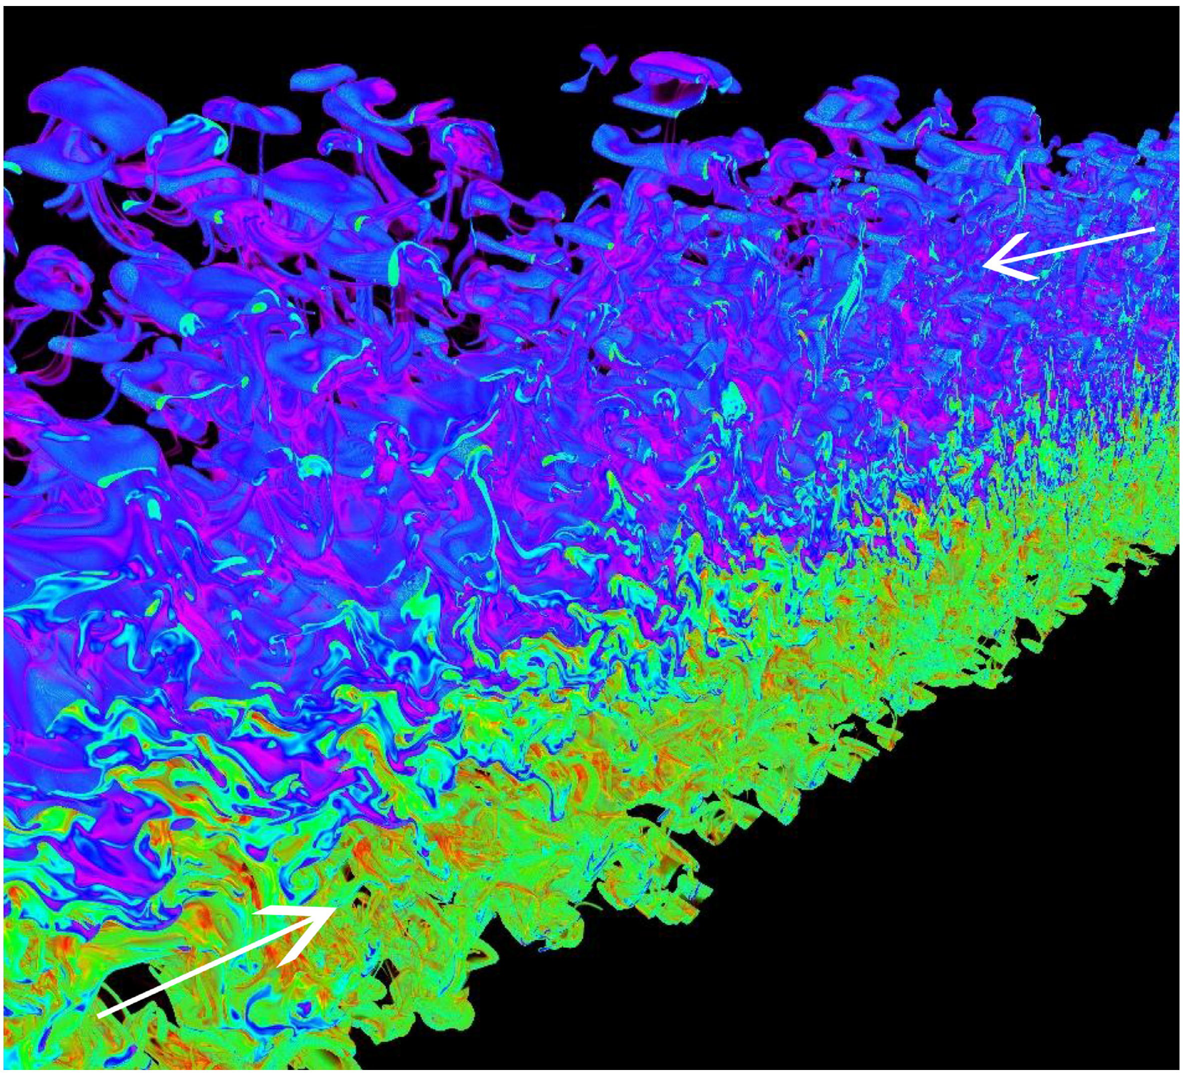

Figure 2.18. Salinity in a direct numerical simulation of saltwater turbulence caused by the mixing of two water layers. The upper (blue-purple) layer is relatively warm and salty, and is flowing to the left. The lower (yellow-red) layer is cooler and fresher, and is flowing to the right. The kinetic energy of the currents – as well as the vastly different molecular properties of heat and salt – generates turbulence that results in enhanced fluxes of momentum, heat, and salt between the water layers. The simulated volume is 2 m high x 4 m long x 10 cm wide.

atmosphere. Warming, however, is accompanied by evaporation, which increases the salinity of the surface water (because evaporating seawater leaves its salt behind). Increased salinity makes the surface water denser, and can therefore cause it to sink. Depending on the action of saltwater turbulence, the heat it contains may be released to the atmosphere the next night, or not until many centuries later.

Finally, desalination facilities extract freshwater from the ocean, and in the process generate plumes of hypersaline water that spread into the surrounding ocean, with potentially devastating effects on sea life. As understanding of saltwater turbulence improves, it will become possible to predict the time needed for salinity to mix down to tolerable levels, and possibly to direct the saline discharge to minimize harmful effects in the meantime.

2.7 Modeling Multiphase Subsurface Flow for Carbon Sequestration

Development of high-fidelity geological reservoir simulation tools represents a longstanding problem of particular interest to both the University of Wyoming and the State of Wyoming. Much of the state's oil reserves lie in the so-called residual oil zone, below the traditionally defined base (oil-water contact) of a reservoir. Enhanced oil recovery techniques based on CO_2 injection are currently under investigation for many of these reservoirs.

Oil, gas, and coal will continue to be important elements of the U.S. energy portfolio until other energy sources are developed to meet demand. This will take decades. Our immediate challenge is to handle the carbon emissions associated with use of fossil fuels. Underground carbon dioxide sequestration based on recaptured utility emissions increasingly is being advocated as a viable approach for mitigating global climate change. Several large-capacity reservoirs are being characterized for sequestration potential and the monitoring tools are under development by UW researchers. The State of Wyoming now leads the nation in developing legal and regulatory framework for carbon dioxide sequestration because they see such sequestration as an imperative. These tools (including advanced computer simulations), strategies, and legal frameworks will be necessary if the U.S. is to meet the President's climate and energy goals.

Although increasingly sophisticated simulation models have been developed over the years to evaluate quantitatively various subsurface flow/transport phenomena such as aquifer management, contaminant cleanup, hydrocarbon extraction, and CO_2 sequestration, significant gaps remain between the predictive capabilities of these models and observed subsurface processes. On one hand, the wide range of scales involved in these problems, from the media pore size to the scale of the entire reservoir, leads to an insatiable demand for increased resolution, while at the same time requires the development of upscaling methods for practical simulations. Additionally, the general inability to fully characterize the heterogeneity of subsurface geological strata leads to large simulation uncertainties that must be characterized using sensitivity analysis and uncertainty propagation methods. High-resolution reservoir simulations with sensitivity analysis or uncertainty management techniques will be instrumental in evaluating CO_2 injection, determining optimal injection rates, and assessing CO_2 storage capacity and potential leakage. A petaflops facility will enable the development, validation, and demonstration of useful high-resolution simulations, while at the same time enabling the development and deployment of novel but computationally intensive techniques incorporating more accurate heterogeneity, time-dependent sensitivity analysis, and stochastic uncertainty propagation methods. Many of these methods, which are currently considered computationally intractable, have the potential to revolutionize the current limits of accuracy and reliability for simulating subsurface phenomena.

Following are six components of current research at UW involving numerical simulation of subsurface flow that will make use of large-scale computing. Together, these projects provide critical elements of sequestration site characterization, monitoring, verification and accounting required for successful carbon dioxide injection and storage.

2.7.1 Multiscale Modeling and Numerical Simulation of CO_2 Injection

A multidisciplinary collaboration at the University of Wyoming is developing new, integrated software models that will take advantage of new hardware and new numerical methods for partial differential equations, aiming at producing scientifically correct, fine-grid, accurate numerical simulations of multiphase, multicomponent flows in heterogeneous porous media. The new models will exploit the capabilities of new heterogeneous multicore, hybrid CPU/GPU hardware that requires the investigation of new numerical kernels. These optimized kernels will serve as building blocks for applications deployed on future heterogeneous petaflops/exaflops systems currently under development.

The new numerical kernels will take advantage of recent developments in numerical methods for the approximation of multiphase, multicomponent flows in porous media, including operator-splitting techniques with multiscale finite volumes for elliptic and parabolic problems and high-resolution central schemes for hyperbolic systems of conservation laws.

First-principles simulations are of utmost importance in producing reliable predictions of the migration, trapping, and possible leakage of CO_2 plumes in the subsurface. The primary research goal is to continue developing a new simulation tool for compositional (multiphase, multicomponent) flow in multiscale heterogeneous porous media. Traditional simulations of multiphase flows rely on ad-hoc upscaling techniques. The models for multiphase flows are defined at the lab scale (a few centimeters) while simulations of interest to important problems (such as CO_2 sequestration) have to be performed in the field scale (a few kilometers). Such ad-hoc techniques, frequently developed and tested for single or two-phase flows, may produce serious errors when applied to more complex compositional flows.

UW plans to use state-of-the-art numerical techniques to solve large-scale compositional flows directly in fine computational grids, which properly resolve the underlying physical heterogeneities. The new multiscale parallel software will use MPI, OpenMP, and the new standard OpenCL, and it will adapt itself to the type and number of available processing cores. Distinct, multiscale domain decomposition strategies will be used in the parallel solution of elliptic, parabolic, and hyperbolic systems. Large global linear or nonlinear problems will never be constructed. The combination of multiscale finite volumes for elliptic and parabolic systems with semi-discrete central finite volume schemes for hyperbolic systems will allow the group to produce new, very accurate simulations of multiphase flows of economical and/or environmental interest.

The group at UW has recently built small GPU servers that can be used for preliminary tests of the new kernels. The UW team is currently building a new GPU server with 15,000 cores. Research goals include:

- Development of optimized kernels will serve as building blocks for applications on future heterogeneous petaflop/exaflop systems, such as the NWSC.
- Very-high-resolution numerical simulation of the injection of CO_2 into brine aquifers.
- Preliminary Monte Carlo studies of the injection process, aiming at the quantification and reduction of uncertainty.

2.7.2 Optimal Geological Model Complexity in Carbon Geostorage Modeling

Research is underway that integrates multiscale permeability upscaling, parallel computing in multiphase flow simulation, and a computationally efficient sensitivity analysis to address the question of optimal geological model complexity in building carbon storage reservoirs for performance assessment of CO_2 sequestration in saline aquifers. The workflow consists of the following steps:

Multiphase flow simulations are conducted in a multi-million-cell three-dimensional high-resolution sedimentary dataset (see Figure 2.19) containing a geologically realistic heterogeneity pattern. Results will define a geological "ground-truth" model in terms of describing the detailed locations of CO_2 flow, storage, and leakage.

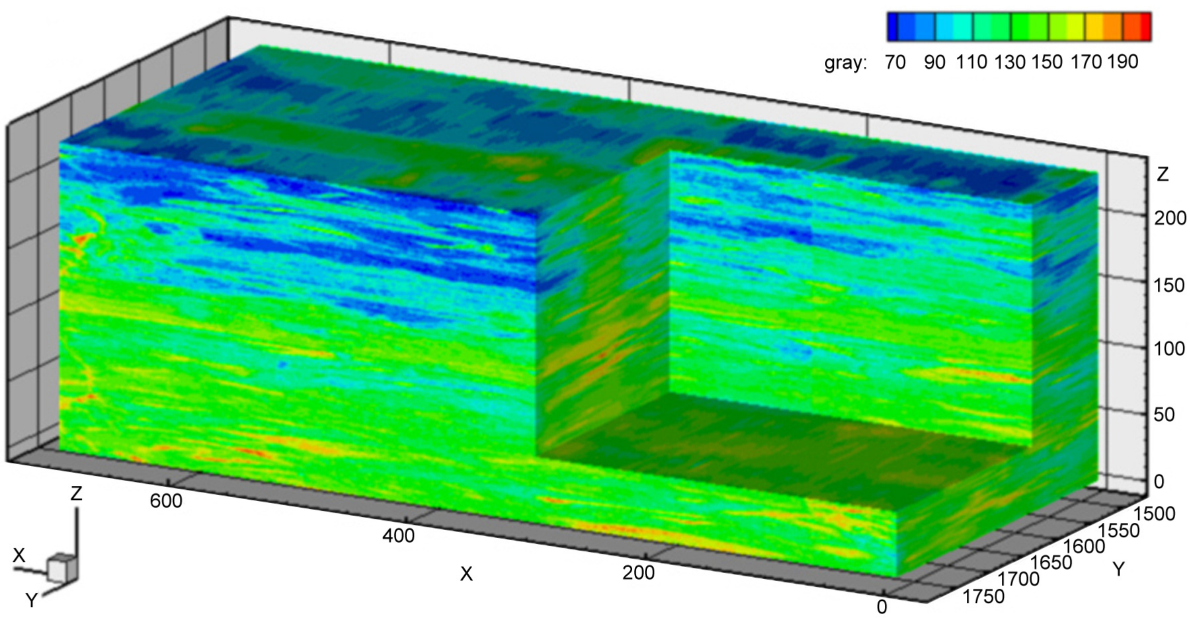

Figure 2.19. Block diagram illustrating a set of deposits created in the Experimental EarthScape Facility, San Anthony Falls Hydraulic Lab, University of Minnesota. This experiment provides the opportunity to understand the impact of heterogeneity on flow/transport, and the impact of using effective parameters (effective conductivity, macrodispersivity) to represent the "missing" or unresolved heterogeneity on model predictions. We will explore the impact of three-dimensional heterogeneity on flow/transport (a cutaway reveals the internal sedimentary structure which controls permeability heterogeneity) and assess changes in the prediction envelope and model sensitivity due to effective parameterizations in the upscaled models.

Based on a new permeability upscaling method, increasingly simpler, homogenized models can be created corresponding to decreasing model complexity, thus reducing the cost of reservoir characterization (see Table 2.1).

A numerical sensitivity analysis will be conducted in all models created from the large sedimentary dataset. Equivalent homogenized models will be developed similarly to those shown in Table 2.1 below. The simulations will be conducted using PFLOTRAN, a parallel multiphase simulation code developed by the Los Alamos National Laboratory for predicting CO_2 storage in saline aquifers. Results will: a) identify the most important parameter uncertainties impacting CO_2 flow in each model; b) create a prediction envelope for each model using the computationally efficient Response Surface approach; and c) identify the optimal model complexity by comparing results of a) and b) predicted by the homogenized models with those of the ground-truth model.

Insights from c) will help guide the development of a large storage reservoir model for potential receiving formations, including the Tensleep/Weber Sandstone, Madison Formation, and Bighorn Formation, deep saline aquifers in southwestern Wyoming. A model of appropriate geological complexity will be built for this formation. Using the parallel flow code, CO_2 injection will be simulated in this model. The prediction envelope will be compared to field responses to injection.

Table 2.1. Geologic models to be evaluated for optimal model complexity.

Geological Ground-truth Model	**Increasingly Homogenized Models:**		
	Facies Scale	Facies Assemblage Scale	Formation scale
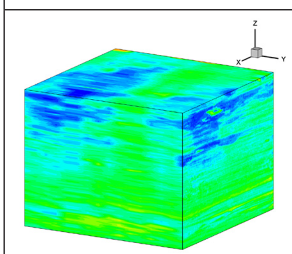	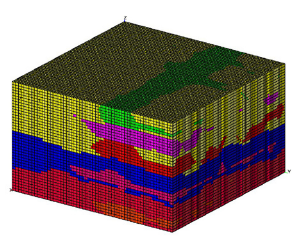	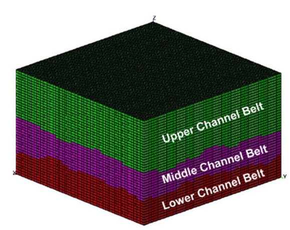	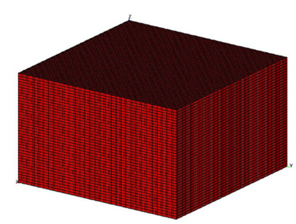
Permeability varies in three dimensions	Permeability varies among facies	Permeability varies among depositional environments	Permeability is uniform within a storage reservoir
Model conditioned to high resolution 3-D seismic, well logs, outcrop analogs, and core measurements	Model conditioned to 2-D seismic reflection data, well logs (from which facies are identified), outcrop analogs	Model conditioned to 2-D seismic reflection data, well logs, outcrop analogs	Model conditioned to 2-D seismic reflection data or existing geological structure or isopach maps

The workflow is first being tested on a multi-million-cell high-resolution sedimentary dataset before being validated for a large CO_2 field test in one of Wyoming's proposed carbon dioxide storage reservoirs.

2.7.3 Evaluating the Impact of Large-scale CO_2 Storage on Freshwater Resources

Commercial-scale carbon geostorage requires thick, regionally extensive deep saline aquifers that only exist in Earth's sedimentary basins. In the western USA, sedimentary basins occur throughout the Rocky Mountain region, and they can provide storage sites for major CO_2 point sources. In the arid west, however, fresh groundwater is an important resource in many communities that rely on it for sustainable domestic use, agriculture development, and industry. Such a resource, often found in shallow aquifers of the basins, needs to be protected against possible contamination due to saline brine displacement by large volumes of CO_2 injected into deep geological formations.

Researchers plan to integrate geological model building at sedimentary basin scale, regional water chemistry data, and large-scale scientific computing resources to address this critical issue in balancing carbon geostorage with freshwater protection in the Rocky Mountain

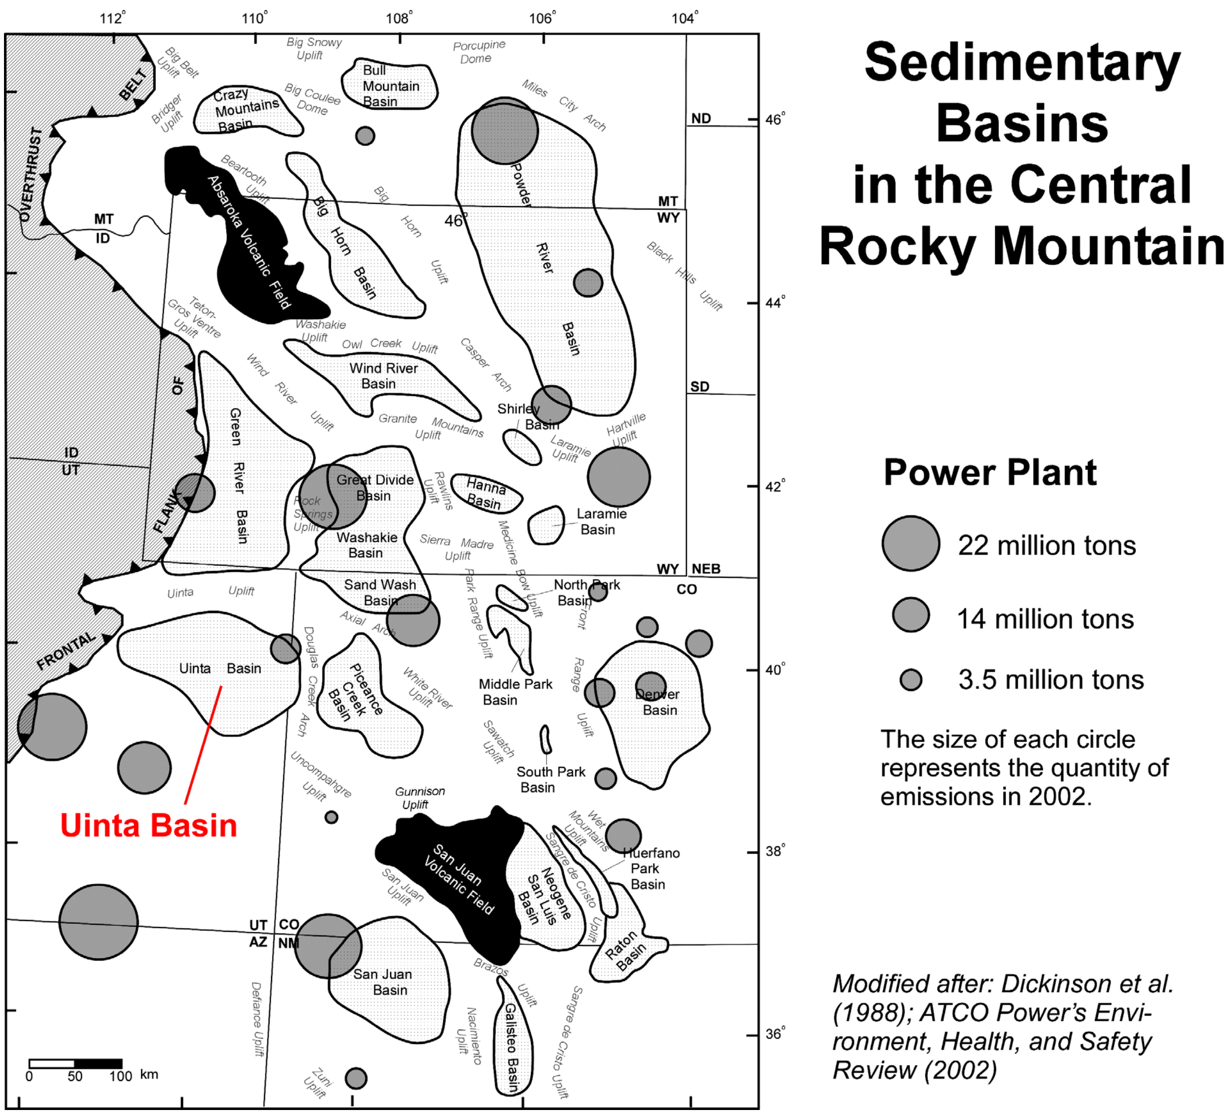

Figure 2.20. Major sedimentary basins in the central U.S. Rocky Mountains containing deep saline aquifers. Many are located near major sources of CO_2 emissions.

region. The Uinta Basin in Utah is a pilot study site for which the freshwater saline water interface has been mapped at the basin scale. Against this background, commercial-scale sequestration in the deep formations of this basin will be simulated using parallel computing to make large-scale as well as detailed predictions of CO_2 flow and migration paths of displaced brines toward the surface. Results will be used to develop a site-specific strategy for CO_2 storage in this basin that can then be extended to evaluating other basins in the Rockies and around the world. Based on the calibrated model, commercial-scale multiphase CO_2 flow in deep formations of the Uinta Basin will be simulated using large-scale scientific computing. Brine displacement at the regional scale (a prediction envelope will be created for each injection scenario) will be modeled. An optimal CO_2 injection strategy (location, rate, duration) will be designed to minimize the impact of large-scale injection on shallow freshwater resources of the basin.

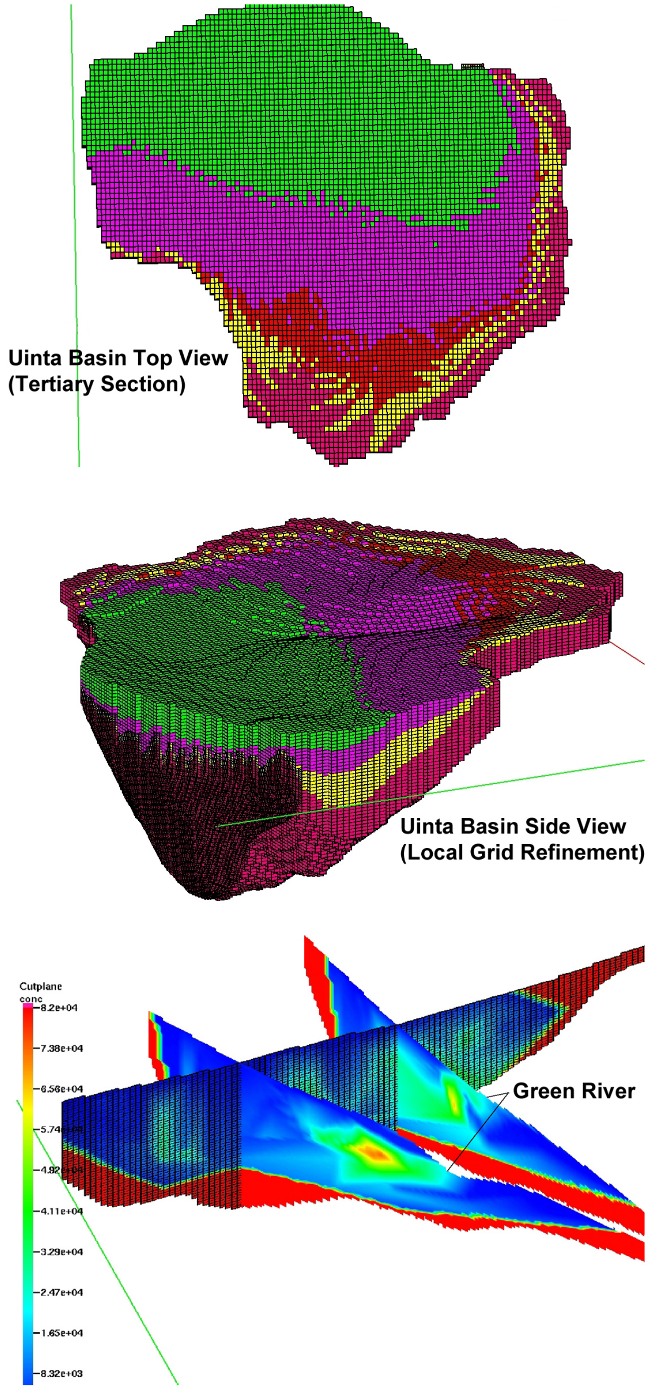

Figure 2.21. The 3-D distribution of geologic facies and the total dissolved solids (TDS) of the ground water associated with the various geologic facies.

2.7.4 Monitoring Carbon-sequestered Saline Aquifers

Monitoring is required by the proposed EPA regulations for Class VI carbon dioxide injection wells. Monitoring of carbon dioxide injection wells requires both robust seismic inversion methods to provide accurate images of the subsurface and also flow simulation models to predict the post-injection CO_2 saturation within the aquifer and to calibrate seismic inversion with such simulations. Scientists at UW have research objectives to combine reservoir flow simulation with multi-component seismic waveform modeling and inversion. In particular, they plan to combine the models predicted from reservoir flow simulations with advanced seismic imaging techniques. Using a three-dimensional model of seismic (P- and S-wave velocity and density) and flow (porosity, permeability, etc.) properties from existing well logs from a proposed sequestration site at the Moxa Arch in southwestern Wyoming, they will synthetically sequester

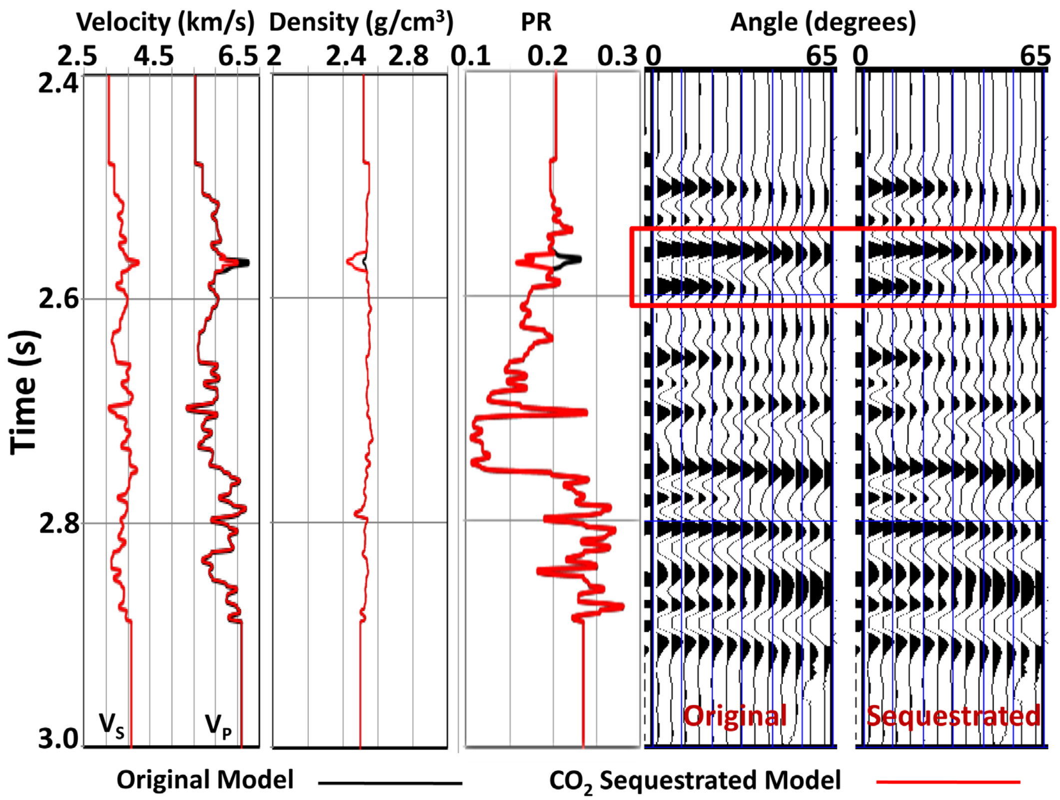

Figure 2.22. Real well-log data from Moxa Arch, southwestern Wyoming, shown in two-way P-wave time and computed vertical component full waveform synthetic seismograms. The original well-log data is in black. A brine-saturated sand formation immediately above 2.6s is replaced with 50% brine and 50% CO_2, and these CO_2-sequestrated well-log curves are in red. Notice that the red and the black curves are identical except at the sandstone formation where brine is replaced by CO_2.

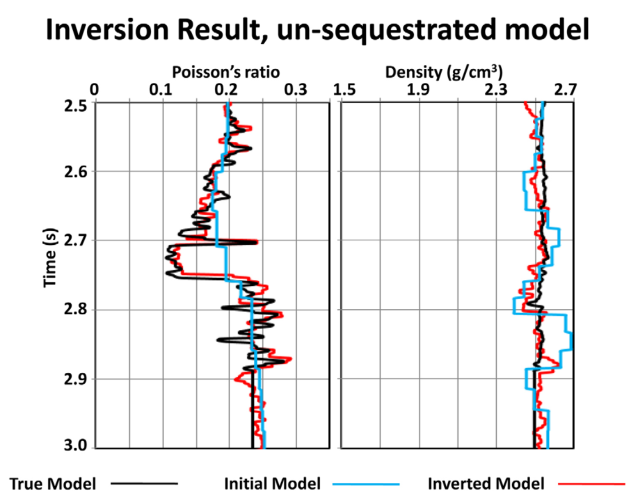

Figure 2.23. Inversion result of the unsequestrated synthetic seismograms of Figure 2.22. The inverted Poisson's ratio and density are shown and compared with the true model. The initial model, shown in cyan is chosen from the P-wave stacking velocity and average Vp-Vs and Vp-density relations.

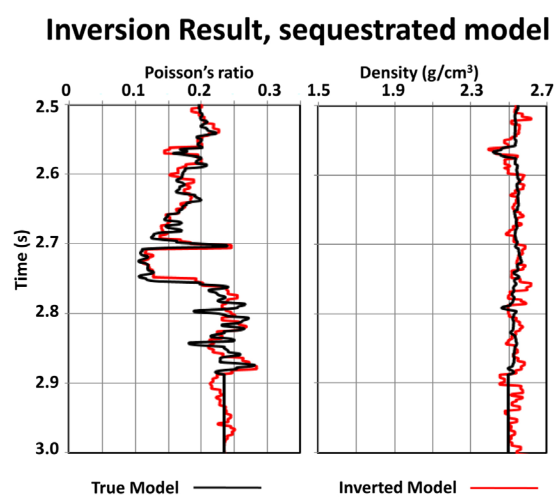

Figure 2.24. Inversion result of the CO_2-sequestrated synthetic seismogram of Figure 2.22. The initial model for this inversion is the final inverted model (red curves) of Figure 2.23.

CO_2 from one of the wells and run reservoir flow simulation models to study the post-injection movement of CO_2 as a function of time over the 3-D model. Using 3-D seismic property models from the simulation at different time intervals of the flow simulation, they will compute finite-difference synthetic seismic responses for the original unsequestrated and post-injection models, process the computed responses, and apply prestack multi-component waveform inversion to predict if post-injection CO_2 saturations could be predicted from seismic inversions. This work will provide optimum parameters for CO_2 sequestration and time-lapse seismic data acquisition.

Although the flow simulations can be performed using a variety of available flow simulation software, this project will benefit from a finite-element multiphase fluid flow model that is being developed. This approach, with fine-scale gridding, is well suited for petascale computing. An angle domain waveform inversion method is also being developed that is computationally intensive and would be of appropriate scale for the NWSC facility.

2.7.5 Improving Estimates of CO_2 Sequestration Potential

Deep saline aquifers in two huge geologic structures in southwest Wyoming are the most promising targets for geologic CO_2 sequestration in Wyoming. The carbonate and sandstone

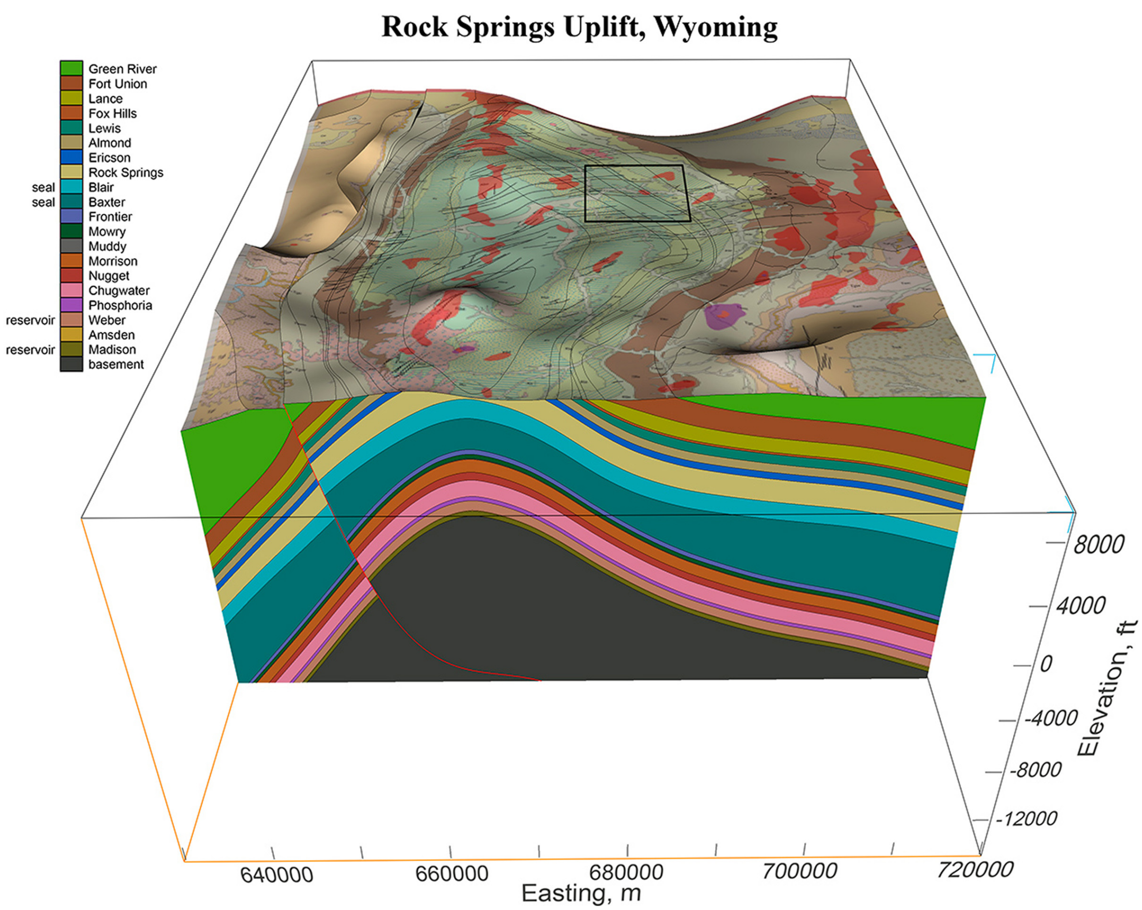

Figure 2.25. Three-dimensional geologic framework model of the Rock Springs Uplift constructed using formation tops picked from well logs. EarthVision software was used to create the 3-D geologic model. The targeted CO_2 sequestration reservoirs, the Weber and Madison formations, are characterized by four-way closure and up to 1500+ meters of overlying, low-permeability Cretaceous shales. The black rectangle on the surface is the 16 km by 16 km domain depicted in Figure 2.26.

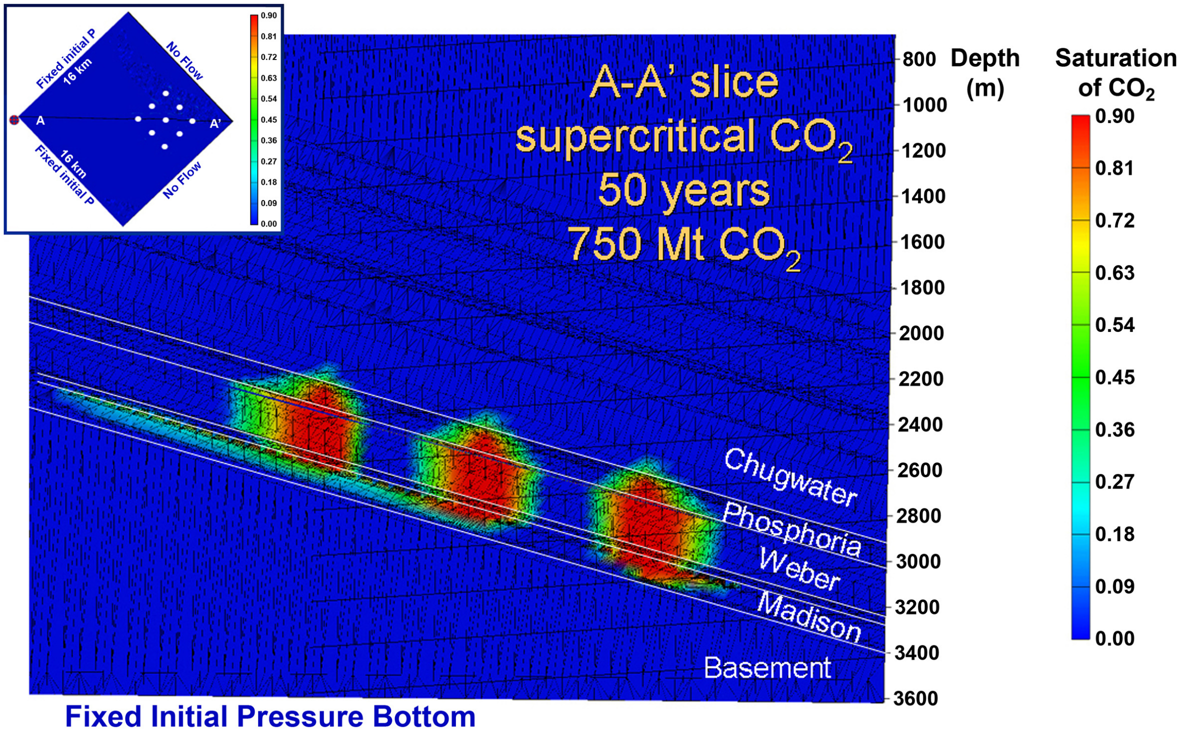

Figure 2.26. Cross section through three of the injection wells showing the dispersal of injected CO_2 after 50 years of injection at a rate of 1.67 million tons CO_2/year in each injection well. Modified from Stauffer et al., 2009.

aquifers lie at depths below drinking water sources and at pressures and temperatures for which CO_2 will be supercritical. Preliminary data obtained by UW and the Wyoming State Geological Survey from prior research funded by DOE suggest that the Rock Springs Uplift and Moxa Arch can safely store Wyoming's current annual CO_2 emissions for more than 750 years. Three-dimensional analysis and numerical modeling that incorporate detailed geologic data is required to provide complete performance assessment of these sequestration sites prior to actual sequestration activities.

High-resolution multiphase models are required to refine the initial capacity estimates. Current models have been run using FEHM, the Los Alamos National Laboratory multiphase porous flow simulator. The model's grid spacing is 200 meters in the x and y directions and with variable spacing in the vertical direction. Future, higher-resolution simulations using NWSC resources will enable more detailed geologic data to be included in the models, leading to more accurate placement and precise optimization of the number of injectors, precise depths for injection, and more detailed simulation of the movement of the CO_2 plume and displaced fluid.

2.7.6 Adaptive Particle-based Pore-level Modeling of Incompressible Fluid Flow

UW researchers use a three-dimensional fully dynamic parallel particle-based model for direct pore-level simulation of incompressible viscous fluid flow in disordered porous media. This model takes directly as input three-dimensional high-resolution microtomography images of naturally occurring or man-made porous systems. These images provide the most faithful representation of the pore space in the porous medium. Over the last several decades, scientists have used, for instance, various pore networks to represent the pore space of the porous systems in their flow models. The majority of these models, however, have used regular two- or three-dimensional networks that do not faithfully represent the random nature of the pore space in naturally occurring porous systems.

During the last 5–10 years, random three-dimensional networks have become available mostly for sandstones using x-ray images or thin sections. These models have been used successfully to predict two- and three-phase flow properties in sandstones. But even in these networks, the pores of the rock sample are represented by idealized geometries. Transformation of pores from an x-ray image, for instance, to idealized geometries inevitably leads to loss of many topological features and thereby introduces uncertainties into the predicted flow properties. Also, the methodologies for construction of reliable pore networks for other rock types, e.g., carbonates, are not fully developed yet. Finally, the existence of very large pores can lead to formation of dynamic effects that are not reliably addressed by dynamic pore network models.

The uncertainties associated with non-direct pore-level models, e.g., pore network models, provides strong motivation for modeling the fluid flow directly in three-dimensional high-resolution voxel images of the porous medium (without converting it into a network). In this method, the entire medium, including both solid and fluid phases, is discretized using particles. The model is based on the Moving Particle Semi-implicit (MPS) technique and modified to improve its stability. It handles highly irregular fluid-solid boundaries effectively, accounts for viscous pressure drop in addition to gravity forces, conserves mass, and can automatically detect any false connectivity with fluid particles in the neighboring pores and throats. The model also includes a sophisticated algorithm to automatically split and merge fluid particles to maintain hydraulic connectivity of extremely narrow conduits. Finally, it uses novel methods to handle particle inconsistencies and open boundaries.

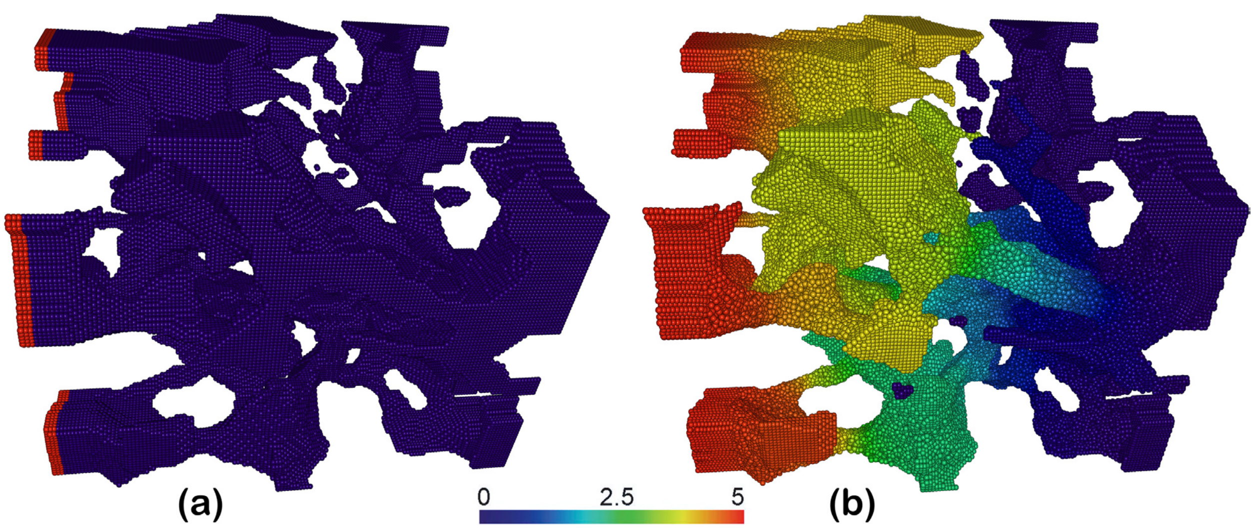

Figure 2.27. Static pressure profile for a flow simulation in a high-resolution microtomography image of a naturally occurring sandstone with the upstream and downstream parts added. (a) t=0. (b) tss=0.1743 s. The system is originally stagnant at atmospheric pressure and at t=0 the upstream pressure is increased to 5 Pa while the downstream pressure is kept constant at zero. Notice the unconnected fluid particles remaining at the initial condition. Total number of particles: 2,016,000; Resolution: 3.398 μm; Number of voxels: 112×112×134; Sample size: 0.394 mm × 0.394 mm × 0.475 mm; Sample volume: 0.0739 mm^3; Sample porosity = 0.222.

Clearly, the computational cost required by this numerical algorithm is enormous. To handle the computational load, we have developed a fully parallel version of the model that runs on distributed memory computer clusters and exhibits excellent scalability. However, even with this large cluster system we are limited in sample size and resolution. Petascale computing will significantly advance our ability to model and understand multiphase fluid flow in real geologic systems.

As an example of how petascale computing will enable us to accelerate our research in this area, Figure 2.27 shows the static pressure profile for a flow simulation in a high-resolution microtomography image of naturally occurring sandstone. The static pressure profile is a direct function of the connectivity and pore structure in the sample. Here, the disconnected pore bodies do not experience any of the applied pressure difference and remain at their initial conditions throughout the simulation. The dynamic pressure profile is a function of velocity and follows its transient behavior (see Figure 2.28). This simulation took about 48 hours on 216, 2.3-GHz processors.

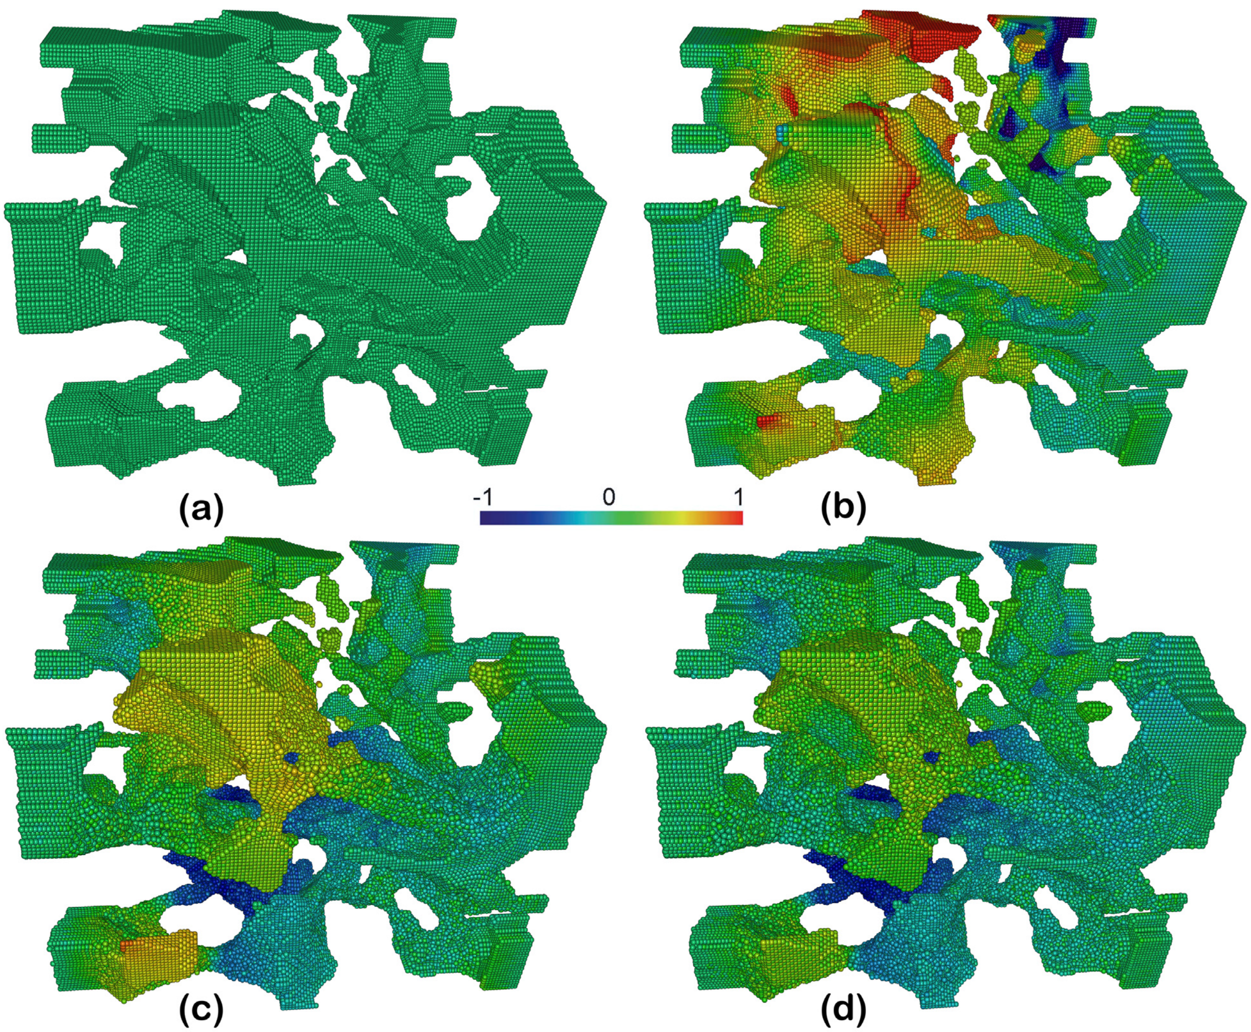

Figure 2.28. Dynamic pressure profile for a flow simulation in a high-resolution microtomography image of a naturally occurring sandstone. (a) t=0. (b) t=0.0045 s. (c) t=0.0879 s. (d) tss=0.1743 s.

These results are from a flow simulation in a 0.065 mm^3 piece of sandstone. We need to carry out direct flow simulations in much larger samples, e.g., 114,000 mm^3, a typical core sample. Many of the macroscopic flow properties, e.g., capillary pressures and multiphase relative permeabilities, are often measured at core scale and used in the prediction of improved hydrocarbon recovery schemes and CO_2 sequestration scenarios. Measurement of these properties, however, is often difficult, time consuming, and expensive. For instance, in three-phase (oil, water, and gas) flow problems, it is impractical to measure these flow properties. This is because experimental measurements of three-phase relative permeabilities and capillary pressures are extremely difficult to perform, and at low saturation the results are very uncertain. Two independent fluid saturations are required to define a three-phase system, and there are an infinite number of possible fluid arrangements, making a comprehensive suite of experimental measurements for all three-phase displacements impossible. This is why numerical simulations of three-phase flow almost always rely on available empirical correlations. These empirical models may give predictions that vary from each other or from direct measurements by as much as an order of magnitude because they have little or no physical basis. Therefore, they introduce significant uncertainties into the predicted reservoir performance.

It is therefore imperative to develop a reliable physically based tool – such as the particle-based model presented here – that can provide plausible estimates of macroscopic properties. However, these models have to carry out simulations at larger scales, e.g., core-scale, so their predicted flow properties can be used in reservoir-modeling simulators. To do this, much larger supercomputers are needed. Availability of petascale computational resources will allow us to perform flow simulations at core-scale with the number of computational units (particles) on the order of 1×10^{12}. This means that we will be able to bridge the gap between the pore-scale and the core-scale without needing ad-hoc up-scaling techniques. This will allow us to directly investigate the impact of micron-level transport phenomena on the core-scale flow behavior and the predicted flow properties. This can significantly reduce the uncertainties associated with our prediction of performance in improved oil recovery schemes and CO_2 sequestration scenarios through utilization of more physically based flow properties predicted by the direct flow simulations at core-scale.

We plan to use petascale computational resources to carry out direct flow simulations in various rock samples, e.g., carbonates and tight gas sandstones, at increasingly larger scales. These examples show that the computational intensity necessary to model complex multiphase fluid flow in the Earth's crust is comparable to that developed for the atmosphere. Currently, the actual structure of the subsurface is not well depicted in numerical simulations. Petascale computing will allow geoscientists to more accurately describe the complexity of geologic structure and thereby improve model predictions. This will become increasingly important as regulation demands accurate monitoring of carbon dioxide plumes in the subsurface. Proposed regulations require that injected carbon dioxide be monitored for decadal-to-century time scales and over very large volumes of the Earth's subsurface. Core-scale flow simulations could potentially save tens of millions of dollars per year in the costs of special core analysis tests.

2.8 Crosscutting Statistical, Mathematical, and Computational Science Techniques

Addressing the modeling and forecasting needs articulated in the previous Earth System science sections will require both the availability of more computer power and commensurate advancements in the relevant computational sciences disciplines, namely computer science, applied mathematics and numerical methods research, geostatistics and data assimilation, and geoturbulence. An important component of future Earth System science research programs is an ecology and environment that promotes the development and integration of more accurate numerical methods, more scalable and more efficient algorithms, better statistical models, and better physical representations of sub-grid-scale phenomena into models of the Earth System.

The crosscutting computational sciences, working together with hardware improvements, can radically accelerate simulation capabilities. For example, Collela et al., 2003 use the example of the history of the solution techniques of the electrostatic potential on a uniform cubic grid to show that "algorithms yield a factor comparable to that of the hardware, and the gains typically can be combined (that is, multiplied together)." Numerical techniques that may have shown promise (i.e., asymptotic error behavior, scalability) but were previously considered noncompetitive may become the techniques of choice as computational hardware evolves toward higher levels of parallelism (Collela et al., 2003: p. 32, Fig. 13). Techniques such as discontinuous Galerkin methods (Cockburn et al., 2000, Nair et al., 2005), time-implicit integration schemes (Staniforth and Côté, 1991, Rizzetta et al., 2003), and adaptive-mesh refinement techniques (Fulton 2001, Jablonowski et al., 2006) – once considered either too complex or too costly – may become indispensable in the future due to their scalability and the flexibility they offer for adaptivity in spatial and temporal resolution. The ability of investigators to explore the potential advantages of such new approaches and bring them to fruition will be greatly accelerated by improving the availability of high-end computing resources to the geosciences community.

2.8.1 Computational Science to Support Petascale Application Development

Over the last several years, computer scientists and software engineers have made concerted efforts to significantly improve the scalability of the various components of the upcoming Community Climate System Model version 4 (CCSM4). This work has included improvements to the partitioning algorithm in the Parallel Ocean Program (POP) and Community Ice CodE (CICE) at 0.1° resolution to increase efficiency on large processor counts (Dennis, 2007, Dennis and Tufo, 2008). A parallel I/O (PIO) library was developed for use by all component models in addition to improvements in the scalability of the memory footprint. The attention to the scalability of the entire CCSM system has enabled the first-ever coupling within the U.S. of an eddy-resolving

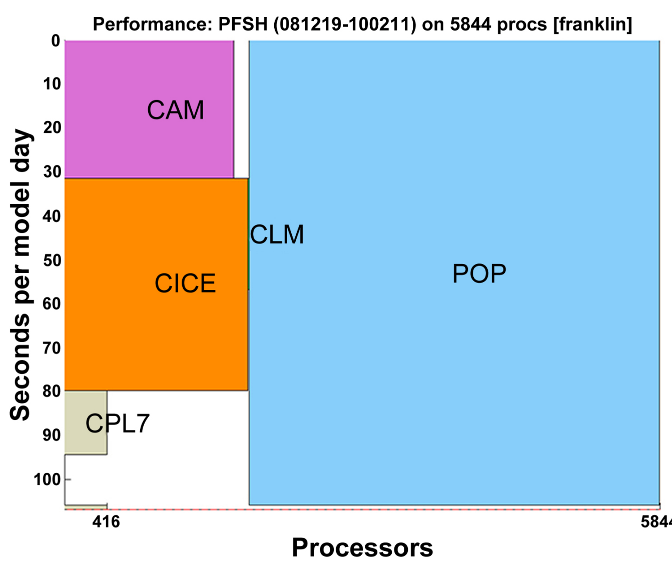

Figure 2.29. Performance of high-resolution CCSM with 0.5° atmosphere and land models coupled to 0.1° ocean and sea ice models on 5,844 processors of Franklin, a Cray XT4. The area of each box represents the relative computational cost (number of processors times seconds per model day) of each component model.

ocean and sea ice model at 0.1° degree to an ultra-high-resolution atmospheric model and land model at both 0.5° and 0.25° resolution.

Figure 2.29 illustrates the computational cost, excluding disk I/O costs, to simulate a single model day of a CCSM configuration with 0.5° atmosphere and land models coupled to 0.1° ocean and sea ice models. The x-axis shows the number of processors, and the y-axis shows the number of seconds each component takes to execute a model day. Figure 2.30 is a plot of the scalability of the simulation rate of high-resolution CCSM for the NERSC Franklin system, the NICS Kraken XT4 system, and the NICS Kraken XT5 system. Note that we do not include the time to write history and restart files.

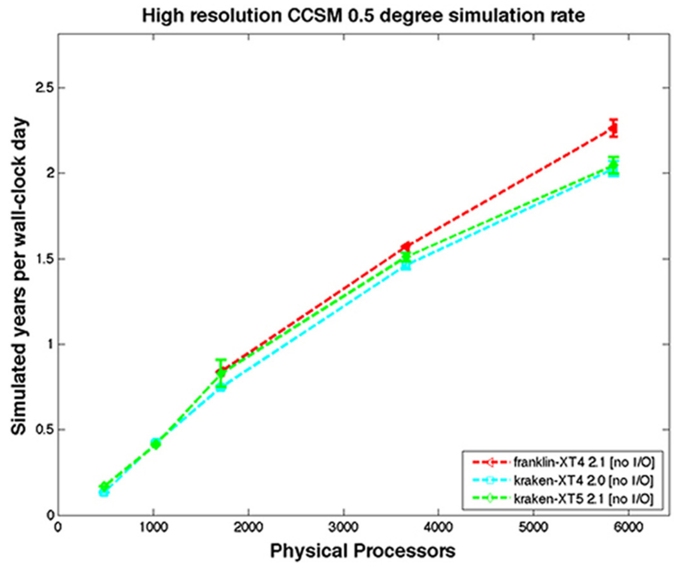

Figure 2.30. Simulation rate of high-resolution CCSM4 with 0.5° atmosphere and land models coupled to 0.1° ocean and sea ice models on Franklin, a Cray XT4, and the Cray XT4 and XT5 versions of Kraken.

2.8.2 Ensemble Data Assimilation

Numerical weather prediction has been one of the great success stories of geophysics. By combining ever-increasing numbers of atmospheric observations with progressively larger and more realistic models, prediction centers have produced increasingly skillful forecasts since initial efforts in the 1960s. Data assimilation is the process of combining observations with short-term model predictions to get optimal estimates of the atmospheric state. Assimilation not only produces the initial conditions required for forecasts, but by confronting atmospheric models with observations it is also a central part of the model improvement process.

Data assimilation has not been as prevalent for models of other components of the Earth System. This is due to the historical dearth of observations and the enormous personnel costs of developing traditional data assimilation systems. However, a host of new *in situ* and remote sensing observations are now available for the ocean, cryosphere, land surface, biosphere, and near-earth space.

New ensemble data assimilation algorithms are now available that are simple to develop for any Earth System model. The time is ripe to extend the experience gained in weather prediction to models of the entire coupled Earth System. Ensemble assimilation for both component models and coupled models can generate high-quality initial conditions for predictions. At the same time, the ability to confront models with observations can greatly accelerate the process of improving models.

Ensemble assimilation algorithms use a set of model forecasts, called an ensemble, to sample the probability distribution of a model given the available observations. The ensemble can be used to generate mean values of quantities like temperature, but also variances and covariances. These covariances are used to determine how observations adjust the ensemble

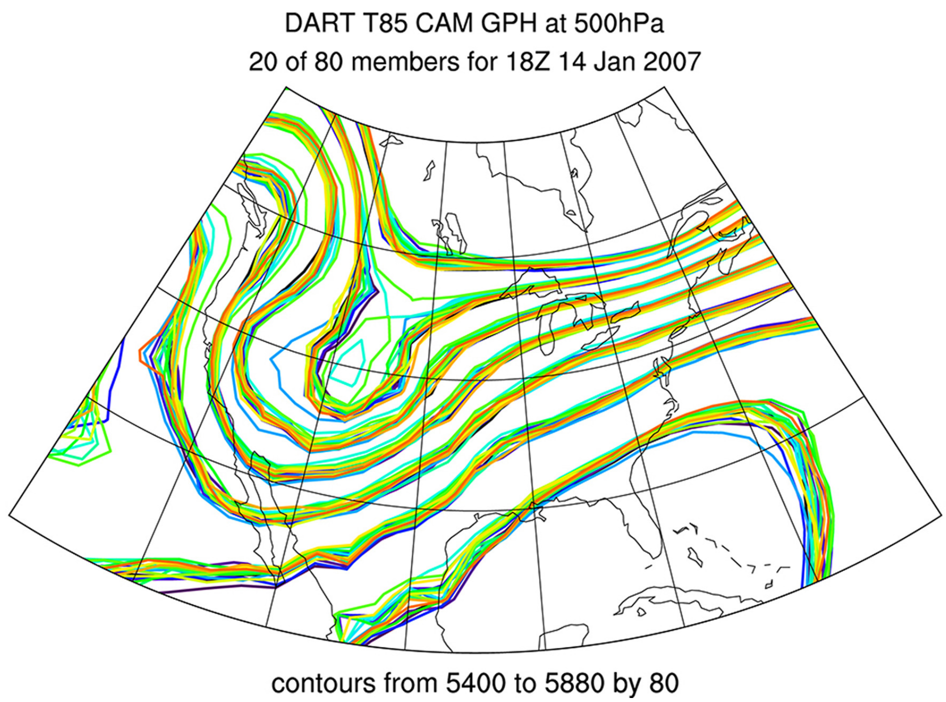

Figure 2.31 An ensemble 6-hour forecast of the 500-hPa height field for 18 GMT on 14 January 2007 produced using the Data Assimilation Research Testbed (DART) and the Community Atmosphere Model (CAM). Contours from the same ensemble member are shown in the same color. Forecasts are more certain where the contours are more similar and less certain where the spread is greater. Being able to sample a forecast probability distribution like this also facilitates powerful tools for model diagnosis and improvement.

estimate. They can also be used as a powerful analysis tool to increase understanding of the dynamics of both the model and the Earth System.

Ensembles of nearly 100 members are required for assimilation in Earth System models. This means that assimilation systems require 100 times the storage and computation of a single model forecast. As Earth System modelers press on to more complex simulations at higher resolutions, only the largest of next-generation computing facilities will be able to convert climate models designed for simulation into predictive systems. The results will be improvements in models, prediction, and our understanding of the Earth System.

2.8.3 Turbulence Modeling

Turbulence is crucial to comprehending a broad variety of effects in meteorology, climatology, oceanography, and ecology, as well as in solar-terrestrial physics, astrophysics, and of course, engineering and technology. Turbulence arises from nonlinear interactions at all spatial and temporal scales present in the system, and it is therefore a quintessential petascale application involving a very large number of degrees of freedom. Rotation, stratification, and magnetic fields can be among the many components of turbulent flows; understanding all of these will be enhanced by substantial increases in our numerical capacity. The interactions of nonlinear eddies and waves (e.g., inertial, gravity, Alfvén) lead to a slowing down of energetic exchanges among scales, a geometric ordering, and possibly to a better predictability of such flows.

Another aspect of turbulent flows that may be studied using petascale computing is the notion of intermittency. It plays a role in many contexts: reactive flows, convective plumes, combustion, the chemistry and dynamics of the atmosphere, heating of the solar corona, or reconnection events in the terrestrial and planetary magnetospheres; it may also be at the origin of the random reversal of the magnetic field of the Earth. Intermittency is thought to arise from the interplay between turbulent structures and is diagnosed through non-Gaussian statistics of field gradients. At present, the precise role of these structures is not understood, but our capacity to analyze their statistics will be significantly enhanced by the resolution provided by petascale direct numerical simulations (DNS), possibly helping to predict the occurrence of extreme events (e.g., the maximum velocity attained in a hurricane or a tornado). Together with petascale computations, we need to develop algorithms for data reduction and structure extraction.

Furthermore, turbulence modeling skill, in engineering as well as for geo- and astrophysics, will continue to be needed simply because the number of excited modes in such flows far exceeds the capacity of computers in the foreseeable future. As the Reynolds number of DNS grows, tests can be devised that study the properties of turbulence models in detail and thus allow improvements—or generalizations—to handle more complex flows. For example researchers need to account for complex boundaries and geometry, or anisotropy in the presence of either rotation, stratification, or magnetic fields. But because realistic parameters are still well out of range for today's DNS, one can foresee complementary roles for DNS and modeling together with experiments and observations. Savings at given values of the parameters (e.g., Reynolds, Prandtl, and Rossby numbers) can be substantial, so such models should prove useful in exploring parametrically dynamical regimes of geophysical and astrophysical turbulence in a variety of conditions resembling how they arise in nature.

To make such improvements, we will have to develop and test extensible and versatile simulation codes, e.g., by evaluating new scalable programming paradigms and introducing new domain decompositions and new performance optimization methods that will scale to the petascale. We need to combine these advances with data reduction techniques for embedded and non-embedded analysis and three-dimensional visualization (based on tools such as the NSF ITR-funded VAPoR software) (Clyne and Rast, 2005). Data analysis and visualization also needs automated coherent structure extraction. Such improvements can lead to applications that will improve numerical weather prediction, enhance accuracy of remote sensing algorithms for aircraft safety, and improve our understanding of space weather. Petascale applications will also allow for graduate and post-doctoral training, and will have many educational impacts.

An advanced application example is shown in Figure 2.32. This rendering shows a strongly helical turbulent flow undergoing rotation at a Rossby number close to that of the Earth's atmosphere. The computation is performed on a regular grid of $1{,}536^3$ points (in excess of 3.6 billion). The most striking aspect of this discovery of new structures is the juxtaposition of complex vortex filaments and laminar columns, viewed here in the vorticity field. If helicity (i.e., the correlation between the velocity and its curl) were absent, the laminar column (at right) would be absent as well. It is well-known in meteorology that helicity plays a role in the dynamics of atmospheric flows, but recent modeling advances allowed us to fully realize that helicity leads to such strong and persistent structures at intermediate to small scales in the presence of rotation.

When examining the behavior of flows in which there is an interplay between different effects leading to different spatio-temporal regimes, petascale computing capacity will allow for resolving some of these regimes; at present they are smeared. For example, we will be able to unravel the interplay between inertial waves, gravity waves, and turbulent eddies in rotating stratified turbulence in the ocean, as well as the scale-by-scale energy transfer between quasi-geostrophic and sub-mesoscales. Similarly, remote sensing and aviation safety can potentially be enhanced by detailed studies of stratified turbulence allowing, again, for scale separation between different phenomena.

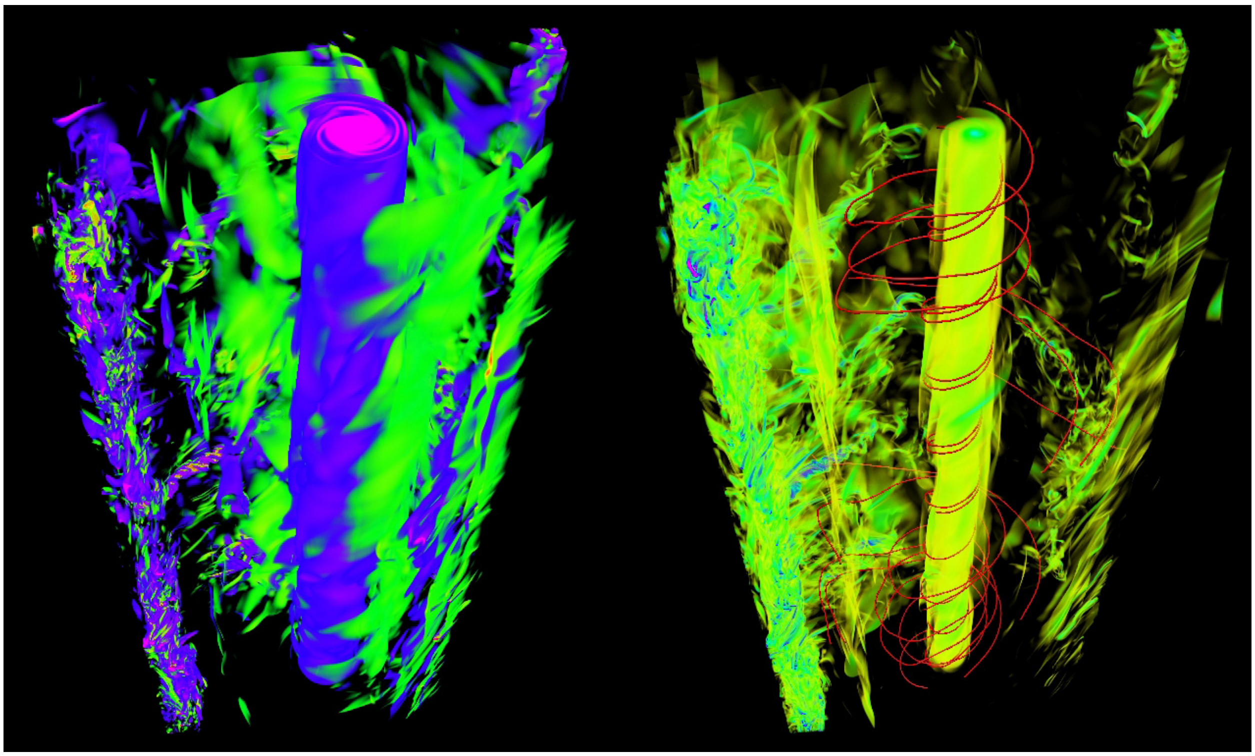

Figure 2.32. Helicity (correlations between the velocity and its curl, the vorticity) is shown on the left in a zoom of a DNS computation performed on a regular grid with more than 3.6 billion points in a turbulent flow at a Rossby number close to that of the atmosphere of the Earth. The NSF ITR-funded VAPoR software tool is used for visualization. Helicity measures the strength of the nonlinear terms, which are very weak in the vertical laminar column, which is a co-rotating cyclonic event. At right is a zoom performed on the vorticity and displaying a complex assemblage of intense vortex filaments whose radius is close to the dissipation length. Juxtaposed to these filaments is a wider laminar column in which the trajectories of fluid particles (represented as red lines) are regular. These columns are long-lived and fully helical, i.e., with a strong alignment of the velocity and its curl.

3. Technical and Research Infrastructure Requirements

This section derives the technical requirements for the petascale system, data infrastructure, and the organizational, operational, and support requirements from the science objectives described in Section 2.

The overall justification of the NWSC project is based on the evidence presented in Section 2 that making dramatically increased computational resources available to the Earth System science community will enable transformative science. Further, Section 2 shows that the impact extends across a number of related scientific disciplines. More computational power allows scientists to move model configurations to higher resolutions, lengthen the timescales or the size of the regions simulated, include new and better science, assimilate observations for new predictive capabilities, or expand ensemble sizes for better statistics. However, a big increase of computational power alone is not sufficient to guarantee success: the facility must be designed with the requirements of a successful petascale enterprise for the Earth System sciences in mind. In such an enterprise, the petascale system is only the centerpiece, surrounding it must be the all-important data infrastructure: the disk, archive, servers, and networking that are essential to allow scientists to manipulate their results and share them with collaborators. In addition, the infrastructure must be operated by an organization focused on and committed to the success of the Earth System science community, and capable of providing the software tools and services needed by the community to succeed.

Commonalities. The following commonalities in the Earth System sciences are readily apparent from the science objectives described in Section 2:

Algorithmic similarities. Many computational geoscience simulation problems involve the basic mechanics and thermodynamics of fluids and are thus united by a family of common solution techniques for the equations that govern these phenomena.

Multiscale structure. Geofluids dynamics typically give rise to turbulent phenomena that produce multiscale structures. The multiscale nature of the Earth System presents modelers with sub-grid-scale parameterization and closure issues. Therefore, turbulence research is fundamental to making progress in many of the science frontiers in atmospheric, oceanographic, and space weather research.

Interdisciplinary problems. Earth System modeling is the ultimate interdisciplinary problem, requiring the close collaboration of a wide variety of Earth System scientists, physical scientists, biologists, social scientists, computer scientists, applied mathematicians, and software engineers to harness massive amounts of computing power to tackle these problems.

Multiple components. Many parts of the Earth System are multiphase, or multi-component, or can be represented by models running on a hierarchy of scales – thus the models themselves tend to have multiple components. Therefore, model-coupling software infrastructure is critical to scientific progress.

Predictive requirements. Numerous multi-component problems are attempting to develop new predictive capabilities via data assimilation. This is in part in response to the explosion of new observation systems and the vast observational data streams they represent.

System Technical Requirements Summary. The following points summarize this section's requirements analysis. Parentheses indicate the primary source of the requirement:

- Based on Earth System science research objectives, a dedicated computational system in the range of 1-1.5 peak petaflops is required in 2012. The requirements for the system include:
    - The system is estimated to require O(100,000) cores to achieve this level of performance (based on conventional multicore processor design trends).
    - The possibility of deploying accelerator co-processors in part of the system should be explored (based on subsurface multiphase flow modeling objectives).
    - The processor supporting the fastest executing thread available is desirable (e.g., for the paleoclimate simulation throughput requirement).
    - The minimum per-core memory requirement for the system is 2 gigabytes per core (based on high-resolution climate simulation). A portion of the computational system may contain at least 2–3 times this minimum to support more memory-intensive codes (such as MHD turbulence). Total system memory required is estimated to be at least 200 terabytes for a 1 petaflops system in 2012.
    - The interconnect bisection bandwidth required for turbulence simulation is estimated at 250 megabytes/sec/core for a core capable of 1 gigaflops sustained (based on spectral turbulence models).
- Approximately 10% of the system resources should be devoted to data analysis and visualization (DAV) nodes (based on the study of Loft et al., 2005). These nodes should have several times more memory per processor than the compute nodes.
- Both the system and its DAV component should be attached and share a high-speed disk system. The disk system should have the following characteristics:
    - Disk bandwidth of up to 300 GB/sec (aggregate) (tornado simulation/global ocean modeling).
    - Assuming a nominal bandwidth of 100 megabyte/sec/disk, the system in 2012 is expected to hold 6 to 15 petabytes of disk.
- It should have an archive system sized to support an initial data production rate in 2012 of 23–35 petabytes/year, and a minimum aggregate sustained I/O rate of 8 GB/sec, with an expected growth rate of 34% per year thereafter.
- The system and its data holdings should be connected to other national cyberinfrastructure via wide area networks capable of supporting bandwidths equivalent to multiple 10 gigabit/sec links, and integrated via grid technology (based on the needs of geoscience workflows, including large computational campaigns such as IPCC).

3.1 Requirements Analysis Details

This discussion covers the details of how the petascale system and data infrastructure technical requirements were derived from the science objectives in Section 2, as well as the similar input and analysis in Bryan et al., 2005 and Loft et al., 2005.

Computing scale. The petascale is a transformative computational scale for Earth System research: many of the longstanding questions and long-sought capabilities first become accessible there. However, the question remains how much Earth System "science " would a petascale system produce? A conservative estimate for a 100,000-core petascale system is that it could deliver as much as 750 million allocable CPU hours per year. While it is difficult to answer precisely what the use pattern of these resources would be, an estimate can be obtained from mapping the 2009 NCAR usage of the Bluefire system, by discipline, onto the petascale system. The usage at NCAR was distributed in the following way: roughly 54% for climate science, 22% for special science campaigns, 10% for meteorology, 5% for space weather, 4% for oceanography, 3% for basic fluid dynamics, and 2% for a reserve. If such use patterns held up at the NWSC, the following simulations could be realized each year:

- For climate science, 2,500 simulated years of high-resolution CCSM (0.5° atmosphere and land with 0.1° ocean and sea ice) could be run plus 200 years of nested regional climate (assuming a high-resolution climate and 1% refinement with WRF running at 4 km) plus 3 simulated years of 4 km global-cloud-resolving atmospheric models for process studies.
- Each year, 57 times as many CPU-hours than are currently allocated on Bluefire could be devoted to special scientific campaigns.
- For weather, 20,000 simulated days (54 years) of Continental United States-sized WRF weather simulation runs at 2.5 km could be completed.
- For oceanography, 4% of the resource would allow 6,000 simulated years of 0.1° global ocean or 750 simulated years of 0.05° global ocean to be completed.
- For solar physics, scientists could simulate active regions on the Sun over the lifecycle of a sunspot assemblage (1 week timescale) and still have another 13M CPU-hours/ year for additional space weather research.

An important technical goal of the petascale system housed in the NWSC is that it must allow U.S. Earth System science research to remain competitive internationally. For example, users of the NWSC system will have to compete intellectually with scientists with access to dedicated systems with similar missions, like the Japanese Earth Simulator-2. The ES-2 is currently being placed in operation at the Yokohama Institute for Earth Sciences. The ES-2 consists of 1,280 vector processors with 20 TB of main memory and a theoretical performance of 131 teraflops. It is composed of 160 processor nodes connected by Fat-Tree Network. Each processor node has 8 vector processors and 128 GB of main memory. The peak performance of each vector processor is 102.4 gigaflops. For comparison, the peak theoretical speed of the IBM POWER6 core in the NCAR Bluefire system is 18.8 gigaflops, 5.5 times slower than the ES-2 vector processor. In addition, the efficiencies achieved on vector processors are generally higher than on microprocessors, with an efficiency on one atmospheric application of 64.9% reported, and 30-50% efficiencies more typical; while single-digit efficiencies are usually realized with microprocessor systems.

Two things may be concluded from these observations: first, it will take many more microprocessor cores to match the performance of a single vector processor: perhaps as many as 50 times more. Second, researchers using microprocessor-based systems have to build significantly more explicit application parallelism into their codes in the form of MPI and OpenMP calls.

Processor speed. The highest possible single-processor speed is required by climate (and especially paleoclimate) models to achieve integration rates of a few simulated years per wall clock day (Benioff, et al., 2005, Ramanathan, et al., 2009). This requirement has driven this field's affinity for vector processors in the past, and in the future may require these researchers to exploit the Single Instruction Multiple Data (SIMD) vector capabilities of microprocessors (e.g., Intel's SSE3 or IBM/Apple's AltiVec) or perhaps accelerator co-processors (e.g., Cell or GPUs), if the programming hurdles of these alternatives can be overcome for climate applications.

Accelerator co-processors. Section 2.7.1 describes the development of subsurface modeling applications that are targeting future 64-bit-capable Graphics Processing Units (GPU) as application accelerators. Developers of WRF and other atmospheric science application have also been experimenting with accelerators. It is clear that the capabilities of accelerators are rapidly improving and they beginning to appear in operational supercomputer configurations offered by vendors. For these reasons, an accelerator option for at least part of the system should be investigated for the NWSC system in 2012.

Large system memory size. Memory requirements in geoscience applications are generally driven by either the high degree of model complexity (which results in a very large number of variables and a large working set size), or in the case of models of underlying turbulent phenomena, the large numbers of degrees of freedom in the system being simulated. Lack of sufficient memory per core can necessitate running such applications in inefficient configurations, thereby wasting resources, or even worse, not being able to run them at all.

For climate simulation, memory use is driven by the trend toward adding more phenomenology to models. This trend has been in place for over a decade and is expected to continue. For example, extensive networks of chemical and biological feedbacks have been added to climate and ocean models in recent years (Bryan et al., 2005, Fennel et al., 2006, Batchelder et al., 2002). Even after extensive memory optimization, current memory requirements of high-resolution instances of CCSM have been measured at 1.4 GB/core. This value is expected to continue to grow as still more climate phenomenology is added.

Turbulence research is at the foundation of many Earth System science problems. The computational requirements of turbulence applications, many of which are based on 3-D spectral methods, have been analyzed in Loft et al., 2005. That study found the interconnect requirements of turbulence applications to be extreme (see below). For a fixed problem size, the communications overhead limits scalability, and system memory size becomes the limiting factor in optimizing such applications: for example, a $2{,}048^3$ problem can require about 6 GB/core to run efficiently on about 1,024 processors.

Interconnect bandwidth and latency. Two examples of Earth System science applications that constrain interconnect design are the barotropic conjugate gradient solver in the Parallel Ocean Program (POP) and the inter-process communication requirements of turbulence models employing 3-D spectral (3D-FFT) methods. The solver used in the barotropic component of POP is based on the single inner product preconditioned conjugate gradient method (D'Azevedo et al., 1993). Nevertheless, the scalability of global reductions (i.e., MPI_ALL_REDUCE) and the 2-D halo update within the conjugate gradient solver is very sensitive both to interconnect latency and to operating system noise at high processor counts (Ferreira et al., 2008). In the case of 3-D spectral methods, for an architecture capable of sustaining 1 gigaflops per core for an FFT, the per-processor bisection bandwidth required to reduce the overhead of the transposition-based communication algorithm to acceptable levels is in the range of 250 MB/sec/core – a challenging figure to maintain for systems with large 3-D torus interconnects.

Robust I/O capability. Resolving the disparate temporal scales in Earth System applications often involves saving long time series of data at high output frequency. For example, ultra-high-resolution tornado simulations are projected to require robust I/O capabilities of up to 300 GB/sec (Loft et al., 2005). Also, five instances of a 100-level 5-km global ocean model, running simultaneously on an O(100,000)-core petascale system while performing daily history file output, will require 300 GB/sec of aggregate bandwidth to keep I/O overhead from impacting system performance. The aggregate I/O bandwidth determines the minimum number of disks in the attached storage. Using the canonical figure of 100 megabytes/sec/disk, a minimum of 3,000 disks are required in the attached storage system. To achieve high bandwidth, applications must perform parallel I/O. Good parallel I/O performance is a challenging task for application developers, and it is another example of the benefits of close collaboration between system administrators, computer scientists, and application experts.

Disk space. Individual disks are expected to hold 2–5 terabytes in 2012, so the 3,000 disks yield a disk space requirement of 6-15 petabytes for the system. This is a disk space-to-system memory ratio of 30:1 to 75:1.

Data analysis and visualization. Because of wide area bandwidth restrictions, data should be kept and analyzed as near to the petascale system as possible (Washington et al., 2008). Based on a review of peer atmospheric research centers around the world, the Loft et al., 2005 study recommended that 10% of the system computational capacity (not including GPU computational capabilities) should be dedicated to the analysis and visualization (DAV) of data. To avoid users having to unnecessarily duplicate data, these DAV resources should share disk storage in a common pool with the system's compute nodes. A robust stack of discipline-specific analysis and visualization tools and services are required. Examples include GrADS data servers[16], network Common Data Form[17] (netCDF) data libraries, netCDF Operators[18] (NCO), VIS5D[19], and NCAR Command Language[20] (NCL) graphics libraries. This tool chain will have to be adapted to handle very large datasets or supplemented by new tools with that capability (Washington et al., 2008).

---

16 See http://www.iges.org/grads/gds/ .
17 See http://www.unidata.ucar.edu/software/netcdf/ .
18 See http://nco.sourceforge.net/ .
19 See http://vis5d.sourceforge.net/ .
20 See http://www.ncl.ucar.edu/ .

Data management services. Washington et al., 2008 succinctly describes the set of data management services required, including collaborative visualizaton; distributed data discovery through ontologically based metadata (e.g., as is provided by the Earth System Grid project); distributed computation, analysis, and management of data; and support for application/module sharing.

Archive. To estimate the tape archive space requirements needed for a petascale facility, we can simply scale up the current data archive rate of Bluefire to that of a 1 petaflops system. The currently measured rate on this 70-teraflops batch system is about 4.4 terabytes/day or 1.6 petabytes/year. Thus, we estimate an archive system scaling factor of 23 petabytes/year/petaflops – in other words, a 1 to 1.5 petaflops system is estimated to produce 23 to 35 petabytes/year. However, this methodology assumes no changes in workload archive characteristics between now and 2012. This may not accurately reflect the collective behavior of a group of scientists doing new types of research on a petascale system. Also, this figure may be an overestimate if increased model resolution dominates the direction of future modeling efforts because in that case, computational complexity will grow faster than the size of data output files. Alternatively, this may be an underestimate if the demand for the archival of new observational and reanalysis data outpaces existing levels.

Networking. Many geoscience applications are interdisciplinary and have complex workflows: some, such as numerical weather prediction, also have real-time constraints. For climate, large computational campaigns such as the IPCC Assessments are run on many systems worldwide, and the data output must be federated to facilitate its analysis as a single cohesive dataset. For such problems, grid technology is a key enabler to support collaboration across centers via distributed scientific workflows. In terms of performance, GridFTP and file transfers initiated by shared, wide-area parallel filesystem technology are needed to facilitate these workflows and achieve the necessary sustained transfer bandwidths. Currently, single user transfer bandwidths of 300 megabytes/sec are considered adequate by high-resolution climate researchers. These requirements will likely go up as the data volumes produced by future experiments increases. In addition, the NWSC's networking infrastructure will likely need to support multiple users performing simultaneous transfers. Thus, even higher aggregate bandwidths will be required by the system, and WAN network bandwidths equivalent to multiple 10 gigabits-per-second links will be required.

3.2 Operational, Organizational, and Support Requirements

An ecosystem is perhaps the most useful analogy to describe the necessary organizational structure for the supercomputing enterprise. As the Graham et al., (2004) *Future of Supercomputing* report stated:

> *Supercomputing is not only about technologies, metrics, and economics; it is also about the people, organizations, and institutions that are key to the further progress of these technologies and about the complex web that connects people, organizations, products, and technologies.*

Similarly, the Earth System science facility and the HPC CI it contains, to be most effective, must be operated by a deeply focused organization that has the specialized knowledge and understands the specific niche requirements of Earth System science applications. Yet it must also be broad enough to encompass the multidisciplinary research involved to support it, and also be well connected to other organizations and facilities that serve related science missions in the wider national and international context.

Summary of operational, organizational, and support requirements. The following points summarize the analysis of the operational, organizational, and support requirements.

- Earth System models are complex. The computing system must be able to compile and run existing applications without extensive modifications.
- The Earth System sciences require a dedicated system operated by a domain-specific organization.
- The system must be well integrated with other national cyberinfrastructure.
- To effectively exploit petascale systems, science application developers need support from computational scientists. Resources should be designated to create and support the interdisciplinary teams needed to make applications – particularly community models – more scalable and efficient.

Why a discipline-specific center? Earth System sciences are a unique collection of scientific disciplines united by deep commonalities and a compelling mission that warrants a discipline-specific approach. The rationale for discipline-specific infrastructure follows those set forth by the Kinter and Seidel (2008) NSF panel as well as voices from the scientific community. This rationale includes the following points:

- *Unique technical requirements.* The technical requirements of important scientific applications in this domain are unique. In particular, climate simulation involves long-running, complex models that have specific requirements for the highest possible processor speed, large system memory size, robust I/O capability, and discipline-specific data analysis and archival services. The discipline-specific approach to computing allows services and system design to be tailored specifically to meet these requirements.
- *Expensive validation and verification procedures.* Climate models require controlled development environments because they are exceptionally sensitive to small numerical changes, such as changing the order of calculation of results, which can result from compiler or supporting run-time library upgrades. Verification and validation requirements of climate models demand exceptionally long simulations, making them extraordinarily costly to revalidate during development. Access to massively parallel development systems must also be available for development work to effectively target future leadership-class systems for science campaigns. The operation of such resources by a discipline-specific HPC organization ensures that model development and related research in Earth System processes can occur in a controlled, yet responsive, environment where researchers can prepare and test complex models. In this role, the discipline-specific HPC organization serves as a developmental staging ground for runs on even larger HPC resources elsewhere.

- *On-demand computing and support of field and computational campaigns.* A discipline-specific HPC organization provides the organizational focus, capabilities, and skill-sets required to provide on-demand and real-time computing resources and data services required by natural disaster forecasting, modeling support for field campaigns (cf., Appendix A.7), support of computational campaigns such as the IPCC Assessments, and to aid workshops and educational programs (cf., Appendix A.8).

- *Successful track record.* The discipline-specific approach is currently demonstrating consistent success supporting the atmospheric and related sciences community. For example, NCAR's scientific computing enterprise has a long and enviable track record of service to the atmospheric and related sciences community (Semtner et al., 2009, Leovoy et al., 2002, Dunning et al., 2001, Semtner et al., 2006, Good et al., 2006). The strongest evidence of this is that despite the availability of enormous computing resources elsewhere, most disciplinary scientists choose to compute, analyze, and archive their data at NCAR.

- *Data assimilation.* Earth System research has a strong predictive component. Predictive systems have workflows consisting of (sometimes large) observational data streams generally feeding into ensemble or variationally based data assimilation systems that require running the underlying model many times: in parallel in the case of ensemble methods, and iteratively in the case of variational approaches. Such systems generally have real-time deadlines and stringent end-to-end performance requirements. Traditionally the domain of operational centers, the assimilation of observations will become increasingly common in future Earth System research environments. For example, data assimilation techniques will be needed for climate models to make skillful climate predictions on inter-annual and even decadal time scales.

- *Geoscience data curation.* Making effective use of geoscience data also calls for a considerable amount of in-house discipline-specific expertise, including knowledge of observational systems, data and metadata standards, and data and metadata formats. Critical components of any Earth System science enterprise are careful curation, secure storage, and efficient dissemination of such data (particularly historical observations).

Relationship to and integration with other CI. As noted in Kinter and Siedel (2008), the petascale facility should complement and be integrated with the rest of the nation's high-end cyberinfrastructure. There are multiple rationales for this position. History has shown that certain computational campaigns – such as the IPCC assessments – are so large that the scope of the resources involved assumes planetary scales. Direct experience working at the Earth Simulator in Japan during the last IPCC conclusively demonstrated that it is counterproductive to create a "computational island." Additionally, it is imprudent to try to meet all the computational demands of the Earth System science disciplines within one facility. Thus, for this category of application, the petascale facility we propose will only be a staging area or at best just one component of a much more ambitious global computational enterprise. To support this sort of activity, the facility must be well integrated with other major centers.

While the TeraGrid is not a suitable proxy for the discipline-specific Earth System science research model, it can be an effective provider of supplementary high-end resources for the community. Therefore, we have concluded that the community is better served if the discipline-specific NWSC facility is well integrated via grid technology and services with the HPC facilities of the TeraGrid XD program, as well as those of the NSF "Track-1" Blue Waters system. In addition, the NWSC should also have and maintain adequate bandwidth access to major DOE leadership facilities such as National Energy Research Supercomputing Center (NERSC) and Oak Ridge National Laboratory (ORNL), as well as other national and international peer HPC centers.

In addition, to foster collaboration and to leverage additional mid-range cyber-resources, the NWSC facility should be well connected to key universities and other disciplinary research organizations. Such connectivity will not only offload many computations from the Earth System science facility, but when linked with the community model and application framework paradigms, it will also enable the upscale migration of discoveries and innovations made by university researchers to the largest HPC resources. In addition, such connectivity will provide a technical mechanism for preserving critical data through cross-site replication.

Impact of computing trends on applications. Computational performance improvements are coming from more cores, not faster cores. This means that scientists must achieve high performance levels for their applications through parallel speedup. Because of Amdahl's Law, which establishes a limit on the amount of speedup based on the fraction of parallelism available in the application, gaining parallel speedup becomes increasingly difficult as more processors are involved in the calculation. For some types of calculations with a scalable algorithm, speedup can be achieved through so-called weak scaling, in which the problem size – in terms of the number of degrees of freedom – is increased in proportion to the number of cores. For example, ensemble-based data assimilation techniques can utilize increased processor counts simply by increasing the ensemble size.

Other applications – most notably climate models, which run very long simulations with a relatively low number of degrees of freedom – are not particularly amenable to this type parallelization solution: model integration rate decreases if resolution is increased without increasing processor speed. For such applications, strong scaling is required for a fixed number of degrees of freedom (resolution) to achieve faster integration rate and higher performance levels through parallelization. This is a much more difficult task; scientists must increase parallelism by eliminating serial bottlenecks in their applications. To make progress via strong scaling, scientists must significantly overhaul their time-tested applications, then reverify and revalidate them.

A key finding of two NSF-sponsored petascale geoscience application workshops was that to effectively exploit petascale systems, application developers need extensive support from computational scientists to help them make these sorts of transformations (Snavely et al., 2008). Further, because of the complexity of the underlying modeling systems, Earth System scientists depend heavily on coordinated teams of computer scientists, system administrators, and application experts who all have experience serving the needs of their scientific discipline. Adequate resources of this type must be allocated to create successful interdisciplinary teams with the appropriate expertise.

4. Facility Requirements

Section 3 used the community input in Section 2, plus other findings in panel studies and workshop reports to establish the need the facility to house a dedicated 1–1.5 petaflops peak system along with the associated data infrastructure, as well as other system connectivity requirements.

4.1 Facility Requirements Summary

The following points summarize the analysis of the NWSC facility requirements based on the technical, operational, and support requirements derived in Section 3.

- Based on current technology trends, the NWSC facility initially needs to support no more than 4 MW to house a 1 to 1.5 petaflops peak system.
- The facility design should support water-cooled systems with high heat densities. Systems at that time are expected to have heat densities in the range of 100 kW per rack and be at least partially, if not fully water-cooled.
- Future science priorities and technology trends are uncertain. To mitigate project risk, the available land, facility design, power substations, and other supporting infrastructure should be flexible and expandable, within feasible limits, without overly impacting the project cost.
- The facility's tape archive area should be designed to support an initial data production rate in 2012 of 23 to 35 petabytes/year. Specifically, the total area reserved for tape archive should be sized based on the 2012 tape archive space required to store between 230 and 350 petabytes.
- The facility should be located near and connected via redundant paths to national-scale wide area networks (WAN) and capable of supporting a bandwidth equivalent to multiple 10 gigabit/second channels.
- Recognizing the environmental mission of the project, as well as community and societal concerns that may arise regarding the planned size and power footprint of the facility, every effort should be made – without unduly impacting project cost – to design, build, and operate it as a "green" facility to reduce the impact on the regional environment and climate.

4.2 Estimating System Power Requirements

A key scale parameter, analogous to a "power efficiency" rating for computers, has emerged for facility design: the number of floating point or arithmetic calculations that can be performed per second per Watt. Surveying the November 2008 list of top 10 supercomputing systems, the average value of this parameter for a very large matrix-matrix multiply[21] was, at

[21] This information is from data published on the Top 500 List. Performance is measured by running the High-Performance Linpack (HPL) benchmark and reporting realizable flops or Rmax value for the system.

that time, approximately 200 megaflops on High-Performance Linpack (HPL) per Watt. This translates to 200 teraflops/MW, or 5 MW/petaflops.

However, the Top-10 list encompasses a significant spread in efficiency: the most efficient system on the list can perform 446 megaflops/Watt (achieved by the IBM Cell processor-based Roadrunner system at LANL), and the least efficient performs 81 megaflops/Watt (attained by the older Cray XT-3/4 Red Storm system). This represents a range of 2.2 MW/petaflops to 12.3 MW/petaflops on HPL in the November 2008 Top 10 list. By excluding the two oldest systems on the list, the remaining eight systems on the 2008 list can achieve better than 6.6 MW/petaflops. Note also that high power efficiency is often inversely correlated with usability: the most notable cases are the Cell-based Roadrunner and Blue Gene/P systems. This suggests that current values for *usable petascale systems* are tightly clustered around the scale parameter of 5 MW/petaflops.

Of course, ongoing shrinkage of CPU transistor elements and other system efficiency advancements between now and 2012 will further improve power efficiency figures for petascale systems. Historically, NCAR has observed an industry trend of a 40% per year improvement of the power efficiency of CMOS-based supercomputing systems. Over the three years between November 2008 and early 2012, this trend line predicts a 2.7-fold efficiency improvement in power efficiency. Another estimate of power efficiency can be obtained by assuming that these improvements will be entirely produced by vendors increasing socket core counts with fixed socket power form factors. This model predicts a four-fold improvement in power efficiency in the two intervening chip generations over the same interval. We choose to use the more conservative historical rate of improvement (2.7x) to get an expected value of

What are the chances of a disruptive low-power computational technology emerging?

No radically new "magic bullet" replacement for CMOS semiconductor technology appears likely to become available in the near future and, with a multi-billion-dollar computing industry relying on CMOS, any potential replacement would face enormous adoption barriers. Despite this, engineers are beginning to incorporate incremental techniques for managing power consumption into the design of CMOS-based high-end microprocessors. These techniques include dynamically varying frequency and voltage and turning off circuits on the chip that are not being used. Additionally, the utilization of CMOS variants – semiconductor materials that are inherently more power efficient – is also being pursued. The effectiveness of these power-saving measures in future generations of systems, particularly those in the HPC applications domain, remains unclear.

Finally, as noted above, accelerator chips such as Cell and GPUs are beginning to be used with increasing frequency in large HPC systems and hold some promise to improve the computation-per-watt ratio. Currently, the use of accelerators in geoscience applications must be considered experimental, as serious barriers exist to their widespread use. These barriers include low memory bandwidths between the conventional processor "front-end" and the accelerator, the lack of support for 64-bit computations, rudimentary support for Fortran compilers and math libraries, and the high levels of accelerator parallelism required by some architectures to achieve good performance. However, rapid advancements are being made in each of these areas, and it is thus the focus of intense experimentation and development by vendors and the computational science community.

1.85 MW/petaflops with a high estimate of 2.4 MW/petaflops and a low estimate of 0.83 MW/petaflops. Using these values to delimit possible future procurement outcomes, it is apparent that in 2012, 1-1.5 petaflops could be accommodated in a facility with an initial power footprint of 4 MW, with at least a factor of two in the most likely case (1 petaflops system with average power efficiency), and a 10% margin of error in the worst case (1.5 petaflops system with the worst anticipated power efficiency). Regardless, the large inherent errors in future system characteristics argue for a flexible facility design (and construction contingency) that allows for an order 10-20% expansion of the initial power footprint.

4.3 Impact of Computing Power Trends on Facility Requirements

For many years, computer improvements came from shrinking transistor feature size and increasing clock speed. However, by 2004, the power consumed (and waste heat produced) by faster-clocked chips finally began to reach engineering limits. This happened about the same time that the increasing power and cooling costs of IT systems began to significantly impact the budgets of IT managers and the limits of their facilities. In response to this and other performance trends, supercomputer vendors switched from faster-clocked single-core processors to flat-clocked multicore processors. However, the impact on facilities has as much to do with economics as with semiconductor technology. Over the last decade, the use of mass-market commodity parts for high-end computing systems has improved the price performance of these systems ($/flops) much faster than shrinking transistor sizes has improved power performance (W/flops).

From the perspective of computer users, the trends toward better, faster, and cheaper have been a positive development. But from the perspective of the facility manager who has to deal with the overall server power heat load per dollar spent on computing, the trends have been decidedly negative. Within a few short years, computer heat densities have increased from a few kilowatts (kW) per 19-inch rack to nearly 70 kW per rack for some of the latest-generation systems. Furthermore, these trends continue: designs for 100-kW/rack systems are in vendor roadmaps for the 2012 timeframe. These trends are also causing HPC vendors to shift from air-cooled to partially and fully water-cooled cabinets. Because of these trends, facilities that have successfully supported IT infrastructure for decades are rapidly becoming obsolete.

To illustrate the effect of the power consumption trend on facilities, Figure 4.1 shows the overall power history of supercomputing systems at NCAR. Since the departure of the Cray C90 – the last non-CMOS[22] supercomputer – power consumption has steadily increased with each new generation of CMOS-based server.

Floor space for data archive. As sustained computer performance increases, the data archive rate will continue to increase as well. However, storage densities are also expected to increase and keep pace with these increases, so that in practice, by migrating old data onto denser media, NCAR's experience suggests that we will only need to provision a total data archive space in the facility that is about 10 times greater than the annual data archive rate would suggest at opening – in other words the space should be sized to support 230–350 petabytes of archive in 2012.

[22] CMOS = Complementary Metal Oxide Semiconductor. The most widely used type of semiconductor technology, CMOS is used in practically every type of computing device.

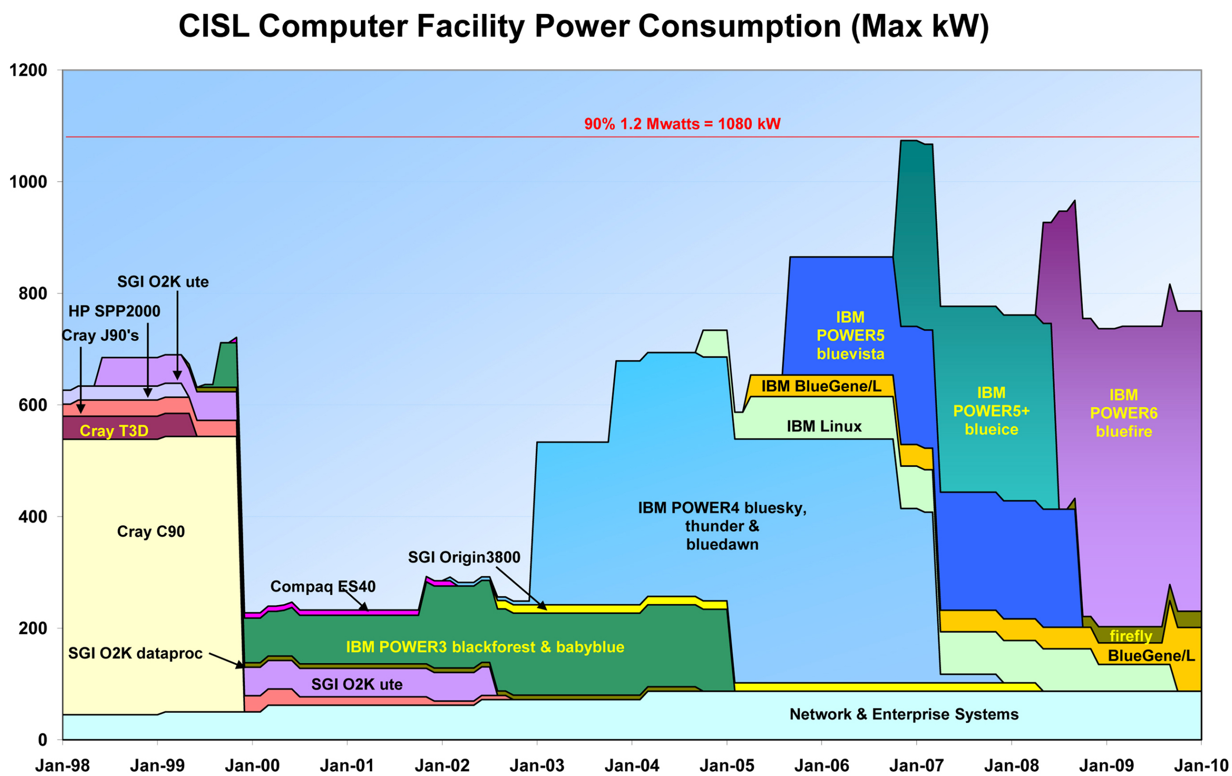

Figure 4.1. Power consumption over the past decade in NCAR's Mesa Lab supercomputing facility. Since the Cray C90 was removed in late 1999, power requirements of CMOS systems at NCAR have steadily increased until facility limits were reached in 2007. NCAR's current program of system consolidation and expansion of the power-efficient Blue Gene/L system will stabilize power usage at around 800 kW in 2010.

Networking. Workflows spanning multiple centers are becoming more common (e.g., the last and current IPCC Assessments), and they have the potential to promote collaborations and discovery by the broader scientific community. Thus, besides archiving data for future analysis, subsets of model output data must be stored on disk and served over wide area networks (WANs). On the TeraGrid currently, 350 MB/sec bandwidths are commonly achievable in parallel (striped) GridFTP transfers between peer centers on a network consisting of single, 10-gigabit/second WAN channels. This requirement will likely grow as larger computing complexes (e.g., the Blue Waters system and the discipline-specific system discussed here) are deployed. Thus, multiple and redundant dedicated 10-gigabit/sec WAN channels will be required to support petascale Earth System research at this facility.

The technological requirements estimated above set the input design parameters of the supercomputing facility and are used to develop the petascale facility's Project Execution Plan (the companion document to this one). These requirements are based on a 1–1.5 petaflops system in 2012, which we think will require a facility initially capable of supporting up to 4 MW of supercomputing equipment, with power densities on the order of 100 kW/rack at that time. The facility should be expandable to mitigate the uncertainties of future system power requirements or an expanded future science mission. It should be provisioned with tape archive space to support an initial data production rate of 23 petabytes/year/petaflops, and an overall archive annual growth rate of 37%. Finally, the facility should be connected to national wide-area networking infrastructure via redundant paths capable of the equivalent of multiple channels of 10 gigabits/sec.

- Future science priorities and technology trends are uncertain. To mitigate project risk, the available land, facility design, power substations, and other supporting infrastructure should be flexible and expandable, within feasible limits, without overly impacting the project cost.
- Recognizing the environmental mission of the project, as well as community and societal concerns that may arise regarding the planned size and power footprint of the facility, every effort should be made – without unduly impacting project cost – to design, build, and operate it as a "green" facility to reduce its impact on the regional environment and climate.

5. Broader Impacts

There is a rising demand by decision makers at all levels of society for improved understanding and prediction of climate and weather phenomena. The decisions to be based on this knowledge have staggering societal implications, and this provides a compelling justification for the NWSC project. Urgent questions about climate change include the impacts of regional change on water supplies, agriculture, and the spread of disease; shifts in the likelihood of extreme weather events such as heat waves, droughts, and floods; and the advisability of specific adaptation and mitigation strategies – the latter including ideas such as removing anthropogenic carbon dioxide from the atmosphere by sequestering it in deep geologic formations. These new issues add to a long list of short-term threats such as severe storms (e.g., hurricanes and tornadoes) and solar flares, for which society has long been demanding improved forecasts. The NWSC is a direct response to these needs, enabling the deployment of high performance cyberinfrastructure that researchers require to perform high-resolution simulations of weather phenomena, global and regional climate, coastal oceans, sunspots, subsurface flow, and more. It will also increase – by over an order of magnitude – the resources available for experimental weather forecasts and support for meteorological field campaigns, thereby accelerating technology transfer between researchers and operational forecasters.

Earth System research and education will be transformed by the creation of the NWSC. The next generation of Earth science researchers and computational scientists will be attracted by the importance of the problem and the scale of the facilities available to them: their talent will be developed through practical research experience gained on petascale equipment from existing and successful education, outreach, and training programs such as UCAR's Significant Opportunities in Atmospheric Research and Science (SOARS), and NCAR's Advanced Studies Program (ASP) and Summer Internships in Parallel Computational Science program (SIParCS). Further, integration of the NWSC facility with other NSF HPCI will provide important cross linkages between these and similar programs at other resource providers. Integration will also broaden the number and types of researchers across the nation making effective use of the facility – particularly for multidisciplinary research – and thereby support the NSF's vision of a transformative national petascale cyberinfrastructure for science and engineering.

A portion of the system will be used by University of Wyoming scientists and their collaborators. This computational resource's impact on the University of Wyoming's academic strategic plans, particularly its growing capabilities in computational science, STEM education, and workforce development will be transformational, attracting faculty candidates, collaborators, and talented students to the state. The NWSC will thus accelerate growth of the University of Wyoming's research capabilities in the related disciplines of atmospheric science, climate science, and ecology.

The NWSC will have also a profound long-term impact on the economy of the State of Wyoming and the region. In particular, the NWSC will provide an anchor to attract industries with complementary computing facility needs. As a catalyst for CI development, IT economic sector growth in Wyoming has the potential to translate into economic benefits to the state such as well-paying jobs.

Both the NWSC and the partnership to build and operate it will be extensible by design. New facility partners can realize economies of scale by investing in allocable portions of large-scale resources typically unavailable to them, or alternatively, by choosing to provide specialized services to the partnership in return for access to NWSC resources. Either way, the growth of the geosciences-based partnership will improve overall partner access to a spectrum of resources and will promote collaboration, reduce duplication of effort, and foster interdisciplinary research projects where none existed before. In particular, through partnership, these beneficial effects can be extended to other EPSCoR states in the region.

Finally, according to Koomey (2007), the total U.S. electricity consumption by servers in the U.S. was estimated to have doubled between 2000 and 2005 to a level of 23 billion kWh. When electricity use for cooling and auxiliary equipment is included, the total for 2005 rose to 45 billion kWh: almost as much energy “overhead” as the servers themselves required. NCAR and Wyoming are vigorously pursuing U.S. Green Building Council (USGBC) Leadership in Energy and Environmental Design (LEED) certification for the NWSC facility. By deploying an innovative, energy-efficient design, the facility represents an engineering advancement that, if replicated throughout the country, could reduce the energy required to operate server technology and increase the competitiveness of U.S. industry. Leading in “green data center design” is also consistent with NCAR’s role as a leading environmental institution.

6. Acknowledgments

The authors want to acknowledge and extend our deepest thanks to the dozens of scientists across many disciplines and at various institutions for their help in articulating the science need for petascale computing resources dedicated to understanding the Earth System. The synopses they generated for this document illustrate the breadth and depth of the issues involved. Their willingness to donate their time in workshops and in writing reports to describe the variety of challenges give this document its strength. Their support is invaluable in communicating the potential that petascale computing holds for advancing science and benefiting society.

We especially want to acknowledge the National Science Foundation for their support and guidance throughout this process. In particular, the NSF supported a series of workshops and reports that advanced our technical progress toward the vision of a petascale scientific computing facility. Specifically, funding for the GARPA and GARPA-2 workshops was provided under NSF-ATM award 0540688 received through subcontract 2005-006559-01 with the University of Illinois; funding for the HPC in Geosciences Workshop was provided under NSF award ATM-0631272; and the Petascale Computing for the Geosciences workshops were supported by NSF-GEO award 0621611.

Appendix A. NCAR Background

The National Center for Atmospheric Research (NCAR) is an NSF Federally Funded Research and Development Center (FFRDC) managed by the University Corporation for Atmospheric Research (UCAR) on behalf of a consortium of 73 research universities in North America. NCAR is already successfully serving the atmospheric and related sciences community by providing supercomputing and data services, including support of computational campaigns and on-demand computing, optimized queue structures with low wait times, a reliable and robust data archive, data curation, data portal services, geosciences-specific data analysis and visualization software, and consulting and training services. NCAR's research portfolio is balanced: it is strongly focused on climate research, yet it also supports a broad cross-section of Earth System science disciplines. NCAR also works in partnership with the broader atmospheric and related Earth sciences to develop community models, modeling infrastructure, and tools in support of the community's research objectives. The following sections expand on these points.

A.1 Track Record

NCAR has provided a discipline-specific computational environment for the atmospheric and related sciences for the past 40 years. The high quality of service provided by NCAR's Computational and Information Systems Laboratory (CISL) has been repeatedly noted by outside panels (Semtner et al., 2009, Leovoy et al., 2002, Dunning et al., 2001, Semtner et al., 2006, Good et al., 2006). NCAR's supercomputing resources currently support over 1,300 users across 700 projects. The demographics of NCAR's supercomputing user community are shown in Figures A.1 and A.2 below.

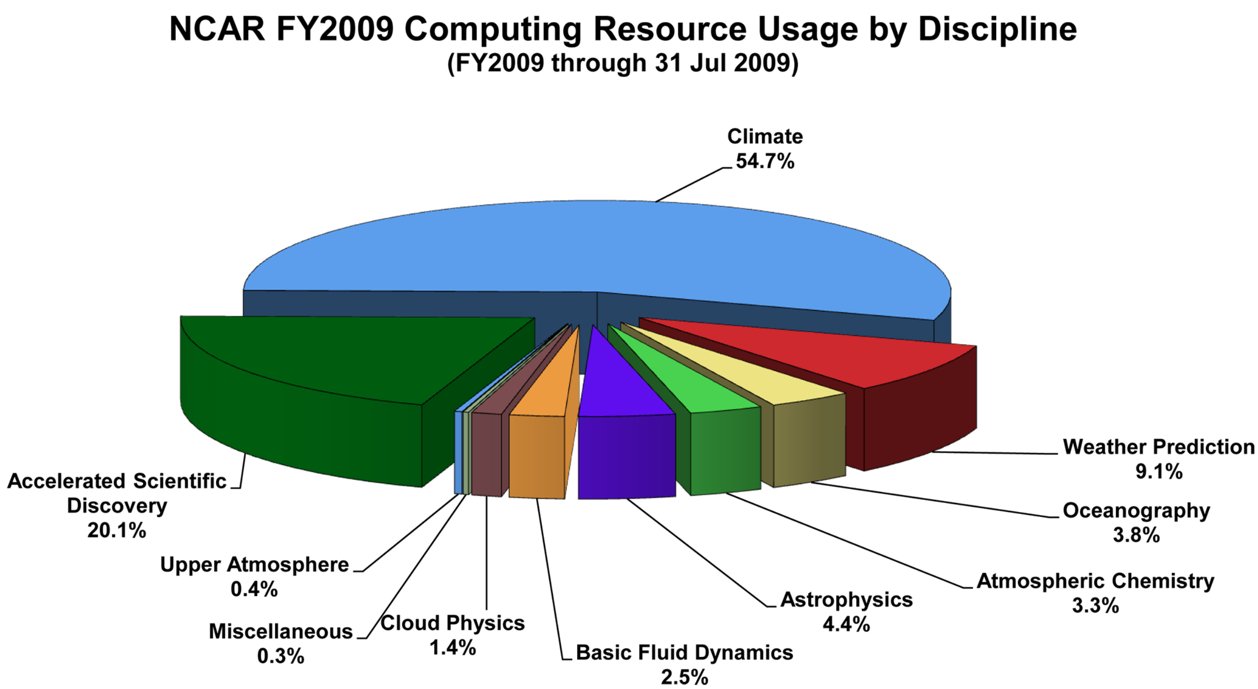

Figure A.1. The 2009 breakdown of resource use at NCAR's current computing facility. The next biggest consumer of resources after climate (54.7%) is computational campaigns for Accelerated Scientific Discovery (20.1%). Weather prediction consumed 9.1%, with all other disciplines combined using less than 20%.

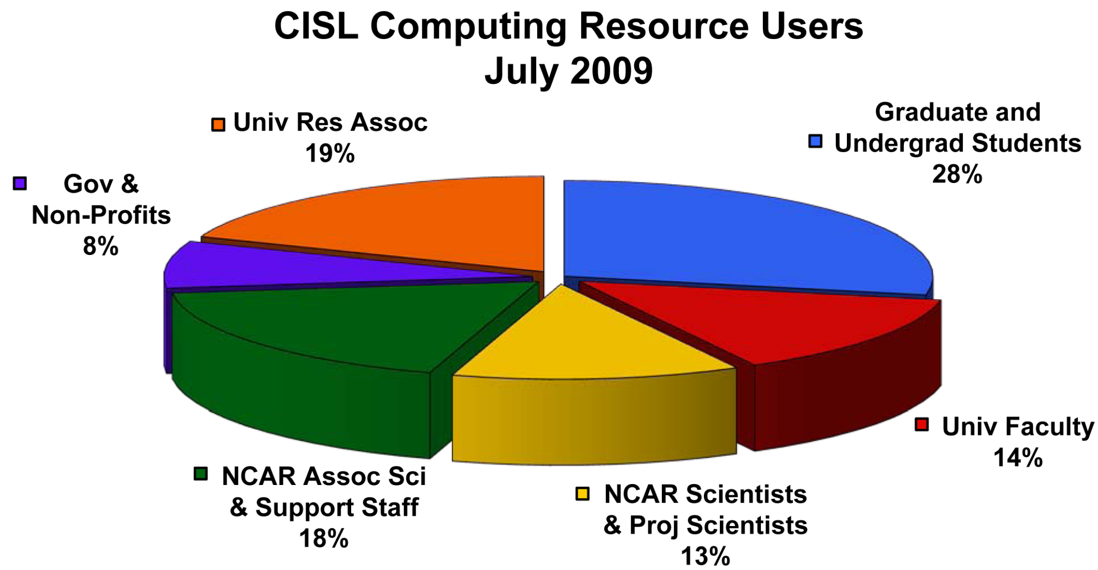

Figure A.2. The breakdown of NCAR supercomputing users by professional category. Over 60% are at universities and 28% are students, highlighting NCAR's valuable role in education.

A.2 Supercomputing Services

With its job scheduling policies, NCAR's CISL provides fair access, maximizes overall system utilization, and minimizes queue wait times. To satisfy these requirements, CISL has developed a priority-based, fair-share job-scheduling algorithm that allows users to specify the importance of their work (priority) while assuring that the entire user community is provided access to system resources (fair-share). As a result, NCAR's high-end production systems routinely deliver in excess of 90% utilization around the clock (see Figure A.3), while maintaining average queue wait times of less than 35 minutes on its IBM P-Series supercomputer Bluefire. Provision is also made for allocating significant portions of the resources for long-running jobs, as in the Breakthrough Science and Accelerated Scientific Discovery campaigns described below, and for supporting dedicated, on-demand and near-real-time field and computational campaigns.

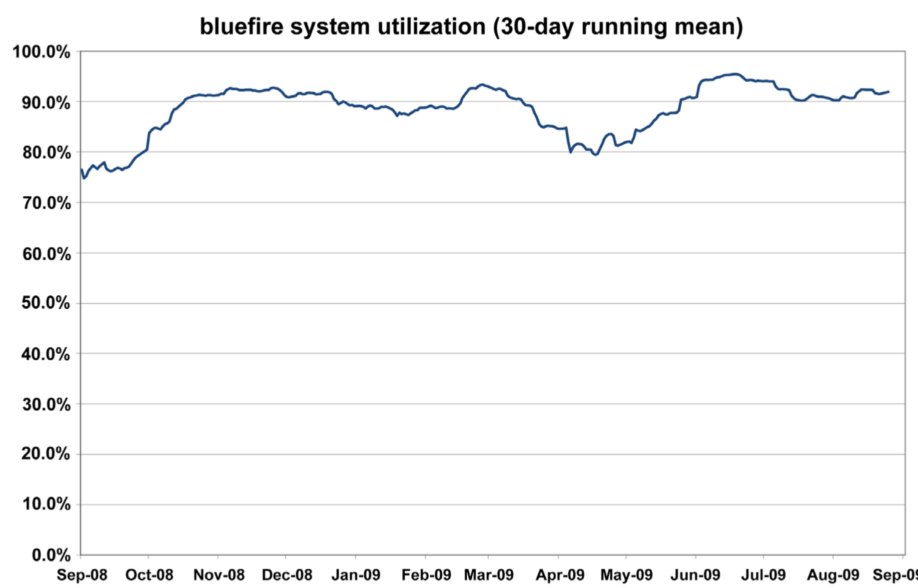

Figure A.3. Thirty-day running average utilization of NCAR's 77-teraflops Bluefire supercomputer over the past 12 months. After a quick ramp-up during the fall of 2008, utilization routinely exceeds 90%.

A.3 Data Services

The NCAR Mass Storage System (MSS) has served the atmospheric sciences community as a high performance, reliable, and accessible archive system for over two decades. The scalability of NCAR's MSS is noteworthy: during 20 years of continuous operation, the number of bytes stored has grown by a factor of 750, maintaining an annual growth rate of about 34%. Figure A.4 shows the MSS growth rate over the past 12 years, both in terms of unique and total (includes duplicate copies) bytes. At the end of FY2008, it surpassed 6 petabytes (PB) of total data (4 PB of unique data), transfers more than 10 terabytes (TB) of data per day back and forth in response to user requests, and transfers another 10 TB of data per day for internal data migration and data movement to new media. Including file duplications for preservation reasons, the MSS is currently growing at a net rate of more than 177 TB per month.

In 2008 CISL began deployment of an IBM High Performance Storage System (HPSS) to provide mass storage for its TeraGrid resource and to test the HPSS as a future CISL data archive solution for the NWSC. HPSS has been deployed by a number of the TeraGrid Resource Providers; thus, the deployment of HPSS at NCAR represents another level of integration with the TeraGrid.

In addition, NCAR provides robust web and grid-based data services. For example, the Research Data Archive (RDA) is a key part of the NCAR strategic priority to build and maintain a large collection of well-curated observational datasets that support scientific studies in

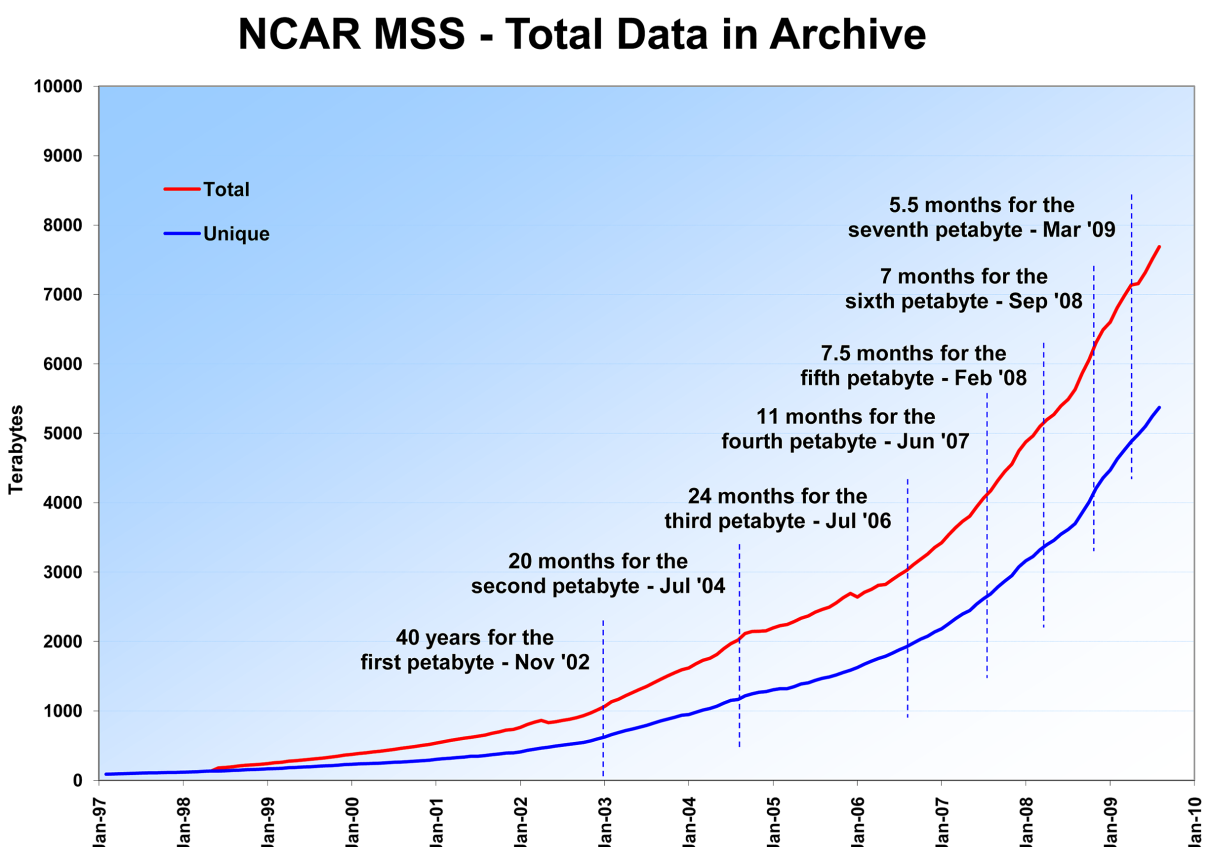

Figure A.4. NCAR's multi-petabyte mass storage system growth over the past 12 years. Unique data is shown in blue, total (including duplicates) is shown in red.

climate, weather, and increasingly, other related disciplines. The RDA was developed to serve the research needs at NCAR and in the associated university community, but since it is an open resource, the worldwide community also uses it. In FY2008, about 6,300 unique persons were provided 110 TB of RDA and related data.

NCAR is also active in deploying data portals for the scientific community. For example, NCAR is a partner in the DOE Office of Science-funded Earth System Grid (ESG). ESG's goal is to use Grid technology to facilitate the analysis of and extraction of knowledge from the O(100s) of terabytes of data produced by global climate models.

A.4 Leadership in Climate Research

One of NCAR's most important missions is supporting climate research. NCAR is arguably one of the pre-eminent climate research centers in the world. The Climate Change Science Program (CCSP) designated NCAR as one of two national modeling centers (Princeton's Geophysical Fluid Dynamics Laboratory – GFDL – is the other) "to develop, evaluate, maintain, and apply models capable of executing the most sophisticated simulations, such as those required for assessments by the [United Nations sponsored] Intergovernmental Panel on Climate Change (IPCC)" (Abraham et al., 2003). There are only a small number of institutions (16 groups from 11 countries participated in the IPCC AR4) with the resources necessary to do high-end climate change research and model development.

A.5 Community Modeling

NCAR supports the development of community models such as the Community Climate System Model (CCSM), the Weather Research and Forecasting (WRF) model, the Whole-Atmosphere Community Climate Model (WACCM), and the Nested Regional Climate Model (NRCM). These models are key enablers of NSF research in the university community.

NCAR also has a similar role in the U.S. Weather Research Program. Thus, NCAR's community models already function as a virtual organization for the Earth System sciences, bringing together a critical intellectual mass at universities and at government laboratories to focus on further developing and improving these models.

Through NCAR, climate researchers have timely access to CISL's stable computational environment in which to make model improvements, run and interpret scientific experiments, and perform long validation runs to check hundreds of code diagnostics and verify model performance.

A.6 Computational Campaign Support

CISL has a demonstrated track record of supporting computational campaigns such as a significant portion of the CCSM's contribution to the last IPCC Assessment. In addition, during FY2007 and FY2008-09, NCAR and NSF redirected significant blocks of time on newly installed supercomputers with the goals of achieving scientific breakthroughs while also broadening the interdisciplinarity of the NCAR user base. In FY2007 for instance, before NCAR's then-newest and most capable supercomputer (Blueice) was released for production use, NSF program managers in OCE and EAR, working with the CISL HPC Advisory Panel (CHAP), invited a small number of researchers to have exclusive use of that system for four months. These researchers were chosen based on their successful scientific track record, their ability to use large amounts of computing resources, and their potential to make significant discoveries through simulation. Six of the eight BTC projects successfully used their large allocations, consuming almost 3 million processor hours.

This practice was repeated in fall 2008 with the Advanced Scientific Discovery campaign using the newest supercomputer, Bluefire.

Resource usage by these Breakthrough Computing and Advanced Scientific Discovery campaigns during FY2007 and FY2008-09 are shown in Figures A.5 and A.6 below.

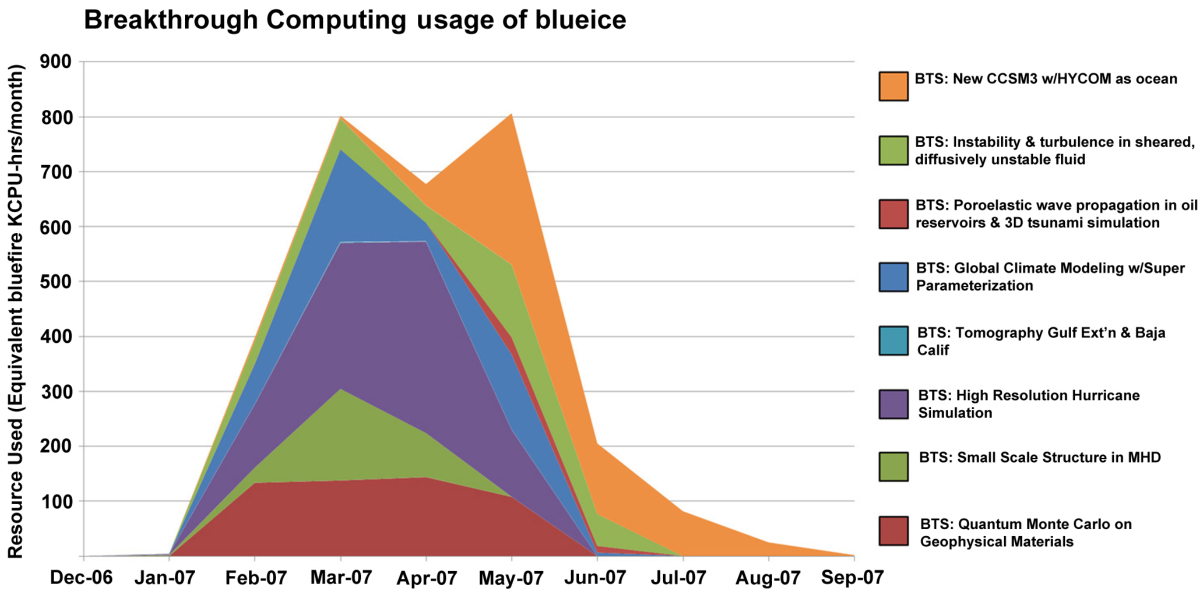

Figure A.5. Computer resource usage during the first scientific breakthrough computing campaign in 2007 in the first months after the IBM POWER5+ system Blueice was installed. Supercomputer systems are generally underutilized during their first few months of deployment. This program allowed these cycles to be maximized for scientific advancement.

Figure A.6. Computer resource usage during the second scientific breakthrough computing campaign (called Accelerated Scientific Discovery) in 2008 in the first months after the IBM POWER6 system Bluefire was installed.

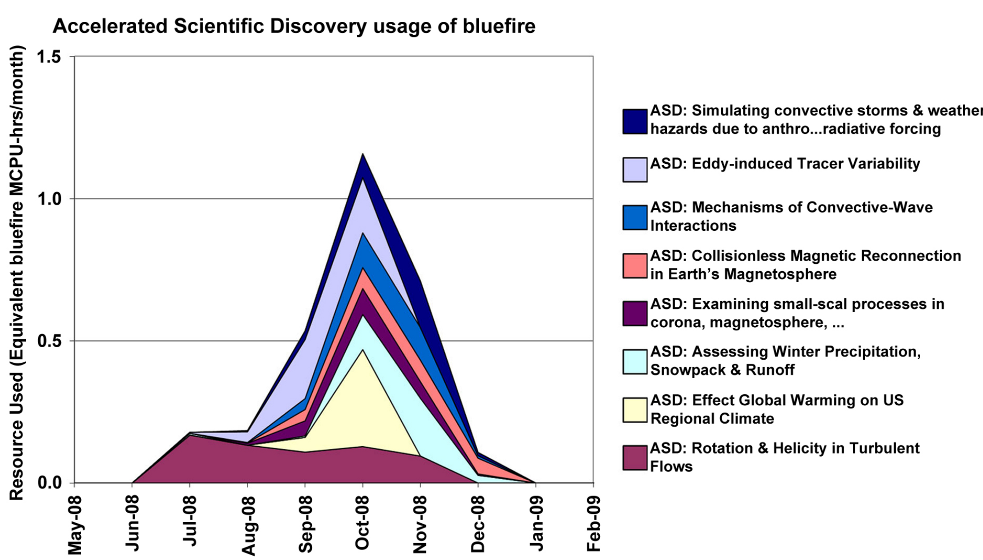

The IPCC AR4 Experience

Development, production, and data distribution of the NCAR Community Climate System Model results for the IPCC made full use of NCAR's dedicated supercomputing and data support infrastructure. NCAR/CISL preproduction risk reduction exercises that addressed a spectrum of potential threat scenarios allowed the IPCC production runs at NCAR to continue uninterrupted while a major international hacker attack was shutting down other supercomputer centers. The end result was that the dedicated CISL computing/data infrastructure-enabled NCAR climate model results to become one of the cornerstones of the Nobel prize-winning IPCC AR4 report that has so dramatically changed the US's perception and response to climate change in the past two years.

For more information on the close CISL/CGD teamwork that went into the IPCC AR4 effort, see http://www.cisl.ucar.edu/news/07/1207.nobel.jsp .

A.7 On-demand Computing Support

NCAR also has a history of providing seasonal, on-demand computing in support of meteorological research and field campaigns. One of NCAR's strengths is its ability to reserve significant computing resources for on-demand computing campaigns while continuing to provide uninterrupted production computing to the community. NCAR has developed an effective system for seamlessly allocating a large fraction of its resources to on-demand campaigns such as observational field campaigns, workshops, and experimental seasonal severe-weather forecasting campaigns for the meteorological research and operational communities. These experimental forecasts help transition research innovations to operational forecast systems. In addition, they are used (for example by meteorologists at the National Centers for Environmental Prediction) to improve ensemble forecasting. Figure A.7 shows the resources (normalized to Bluefire CPU-hours) allocated over the last five years to provide specific on-demand services. Below we summarize significant on-demand computations conducted at NCAR from 2004 to 2009.

Spring severe weather forecasts. These annual seasonal forecasts continue to evaluate and refine new severe weather forecasting applications while testing recent physics and numerics upgrades to WRF-ARW, as a foundation for NCAR collaborations with the National Severe Storms Lab (NSSL) and the Storm Prediction Center (SPC), through their springtime Hazardous Weather Testbed (HWT) Experiment. They have also been regularly used and evaluated by NWS forecasters and university faculty and students across the country. In 2009, these forecasts continue to be a core component of the HWT experiment and will also support the VORTEX-II field experiment.

Real-time hurricane forecasts. Near-real-time forecasts of Atlantic tropical cyclones are performed annually during the hurricane season. In 2009, these forecasts are done twice daily for the Real-Time Demonstration Project coordinated by the NOAA Hurricane Forecast Improvement Project (HFIP). These forecasts then become part of a multi-member ensemble that will be distributed to the National Hurricane Center.

DTC winter experiment. Resources were provided to support the real-time forecast experiment in the area covered by BAMEX (Bow Echo and Mesoscale convective vortices EXperiment). Researchers work with operational forecasters to determine if the value of high-resolution forecasts is becoming time critical.

WRF-MIRAGE. NCAR organized a large, intensive field campaign to study the chemical and physical transformations of pollutants in the outflow of the world's second largest metropolitan area, Mexico City (MIRAGE-Mex: Megacity Impacts on Regional and Global Environments – Mexico City case study). In preparation for the field program, NCAR conducted a real-time forecast experiment during February and March 2005, a full year in advance of the campaign, to identify an optimal configuration of the forecast system and forecast products in cooperation with MIRAGE investigators both inside and outside NCAR.

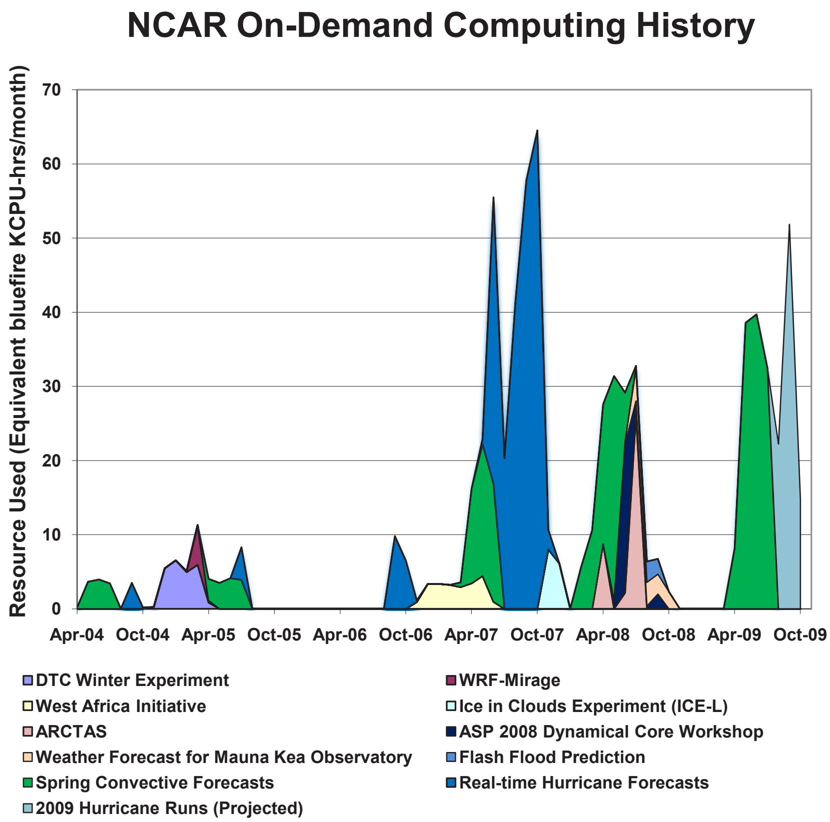

Figure A.7. On-demand resources allocated to a variety of special projects. The resources allocated for on-demand computing appear smaller in the past, as they are normalized to equivalent CPU-hrs of Bluefire, the primary computing system at NCAR.

West Africa Initiative. One component of NCAR's "African Initiative" is the establishment of an operational WRF-based forecasting system for Africa in general, and West Africa in particular, where the forecasters from African nations will use the products through a web interface. While waiting to procure, set up, and stabilize a PC cluster to run these forecasts, NCAR used CISL resources to produce daily forecasts.

Ice in Clouds Experiment. NCAR computing resources supported a field project called ICE-L (Ice in Clouds Experiment – Layer Clouds) that used the NCAR C-130 aircraft and associated instruments to: 1) establish which heterogeneous ice nucleation modes are active and important in clouds; 2) identify ice nucleating aerosols and obtain quantitative measurements; and 3) predict ice concentrations in clouds using numerical models in layer and wave clouds.

ASP 2008 Dynamical Core Workshop. Computing resources were dedicated to support a workshop designed to introduce a multidisciplinary group of graduate students to the latest developments in weather and climate modeling, with special attention to the dynamics component of global atmospheric models, the so-called dynamical core (dyCore). In particular, the colloquium surveyed and evaluated candidates for dyCores that are now being considered for next-generation Earth System models in the U.S. and worldwide. CISL is supporting three workshops for graduate student training in the summer of 2009: the CAM tutorial workshop (organized by NCAR), the Petascale Workshop (organized by NCSA and other TeraGrid sites), and the Summer Ice Sheet Modeling Workshop (organized by J. Johnson, University of Montana). All three of these workshops are sponsored by NSF and will receive dedicated nodes on Bluefire to enhance the students' learning experience.

Weather forecast for Mauna Kea Observatory. The Mauna Kea Weather Center (MKWC) provides custom weather forecasts for the community of astronomers on Mauna Kea in Hawaii. When MKWC no longer had access to computing resources from the National Astronomical Observatory of Japan and the Maui High Performance Computing Center, NCAR provided on-demand resources and helped set up interim operational WRF-based forecast runs for MKWC until they were able to deploy their own computing resources for that purpose.

Flash flood prediction. A research evaluation of an operational flash flood prediction system was performed using a 36-hour 1-km-resolution WRF run. These forecasts are the backbone of 24-36-hour forecast activities and were proven highly useful when there is a significant potential for significant flood-producing rainfall events.

ARCTAS – Arctic Research of the Composition of the Troposphere from Aircraft and Satellites. Computational resources were provided for daily chemical forecasts for ARCTAS, a community field program in which NCAR participated. This experiment focused on wildfires in Canada and Alaska, and provided predictions of distributions of carbon monoxide, aerosols, and ozone.

A.8 Education, Outreach, and Training

NCAR's education, outreach, and training activities are key strategic elements the NWSC proposal:

ASP. NCAR's Advanced Study Program (ASP) supports the scientific community by encouraging the development of early-career scientists in the Earth System sciences, directing attention to timely scientific areas needing special emphasis, helping organize new science initiatives, supporting interactions with universities, and promoting continuing education at NCAR. ASP operates an interdisciplinary Postdoctoral Fellowship Program, a Faculty Fellowship Program, and a Graduate Student Visitor Program. Focused on education, ASP also organizes and hosts an active lecture series, seminars, workshops, and summer colloquia.

TOY. The Institute for Mathematics Applied to the Geosciences (IMAGe) at NCAR sponsors a Theme of the Year (TOY) program to focus on specific areas of research that will benefit from intense collaborative effort. The topics are selected by the IMAGe external advisory panel and are coordinated by a Visiting Co-director.

SIParCS. The Summer Internships in Parallel Computational Science (SIParCS) program in CISL offers graduate students and undergraduate students (who have completed their sophomore year) significant hands-on R&D opportunities in high performance computing (HPC) and related fields that use HPC for scientific discovery and modeling. This program embeds students as summer interns in the Computational and Information Systems Laboratory (CISL), an organization within NCAR charged with provisioning supercomputing and data systems to the geosciences research community, as well as conducting research and development in computational science, data analysis, scientific visualization, and numerical modeling. These twin roles of service and research in CISL support NCAR's broad scientific mission of discovery in the atmospheric and related sciences.

NCL Atmospheric Data Analysis and Visualization Workshops. CISL offers three-to-four-day NCAR Command Language workshops on data analysis and visualization. These are taught three to five times per year, both on and off-site. Previous training venues have included international locations such as Turkey and Germany. Lectures given by NCAR scientists and software engineers familiar with data analysis and visualization and include hands-on labs with personal consulting. It is important to note that NCAR's community contribution in this area is not simply training to use these tools: it is fundamental training in the important practice of Earth System data analysis.

RSVP. CISL regards collaborations with university faculty, practitioners in high-performance computing, new researchers, and students as central to its mission to support the atmospheric and related sciences. CISL's Research and Supercomputing Visitor Program (RSVP) emphasizes visits of several weeks to several months to allow for significant collaboration with CISL staff and seeks to provide partial or possibly full financial support for travel and living expenses related to the visit.

SOARS. Significant Opportunities in Atmospheric Research and Science (SOARS) is an interdisciplinary undergraduate-to-graduate bridge program designed to broaden participation in the atmospheric and related sciences. The program is built around research, mentoring, and community. SOARS participants, called protégés, spend up to four summers doing research in atmospheric and related sciences. SOARS offers comprehensive financial support for summer research, as well as undergraduate and graduate school funding. Over 90% of SOARS protégés have entered graduate school, and many have entered the workforce with an MS or attained their Ph.D.

A.9 Integration with the TeraGrid

NCAR is currently a resource provider on the TeraGrid and has been providing allocable cycles from its MRI-funded Blue Gene/L system since 1 August 2007. NCAR has also provisioned data analysis and visualization resources on its TeraGrid node and has recently placed a 1 petabyte HPSS archive in service. NCAR plans to expand this Blue Gene/L resource four-fold in FY2009 and will continue to provide computing resources at least until the end of the current TeraGrid project in April 2011. NCAR is therefore already demonstrating the experience and in-house expertise to connect a Track-2-scale Earth System science resource to TeraGrid XD systems. These capabilities afford atmospheric scientists opportunities to form strategically important interdisciplinary collaborations via the TeraGrid.

A.10 Integration with Investments in Other Sectors

NCAR's role as community integrator and its leadership responsibilities as an NSF FFRDC means that NCAR has both the experience and expertise to ensure that the facility will be well integrated with national and international HPCI. In particular, NCAR/UCAR has been a leader in developing regional and national-scale networking consortia for a number of years (FRGP, 2009, NLR, 2009). For example, NCAR/UCAR provides the engineering and Network Operations Center support for the Front Range GigaPOP. NCAR has also led in the development of national-scale grid capabilities (Catlett et al., 2007, Bernholdt et al., 2005, Dunning et al., 2001). NCAR/UCAR is actively engaged in international information system integration efforts (Ramamurthy et al., 2007, CDP, 2009): NCAR is collaborating in pilot project activities such as the WMO Information System (WIS) THORPEX Interactive Grand Global Ensemble (TIGGE) project, serving three TIGGE archive databases worldwide, along with data from the Chinese Meteorological Agency (CMA) and the European Centre for Medium-range Weather Forecasts (ECMWF) (TIGGE, 2009).

A.11 Furthering Interagency Collaboration

Since 1995, NCAR has hosted the Climate Simulation Laboratory (CSL), a multi-agency computational resource in support of the U.S. Global Change Research Program (formerly the U.S. Climate Change Science Program). The purpose of the CSL is to provide high performance computing, data storage and supporting infrastructure for large-scale, long-running simulations of Earth's climate system on time scales of seasons to centuries. The CSL provided significant computational resources for the IPCC Fourth Assessment and will support the upcoming IPCC Fifth Assessment. The CSL provides a solid basis for fostering additional interagency collaboration in the Earth System sciences at the new facility.

Appendix B. Wyoming Background

B.1 State of Wyoming

Wyoming has several unique attributes that influence its economy, education, and technology and that guide the research and educational initiatives proposed here. Wyoming has the smallest population (<525,000) and the 2nd-lowest population density among the 50 states. Consequently, a single research institution, the University of Wyoming (UW) is supported by the state. The higher education system consists of UW with 13,000 students, seven community colleges with 23,000 students (66% part-time), and a single tribal college with 46 students. Each of the community colleges and the University are governed separately. With foresight, the state legislature and Governor Dave Freudenthal has created a permanent trust fund to provide full-tuition "Hathaway Scholarships" for students to attend UW or the community colleges if the state's mandated curriculum, grades, and standardized test scores are met. This led to more than a 10% increase in the incoming freshman classes at UW. These students come from highly diverse high school backgrounds: of the 70 high schools in Wyoming, 33% have <100 students and only 10% have over 1,000. The State of Wyoming is particularly focused on knitting together this dispersed population by building distance learning and bringing science education to students and the public through all useful media.

Wyoming's economy is dominated by extractive industries, with natural gas, coal, oil, and soda ash extraction accounting for over 70% of state revenues. The few technology businesses in Wyoming are in environmental engineering, materials, and software companies. Efforts to diversify the economy are centered in the Wyoming Business Council, a quasi-private-government entity, which was created to combine a variety of economic development groups within state government. Interestingly, Wyoming is one of the most connected states in terms of the percent of the population connected to the Internet. In this proposal, economic diversification efforts are designed around new research technologies associated with climate change research and computational power associated with our cyberinfrastructure plan and with outreach and education programs to start-up companies. Also, southern Wyoming has one of the highest concentrations of fiber networks found in the West; numerous carriers have fiber lines running across southern Wyoming. This is a key resource in state economic development plans as they relate to industries requiring a high amount of reliable bandwidth – a fact contributing to their selection by NCAR as a site for the new supercomputing facility.

An important enabling partner in the construction and operation of the NWSC, the State of Wyoming has committed substantial financial resources to the project. As an EPSCoR[23] state, Wyoming has a particularly important role to play in facilitating the NCAR-Wyoming partnership's interaction with other EPSCoR jurisdictions.

B.2 The University of Wyoming

The University of Wyoming (UW) a public land-grant institution founded in 1886 and located in Laramie, Wyoming. UW is a national research university prominent in the fields of environment and natural resource research, specializing in agriculture, energy, geology, and water resource-related fields. In 2008, approximately 13,000 students were enrolled at UW.

23 The Experimental Program to Stimulate Competitive Research (ESPCoR) is a federal program to strengthen research and education in science and engineering throughout the United States. The term "EPSCoR states" refers to those states that have been determined to have historically received insufficient research and education funding.

The University of Wyoming consists of seven colleges: agriculture, arts and sciences, business, education, engineering, health sciences, and law. The university maintains a low student-faculty ratio – one of the lowest such ratios among four-year schools in the West. UW also offers a variety of cultural and social activities. The university offers 86 bachelor's, 66 master's, and 26 doctoral degrees.

The university has a longstanding strategic commitment to investing in the computational sciences and the environmental and natural resource sciences. Building major strength in computational sciences in these areas has been major objectives in the past two UW five-year academic plans, the most recent of which will end in 2009. Academic Plan II was written prior to the formation of the NWSC and the deepened relationship between NCAR and UW, therefore direct references to NCAR will not be found. But the new plan – *Creation of the Future: University Plan III* (the university plan for 2009-2014) – takes full advantage of this developing partnership. UW's goal is to become one of the nation's leading universities in these areas. Specifically, the plan for 2009-2014 states:

- "The University will build an undergraduate program in Earth System sciences expanding on the expertise in Geology & Geophysics, Atmospheric Sciences, Geography, Renewable Resources, Chemistry, Botany, Zoology and Natural Resource Economics. The program will ... have clear connections to research in areas concerning humans in the biosphere, ecology, climatology, biogeoscience as well as the more traditional earth and atmospheric sciences."

- "For the next five years the Office of Academic Affairs will earmark at least one position per year in computational fluid mechanics, computational geosciences or computational-related supporting fields through central position management. Coherent interdisciplinarity and curricular impact will be key factors in the allocations."

Pursuing this strategy, UW has hired 12 new tenure-track faculty members expert in some field of computational science over the past 10 years. More specifically, with these hires UW has built a strong group specializing in computational geosciences and particularly in multiphase fluid flow in porous media – or the movement of water, natural gas, petroleum, and carbon dioxide within sandstone and carbonaceous reservoirs at depth. Thus, UW is at the forefront of research in numerical reservoir simulation, with considerable faculty expertise in the area of porous media flow physics, physical chemistry, geological modeling, and high performance computing, spread over numerous academic departments including Mathematics, Geology and Geophysics, Chemistry, and Chemical and Petroleum Engineering. The recent formation of the School of Energy Resources and the participation of the University in DOE-NETL-funded research on geologic carbon sequestration, have brought additional campus-wide emphasis to these activities. The University will continue even more strongly on its path to develop the computational Earth System sciences and the computational geosciences.

Creation of the Future: University Plan III contains strategic plans for further enhancing computational science at UW, making UW's campus high-performance computing infrastructure operations more coherent, its CI interoperable with the NCAR-UW supercomputer facility, and moving to provide co-curricular educational opportunities for UW students.

UW will help translate the technological and educational impacts of the NWSC into benefits for Wyoming, the region, and especially other EPSCoR states. Specifically, over the past 10 years UW has built a strong group specializing in computational geosciences and particularly in multi-phase fluid flow in porous media – or the movement of water, natural gas, petroleum, and carbon dioxide within sandstone and carbonaceous reservoirs at depth. Now UW has one of the nation's strongest groups in this area. By partnering with the University of Wyoming, NCAR and the NWSC will bring together energy researchers, climate researchers, and experts in modeling carbon sequestration as a climate change mitigation strategy under the same research umbrella.

The relationship between the University of Wyoming and NCAR has existed for many years. UW was one of the first universities to join UCAR and has enjoyed a collaborative relationship with NCAR through the deployment of the Wyoming King Air Atmospheric Science research aircraft, an NSF-funded research facility owned by the university. NCAR assists UW in some mission deployments, NCAR scientists use the Wyoming King Air for their research, and NCAR shares equipment that can be used in this and other research aircraft. More recently, UW partnered with NCAR, Colorado State University, and the University of Colorado at Denver to build the Front Range GigaPOP (FRGP) as Internet-2 was being developed. Now this important point of presence supports Internet-2, Lambda Rail, and TeraGrid. It also includes many more research entities, including universities and federal laboratories along the Front Range. These examples underscore the fact that a UW-NCAR relationship is not new, but that the NWSC initiative is another collaboration between two entities that have been partners for decades.

Appendix C. References

Ad Hoc Committee for Cyberinfrastructure Research, Development and Education in the Atmospheric Sciences, 2004: *Cyberinfrastructure for Atmospheric Sciences in the 21st Century*, 54 pp. Available online at http://gladiator.ncsa.uiuc.edu/PDFs/nsf/cyrdas_report_final.pdf .

Andrejczuk, M., W.W. Grabowski, S.P. Malinowski, and P.K. Smolarkiewicz, 2004: Numerical simulation of cloud-clear air interfacial mixing. *J. Atmos. Sci.*, **61**, 1726-1739.

Anthes, R.A., (co-chair with B. Moore), et al., 2007: *Earth Science and Applications from Space: National Imperatives for the Next Decade and Beyond.* Report of the Space Studies Board, National Research Council of the National Academies, National Academy Press, Washington, D.C., 456 pp.

Atkins, Dan (chair), et al., 2003: *Revolutionizing Science and Engineering Through Cyberinfrastructure*: Report of the National Science Foundation Blue Ribbon Panel on Cyberinfrastructure. Available online at http://www.communitytechnology.org/nsf_ci_report/report.pdf .

Ayala, O., W.W. Grabowski, and L.-P. Wang, 2007: A hybrid approach for simulating turbulent collisions of hydrodynamically interacting particles. *J. Comp. Physics*, **225**, 51-73.

CCSP, 2008: *The effects of climate change on agriculture, land resources, water resources, and biodiversity:* A Report by the U.S. Climate Change Science Program and the Subcommittee on Global Change Research. Lead authors P. Backlund, A. Janetos, D. Schimel. U.S. Environmental Protection Agency, Washington, D.C., 362 pp.

Belady, Christan, PE, *In the Data Center, Power and Cooling Costs More Than the IT Equipment It Supports*, Hewlett-Packard Reprint from Electronics Cooling, February 2007. See Section 3.

Benioff, M.D., Lazowska, E.D., (co-chairs), et al., Presidential Information Technology Advisory Committee, 2005: *Computational Science: Ensuring America's Competitiveness*, Office of the President of the United States, Washington, D.C., 116 pp. Available online at http://www.nitrd.gov/pitac/reports/20050609_computational/computational.pdf .

Bernholdt, D., Bharathi, S., Brown, D., Chanchio, K., Chen, M., Chervenak, A., Cinquini, L., Drach, B., Foster, I., Fox, P., Garcia, J., Kesselman, C., Markel, R., Middleton, D., Nefedova, V., Pouchard, L., Shoshani, A., Sim, A., Strand, G., Williams, D., 2005: The Earth System Grid: Supporting the Next Generation of Climate Modeling Research. *Proceedings of the IEEE*, 93 (3), 485-495. Available online at http://ieeexplore.ieee.org/iel5/5/30407/01398005.pdf?arnumber=1398005 .

Bleck, R., et al., 2002: *An Information Technology Infrastructure Plan to Advance Ocean Sciences.* Office of Naval Research and the National Science Foundation, 80 pp. Available at https://www.paleostrat.org/Documents/oiti%202002.pdf .

Bower, A.S., M.S. Lozier, S.F. Gary, C.W. Böning, 2009: Interior pathways of the North Atlantic meridional overturning circulation. *Nature*, **459**, 243-247.

Bryan, F., (chair), et al., Ad Hoc Committee and Technical Working Group for a Petascale Collaboratory for the Geosciences, 2005: *Establishing a Petascale Collaboratory for the Geosciences: Scientific Frontiers. A Report to the Geosciences Community.* UCAR/JOSS, Boulder, Colo., 80 pp.

Bryan, G.H., J.C. Wyngaard, and J.M. Fritsch, 2003: Resolution requirements for the simulation of deep moist convection. *Mon. Wea. Rev.*, **131**, 2394-2416.

Butcher, J.C., 2003: *Numerical Methods for Ordinary Differential Equations.* Wiley, 440 pp.

Catlett, C., et al., 2007: TeraGrid: Analysis of Organization, System Architecture, and Middleware Enabling New Types of Applications. In Grandinetti, L., ed., *HPC and Grids in Action.* IOS Press, Amsterdam.

Clyne, J., and M. Rast, 2005: A prototype discovery environment for analyzing and visualizing terascale turbulent fluid flow simulations. *Proceedings of Visualization and Data Analysis 2005,* pp. 284-294.

Cockburn, B., G.E. Karniadakis, and C.W. Shu, 2000: *Discontinuous Galerkin Methods.* Springer, 470 pp.

Cohen, R.E., ed., 2005: *High-Performance Computing Requirements for the Computational Solid Earth Sciences.* Geophysical Laboratory, Carnegie Institute of Washington, 94 pp. Available online at http://www.geo-prose.com/pdfs/computational_SES.pdf .

Colella, P., Dunning, T.H., Gropp, W.D., and Keyes, D.E., eds., 2003: *A Science-Based Case for Large-Scale Simulation, Volume 1.* U.S. Department of Energy Office of Science, Washington, D.C., 70 pp.

Collins, W.D., et al., 2006: The Community Climate System Model Version 3 (CCSM3). *J. Climate,* **19** (11), 2122-2143.

Crawford, et al., 2007: *Cyberinfrastructure Vision for 21st Century Discovery***.** National Science Foundation, 24 pp.

Committee on Strategic Advice on the U.S. Climate Change Science Program, 2009: *Restructuring Federal Climate Research to Meet the Challenges of Climate Change.* National Research Council, 266 pp.

Davis, C., W. Wang, S. Chen, Y. Chen, K. Corbosiero, M. DeMaria, J. Dudhia, G. Holland, J. Klemp, J. Michalakes, H. Reeves, R. Rotunno, and Q. Xiao, 2008: Prediction of landfalling hurricanes with the advanced hurricane WRF model. *Mon Wea. Rev.,* **136**, 1990-2005.

D'Azevedo, E.F., V. Eijkhout, C.H. Romine, 1993: *A Matrix Framework for Conjugate Gradient Methods and Some Variants of CG with Less Synchronization Overhead.* PPSC 1993: 644-646. Available online at http://web.tacc.utexas.edu/~eijkhout/Articles/1993-SIAM-PPSC.pdf .

Dennis, J.M., and H.M. Tufo, 2008: Scaling Climate Simulation Applications on IBM Blue Gene. *IBM Journal of Research and Development: Applications of Massively Parallel Systems,* **52** (1/2), 117-126.

Dennis, J.M., 2007: Inverse Space-Filling Curve Partitioning of a Global Ocean Model. *IEEE International Parallel & Distributed Processing Symposium,* Long Beach, Calif., March 26-30, 2007.

Drake, J.B., P.W. Jones, and G.R. Carr, Jr., 2005: Overview of the software design of the Community Climate System Model. *Int. J. High Perf. Comp. App.,* **19**, 177-186.

Data Support Section (DSS) within NCAR's Computational and Information Systems Laboratory (CISL), cited 2007. Available online at http://dss.ucar.edu/ .

Dunning, T., (chair), et al., 2001: The NSF's Panel Review of the Scientific Computing Division, November 8, 2001. See p. 5: "The results of a survey of the users of SCD computers indicates a high level of satisfaction with the high performance services provided by SCD." See p. 8: "Overall, SCD appears to be doing a superlative job providing mass storage services to the atmospheric sciences community." See p. 9: "The MSS continues to be a jewel in the crown for computing at NCAR/SCD, providing capabilities that are difficult to find at other supercomputing centers." See p. 10: "SCD continues to be a world-class leader in the stewardship of atmospheric data and climate observation data—the Panel knows of no other institution that currently does this more effectively." See p. 10: "SCD is to be commended for taking a strong leadership role in the development of the data grid for the Earth Systems community."

Evershed, J., 1909: Radial movement in sun-spots. *Monthly Notices of the Royal Astronomical Society*, **69**, 454-457.

Fan, Y., 2009: The Emergence of a Twisted Flux Tube into the Solar Atmosphere: Sunspot Rotations and the Formation of a Coronal Flux Rope. *The Astrophysical Journal,* **697** (2), 1529-1542.

Fennel, K.F., J. Wilkin, J. Levin, J. Moisan, J. O'Reilly, and D. Haidvogel, 2006: Nitrogen cycling in the Middle Atlantic Bight: Results from a three-dimensional model and implications for the North Atlantic nitrogen budget. *Global Biogeochemical Cycles,* **20**, GB3007, doi:10.1029/2005GB002456, 14 pp.

Ferreira, K.B., R. Brightwell, P. Bridges, 2008: Characterizing Application Sensitivity to OS Interference Using Kernel-Level Noise Injection. Conference paper, *International Conference on High Performance Computing, Networking, Storage, and Analysis (SC'08),* November 2008.

Fiedler, B., 2009: Suction vortices and spiral breakdown in numerical simulations of tornado-like vortices. *Atmos. Sci. Lett.*, DOI: 10.1002/asl.217.

Franklin, C.N., P.A. Vaillancourt, M.K. Yau, 2007: Statistics and parameterizations of the effects of turbulence on the geometric collision kernel of cloud droplets. *J. Atmos. Sci.*, **64**, 938-954.

Front Range GigaPOP (website). Available online at http://www.frgp.net/ .

Fulton, S.R., 2001: An adaptive multigrid barotropic tropical cyclone track model. *Mon. Wea. Rev.,* **129**, 138-151.

Good, M., (chair), et al., 2006: *Report of the NSF Panel to Review NCAR Management,* March 20-22, 2006. 12 pp. See p. 2: "NCAR has pursued successfully its computational and information systems upgrades, with the result that it has multiplied by at least 15 its supercomputing capacity over the last 4-5 years, a quite enviable result."

Grabowski, W.W., X. Wu, M.W. Moncrieff, and W.D. Hall, 1998: Cloud-resolving modeling of tropical cloud systems during Phase III of GATE. Part II: effects of resolution and the third spatial dimension. *J. Atmos. Sci.,* **55**, 3264-3282.

Graham, S.L., M. Snir, and C.A. Patterson, eds., 2004: *Getting up to Speed: The Future of Supercomputing.* National Research Council Computer Science and Telecommunications Board. National Acadamies Press, Washington, D.C., 308 pp. Available online at http://www.nap.edu/books/0309095026/html/ .

Haidvogel, D.B., H.G. Arango, K. Hedstrom, A. Beckmann, P. Malanotte-Rizzoli, and A.F. Shchepetkin, 2000: Model Evaluation Experiments in the North Atlantic Basin: Simulations in Nonlinear Terrain-Following Coordinates. *Dynamics of Atmospheres and Oceans,* **17** (32), 239-281. Available online at http://marine.rutgers.edu/po/Papers/damee_roms1.pdf .

Haidvogel et al., 2008: Ocean forecasting in terrain-following coordinates: Formulation and skill assessment of the Regional Ocean Modeling System. *J. Comp. Physics,* **227**, 3595-3624.

Heinemann, T.,Å. Nordlund, G.B. Scharmer, and H. C. Spruit, 2007: MHD Simulations of Penumbra Fine Structure. *The Astrophysical Journal*, **669** (2), 1390-1394.

Heinz, S., 2003: *Statistical Mechanics of Turbulent Flows.* Springer, 214 pp.

Hewitt, W.T., A.M. Jones, and Z. Chaplin, 2007: *Hadley Centre Review 2006: Supercomputing, Final Report,* University of Manchester, Manchester, UK, 62 pp.

Hiroyasu, H. and S. Emori, eds., K-1 Model Developers, 2004: *K-1 Coupled GCM (MIROC) Description.* Center for Climate System Research, Univeristy of Tokyo. Available at: http://www.ccsr.u-tokyo.ac.jp/kyosei/hasumi/MIROC/tech-repo.pdf .

Hoerling, M., and J. Eischeid, 2007: Past Peak Water in the Southwest. *Southwest Hydrology*, 6, 18-20.

Houze, R.A. Jr., and A.K. Betts, 1981: Convection in GATE. *Rev. Geophys. and Space Phys.*, **19**, 541-576.

Inderfurth, K.F., D. Fabrycky, and S. Cohen, 2005: *The 2004 Indian Ocean Tsunami: 6 Month Report.* The George Washington University, Washington, D.C., 27 pp.

Jablonowski, C., M. Herzog, J.E. Penner, R.C. Oehmke, Q.F. Stout, B. van Leer, and K.G. Powell, 2006: Block-structured adaptive grids on the sphere: advection experiments. *Mon. Wea. Rev.,* **134**, 3691-3713.

Johnson, D.L., et al., 2006: *Service Assessment: Hurricane Katrina, August 23-31, 2005.* U.S. Department of Commerce National Oceanic and Atmospheric Administration (NOAA) and National Weather Service (NWS), Silver Spring, Md., 48 pp. Available online at http://www.weather.gov/os/assessments/pdfs/Katrina.pdf .

Kain, J.S., S.J. Weiss, D.R. Bright, M.E. Baldwin, J.J. Levit, G.W. Carbin, C.S. Schwartz, M.L. Weisman, K.K. Droegemeier, D.B. Weber, and K.W. Thomas, 2008: Some practical considerations regarding horizontal resolution in the first generation of operational convection-allowing NWP. *Wea. and Forecasting,* **23**, 931–952.

Kalnay, E., 2003: *Atmospheric Modeling, Data Assimilation, and Predictability.* Cambridge University Press, Cambridge, UK, 341 pp.

Karl, T.R. and Meehl, G.A., (co-chairs), et al., 2008: *Weather and Climate Extremes in a Changing Climate. Regions of Focus: North America, Hawaii, Caribbean, and U.S. Pacific Islands.* U.S. Climate Change Science Program Synthesis and Assessment Product 3.3, June 2008.

Khairoutdinov, M., D.A. Randall, and C. DeMott, 2005: Simulations of the atmospheric general circulation using a cloud-resolving model as a superparameterization of physical processes. *J. Atmos. Sci.,* **62**, 2136-2154.

Kinter, J., and E. Seidel (co-chairs), et al., 2008: *Recommendations to ATM on Computing for Atmospheric Sciences in the Period 2011-2016.* Report to the NSF, 23 pp.

Kolmogorov, A.N., 1941: The local structure of turbulence in incompressible viscous fluid at very large Reynolds numbers. *Dokl. Akad. Nauk SSSR*, **30**, 301.

Koomey, J.C., 2007: *Estimating Total Power Consumption by Servers in the U.S. and the World.* Lawrence Berkeley National Laboratory, Berkeley, Calif., 27 pp.

Kunze, E., et al., 2006: Observations of Biologically Generated Turbulence in a Coastal Inlet. *Science,* **313** (5794), pp. 1768-1770.

Laursen, K., P. Backlund, L. Buja, P. Fox, R. Loft, M. Wiltberger, 2006: *Report on the High Performance Computing in the Geosciences Workshop.* NCAR, September 25-27, 2006, Boulder, Colo., 14 pp. Available online at http://www.ncar.ucar.edu/Director/dcworkshop/HPCGEO_workshop_report_FINAL.pdf .

Leovoy, C.B., (chair), et al., 2002: *Panel Report on the Review Proposal of the University Corporation for Atmospheric Research Management of the National Center for Atmospheric Research 2003-2008.* See Section H, Facilities. Also see: "One of the major contributions of NCAR to the national atmospheric sciences research effort has been the provision of atmospheric data sets to universities and other research and educational users."

Loft, R., (chair), et al., Technical Working Group and Ad Hoc Committee for a Petascale Collaboratory for the Geosciences, 2005: *Establishing a Petascale Collaboratory for the Geosciences: Technical and Budgetary Prospectus.* A Report to the Geosciences Community. UCAR/JOSS, Boulder, Colo., 56 pp.

Loft, R., et al., 2008: *Workshop Report: Petascale Computing in the Geosciences.* 43 pp.

Lyons, J.W. (chair) et al., 2008: *The Potential Impact of High-End Capability Computing on Four Illustrative Fields of Science and Engineering.* The National Research Council of the National Academies, National Academies Press, Washington, D.C., 142 pp.

Mahoney, J.R., (chair), et al., 2003: *A report by the Climate Change Science Program and the Subcommittee on Global Change Research.* See p. 108.

Mantua, N., D. Haidvogel, Y. Kushnir, N. Bond, 2001: Making the Climate Connections: Bridging scales of space and time in the U.S. GLOBEC program. *Oceanography*, **15** (2), 75-86. Available online at http://www.tos.org/oceanography/issues/issue_archive/issue_pdfs/15_2/15_2_mantua_et_al.pdf .

Masumoto, Y., H. Sasaki, T. Kagimoto, N. Komori, A. Ishida, Y. Sasai, T. Miyama, T. Motoi, H. Mitsudera, K. Takhashi, H. Sakuma, and T. Yamagata, 2004: A fifty-year simulation of the world ocean – Preliminary outcomes of OFES (OGCM for the Earth Simulator). *J. Earth Simulator,* **1**, 35-56.

McClean, J.L., M.E. Maltrud, and F.O. Bryan, 2006: Quantitative measures of the fidelity of eddy-resolving ocean models. *Oceanography*, **19**, 60-73.

Meehl, G.A., et al., 2006: Climate change projections for the twenty-first century and climate change commitment in the CCSM3. *J. Climate*, **19** (11), 2597-2616.

Majda, A.J., and B. Khouider, 2002: Stochastic and mesoscopic models for tropical convection. *Proc. Natl .Acad. Sci.,* **98** (4), 1341-1346.

Molteni, F., R. Buizza, T.N. Palmer, and T. Petroliagis, 1996: The ECMWF ensemble prediction system: methodology and validation. *Quart. J. Roy. Meteor. Soc.,* **122**, 73-119.

Michalakes, J., T. Canfield, R. Nanjundiah, S. Hammond and G. Grell, 1995: Parallel Implementation, Validation, and Performance of MM5. In *Coming of Age: Proceedings of the Sixth ECMWF Workshop on the Use of Parallel Processors in Meteorology*, World Scientific, River Edge, NJ, 266–276.

Michalakes, J., J. Hacker, R. Loft, M. McCracken, A. Snavely, N.J. Wright, T. Spelce, B. Gorda, R. Walkup, 2007: WRF Nature Run. *Proceedings of the 2007 ACM/IEEE conference on Supercomputing,* Conference on High Performance Networking and Computing, Gordon Bell Prize Finalists Session, Reno, Nev., 6 pp.

Miura, H., M. Satoh, H. Tomita, A.T. Noda, T. Nasuno, S. Iga, 2007: A short-duration global cloud-resolving simulation with a realistic land and sea distribution, *Geophys. Res. Lett.,* **34**, L02804, doi:10.1029/2006GL027448, 5 pp.

Moncrieff, M.W., 2004: Analytic representation of the large-scale organization of tropical convection. *J. Atmos. Sci.,* **61**, 1521-1538.

Moncrieff, M.W., M. Shapiro, J. Slingo, and F. Molteni, 2007: Collaborative research at the intersection of weather and climate. *WMO Bulletin,* **56**, 204-211.

Morrison, H., and W.W. Grabowski, 2009: Comparison of bulk and bin warm rain microphysics models using a kinematic framework. *J. Atmos. Sci.,* in press.

Nair, R.D., S.J. Thomas, and R.D. Loft, 2005: A discontinuous Galerkin scheme of the cubed sphere. *Mon. Wea. Rev.,* **133**, 814-828.

NLR, National LambdaRail (website), 2009. Available online at http://www.nlr.net/ .

Nocedal, J., and S. J. Wright, 1999: *Numerical Optimization.* Springer, 656 pp.

Odenwald, S., 2009a: Space Weather (website). Available online at http://www.solarstorms.org/SUtilities.html .

Odenwald, S., 2009b: Space Weather (website). Available online at http://www.solarstorms.org/Spower.html .

Palmer, T.N., 2001: A nonlinear dynamical perspective on model error: a proposal for non-local stochastic-dynamic parameterization in weather and climate prediction models. *Quart. J. Roy. Meteor. Soc.,* **127**, 279-304.

Parry, M.L., O.F. Canziani, J.P. Palutikof, P.J. van der Linden, and C.E. Hanson, eds., 2007: *Climate Change 2007: Impacts, Adaptation and Vulnerability: Working Group II Contribution to the Fourth Assessment Report of the IPCC.* Cambridge University Press, Cambridge, UK, 976 pp.

Peixoto, J.P., and A.H. Oort, 1992: *Physics of Climate.* American Institute of Physics, New York, 564 pp.

Pope, S.B., 2000: *Turbulent Flows.* Cambridge University Press, Cambridge, UK, 771 pp.

Ramamurthy, M.K., T. Yoksas, and L. Miller, 2007: Building and Sustaining International Scientific Partnerships Through Data Sharing. *American Meteorological Society Conference Proceedings,* 6 pp. Available online at http://ams.confex.com/ams/pdfpapers/120079.pdf .

Ramanathan, V. (chair), et al., 2009: *Restructuring Federal Climate Research to Meet the Challenges of Climate Change.* Committee on Strategic Advice on the U.S. Climate Change Science Program, National Research Council of the National Academies, Washington, D.C., 266 pages. See pp. 220-221.

Randall, D.A., M. Khairoutdinov, A. Arakawa, and W. Grabowski, 2003: Breaking the cloud-parameterization deadlock. *Bull. Amer. Meteor. Soc.,* **84**, 1547-1564.

Randall, D., S. Krueger, C. Bretherton, J. Curry, P. Duynkerke, M. Moncrieff, B. Ryan, D. Starr, M. Miller, W. Rossow, G. Tselioudis, and B. Wielicki, 2004: Confronting Models with Data: The GEWEX Cloud Systems Study (GCSS). *Bull. Amer. Meteor. Soc.,* **84**, 455-469.

Rempel, M., M. Schüssler, and M. Knölker, 2009a: Radiative MHD simulation of sunspot structure. *The Astrophysical Journal,* **691** (1), 640-649.

Rempel, M., M. Schüssler, R.H. Cameron, M. Knölker, 2009b: Penumbral Structure and Outflows in Simulated Sunspots. *Science,* **325** (5937), 171-173.

Rotunno, R., Y. Chen, W. Wang, C. Davis, J. Dudhia, and G.J. Holland, 2009: Large-eddy simulation of an idealized tropical cyclone. *Bull. Amer. Meteor. Soc.,* In press, submitted March 19, 2009; revised July 1, 2009.

Rizzetta, D.P., M.R. Visbal, and G.A. Blaisdell, 2003: A time-implicit high-order compact differencing and filtering scheme for large-eddy simulation. *Intl. J. for Num. Methods in Fluids,* **42**, 665-693.

Rodriguez, M.A., and D. Dabdub, 2004: A modeling study of size- and chemically resolved aerosol thermodynamics in a global chemical transport model. *J. Geophys. Res.,* **109**, D02203.

Sagaut, P., 2002: *Large Eddy Simulation for Incompressible Flows.* 2nd ed. Springer, 426 pp.

Satoh, M., T. Matsuno, H. Tomita, H. Miura, T. Nasuno, and S. Iga, 2008: Nonhydrostatic icosahedral atmospheric model (NICAM) for global cloud resolving simulations. *J. Comp. Phys.,* **227**, 3486-3514.

Scharmer, G.B., B.V. Gudiksen, D. Kiselman, M.G. Löfdahl, and L.H.M. Rouppe van der Voort, 2002: Dark cores in sunspot penumbral filaments. *Nature,* **420** (6912), 151-153.

Schüssler, M., and A. Vögler, 2006: Magnetoconvection in a Sunspot Umbra. *The Astrophysical Journal,* **641** (1), L73-L76.

Semtner, A., (chair), et al., 2006: CISL High Performance Advisory Panel (CHAP) Letter to NCAR Director, October 4, 2006. Page: "It is our opinion that NCAR has demonstrated a track record of acquiring and managing supercomputing resources that are second to none."

Shukla, J., J. Kinter, D.M. Straus, 2008: *What are the Modeling Requirements for Seamless Prediction of Weather and Climate from Days to Decades? The Known, the unknown, and the unknowable.* Report of a workshop held at the Royal Society UK, December 1-2, 2007 with support of the Sloan Foundation. COLA Technical Report 263, May 2008, 17 pp. See page 7.

Shukla, J., et al., 2009: *World Climate Research Programme (WCRP) Workshop Report: World Modelling Summit for Climate Prediction.* Workshop held in Reading, UK, May 6-9, 2008, WCRP No. 131, WMO/TD 1468, January 2009, 29 pages. See especially conclusion 5, page 5.

Siegel, A., J.B. Weiss, J. Toomre, J.C. McWilliams, P.S. Berloff, and I. Yavneh, 2001: Eddies and vortices in ocean basin dynamics. *Geophys. Res. Lett.*, **28**, 3183-3186.

Skamarock, W.C., J. Oliger, and R.L. Street, 1989: Adaptive grid refinements for numerical weather prediction. *J. Comp. Phys.*, **80**, 27-60.

Smagorinsky, J., 1963: General circulation experiments with the primitive equations. Part I: the basic experiment. *Mon. Wea. Rev.*, **91**, 99-164.

Snavely, A., R. Loft, R. Pennington, 2008: *Petascale Computing in the Geosciences Workshop Report, 2008*, 43 pages. Report summarizes experiences and lessons learned from three separate workshops held in 2006 and 2007 including Geosciences Application Requirements for Petascale Architectures (GARPA) 1 (June 1-2, 2006), GARPA-2 (February 21-22, 2007), and Petascale Computation for the Geosciences-I workshop (April 4-5, 2006). Presentations available at http://www.sdsc.edu/PMaC/GeoScience_Workshop/. (Funding for GARPA-1 and GARPA-2was provided by NSF-ATM award 0540688 received through subcontract 2005-006559-01 with the University of Illinois; Petascale Computing and the Geosciences was supported by NSF-GEO-0621611.)

Solomon, S., D. Qin, M. Manning, Z. Chen, M. Marquis, K.B. Avery, M. Tignor, and H.L. Miller, 2007: *Climate Change 2007: The Physical Science Basis.* Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge, UK, 996 pp.

Staniforth, A., and J. Côté, 1991: Semi-Lagrangian integration schemes for atmospheric models—a review. *Mon. Wea. Rev.*, **119**, 2206-2223.

Stauffer, P.H., R.S. Surdam, J. Jiao, T.A. Miller, and R.D. Bentley, 2009: Combining geologic data and numerical modeling to improve estimates of the CO_2 sequestration potential of the Rock Springs Uplift, Wyoming. *Energy Procedia*, in press.

Stern, N., 2006: *Stern Review: the Economics of Climate Change.* Available at http://www.hm-treasury.gov.uk/stern_review_report.htm .

TIGGE, THORPEX Interactive Grand Global Ensemble (website), 2009: TIGGE Data Archive Portal. NCAR, Boulder, Colo. Available online at http://tigge.ucar.edu/home/home.htm .

Vaillancourt, P.A., M.K. Yau, P. Bartello, W.W. Grabowski, 2002: Microscopic approach to cloud droplet growth by condensation. Part II: Turbulence, clustering, and condensational growth, *J. Atmos. Sci.*, **59**, 3421-3435.

Walter, J., ed., 2005: *World Disasters Report 2005: Focus on Information in Disasters.* International Federation of the Red Cross and Red Crescent Societies, Kumarian Press, Sterling, Va, 246 pp. Available online at http://www.ifrc.org/publicat/wdr2005/ .

Wang, L.-P., O. Ayala, S.E. Kasprzak, and W.W. Grabowski, 2005: Theoretical formulation of collision rate and collision efficiency of hydrodynamically-interacting cloud droplets in turbulent atmosphere. *J. Atmos. Sci.*, **62** (7), 2433-2450.

Wehner, M., L. Oliker, and J. Shalf, 2008: Towards Ultra-High Resolution Models of Climate and Weather. *International Journal of High Performance Computing Applications*, **22** (2), 149-165. Available online at http://www.nersc.gov/projects/SDSA/reports/uploaded/IJHPCA06_CAM_final.pdf .

Weisman, M.L., C.A. Davis, W. Wang, and K. Manning, 2008: Experiences with 0-36h Explicit Convective Forecasts with the WRF-ARW Model. *Wea. And Forecasting*, **23**, 407-437.

Wyngaard, J.W., 2004: Toward numerical modeling in the “terra incognita”, *J. Atmos. Sci.*, **61**, 1816-1826.