
# Report of the 2013 NSF CyberBridges Workshop:

# Developing the Next Generation of Cyberinfrastructure Faculty for Computational- and Data-enabled Science and Engineering

July 15-16, 2013
Arlington, Virginia

**Suzanne Shontz**
Mississippi State University

**Thomas Hacker**
Purdue University

# Table of Contents

## 1.0 Executive Summary

The second annual workshop for the NSF Division of Advanced Cyberinfrastructure (ACI) CAREER awardees was held July 15-16, 2013 in Arlington, Virginia. Thirty-one attendees and five keynote speakers attended the workshop. The attendees, who were funded by thirteen directorates, were selected from twenty-nine institutions. Sixteen of the attendees from the 2012 workshop returned for this year's workshop. Five keynote presentations were given by nationally- and internationally-recognized leaders in fields relevant to the use and development of cyberinfrastructure in science and engineering research. Each keynote presentation was followed by a discussion session with workshop attendees.

The workshop provided a venue for CAREER awardees to interact and to develop new collaborations with leading researchers and other CAREER awardees. As a result of this workshop, 46 potential new collaborations were identified by attendees.

Attendees expressed several concerns.

1. In the area of how new faculty work and thrive, similar to last year, concerns were focused on issues involved in interdisciplinary research and non-traditional forms of knowledge dissemination. *Disciplinary barriers, along with a lack of coordinated interdisciplinary funding opportunities continue to impede progress.* To address this, a suggestion is to **increase the number of interdisciplinary funding opportunities and to establish a communications process within the funding agencies to assist researchers seeking funding for interdisciplinary projects.** In addition to these continuing concerns, attendees *discussed the need for more pervasive education in computational modeling, and increasing the use of high performance computing in areas of science currently not well served by HPC.* A possible approach to address this problem would be to **promote the teaching and use of computational models and tools across first and second year college courses along with the establishment of an award program to recognize novel effective teaching methods for early career faculty.**

2. In the area of research, attendees felt that *the highest capability supercomputers today tend to be oversubscribed, and that the community could benefit from access to a broader set of heterogeneous computing resources and new software environments that could simplify the use of these resources.* Science communities today that individually have "subcritical" demand may collectively have a large set of common needs that could help to drive new software approaches (e.g., domain specific languages) and hardware architectures (e.g., heterogeneous platforms with accelerators, FPGAs, and GP-GPUS). To address this concern, a suggestion is **to encourage the development of workshops to bring together members of these "subcritical" demand communities to collect needs and requirements as inputs into solicitations for resources that could benefit these communities.**

3. Attendees expressed *a need for access to a diverse and balanced mix of computing resources that range from local small scale clusters and cloud computing systems up to national scale*

*petascale and future exascale systems, along with a coordinated approach to simplify allocation requests across the spectrum of resources.* There is an ecosystem of applications that can thrive, evolve, and grow in scaling and size over time within this balanced mix of capabilities, and not all “big data” problems they encounter require very high end supercomputing resources. To address this need, a suggestion is to **establish a national resource allocation process or program with a single request, review, and award system.** This would provide researchers a “one stop shopping” portal to discover and request access to any federally sponsored shared computational resource.

4. The availability of high-quality sensor data is supporting work in simulation (the *forward* problem) as well as work on *inverse* problems seeking to infer underlying models from observational data. Improving the accuracy of these analyses in less time depends on better ways of quantifying uncertainty in the data as well as verifying and validating the correctness of computational models. In addition, new techniques for reducing the large high-dimensional models and for exploring the resulting high-dimensional space efficiently are also needed**. A suggestion is to encourage the development of new algorithms and capabilities in future solicitations.**

## 2.0 Workshop Overview

In 2010, the NSF Office of Cyberinfrastructure (now the Division of Advanced Cyberinfrastructure) began making full awards in NSF's Faculty Early Career Development (CAREER) program, to support investigators working on interdisciplinary research in cyberinfrastructure and the application of cyberinfrastructure to science and engineering research. We held the first workshop for Office of Cyberinfrastructure CAREER awardees on June 25-26, 2012, in Arlington, Virginia. At the time of the workshop, approximately 50 CAREER projects had been awarded to researchers. The attendees, who were funded by the NSF OCI, BIO, CISE, ENG, HRE, and MPS directorates, were selected from 24 institutions. Five keynote presentations were given by nationally and internationally recognized leaders in fields relevant to the use and development of cyberinfrastructure in science and engineering research. Each keynote presentation was followed by a discussion session with workshop attendees. The workshop provided a venue for CAREER awardees to interact and to develop new collaborations with leading researchers and other CAREER awardees. As a result of this workshop, 55 potential new collaborations were identified by attendees. The workshop provided many opportunities for discussions among attendees and speakers. We received many positive comments and positive survey feedback from the attendees, and encouragement to propose a follow-on workshop in 2013.

To bring together the community of ACI CAREER awardees and to build upon the successes of the 2012 workshop, we proposed and held a workshop on July 15-16 2013 in Arlington, Virginia with several goals: (1) encourage networking and discussion among awardees; (2) provide a forum to facilitate the discovery of new synergies and connections among researchers from the community; and (3) provide inspiration and motivation for new research through a series of keynote presentations by leaders in the areas of Computational- and Data-enabled Science and Engineering, Scientific Visualization, High Performance Computing, Education, and Grand Challenges in Cyberinfrastructure. The workshop provided networking opportunities for attendees to seek out and gain potential collaborators, and included a poster session that allowed poster presenters to solicit additional interest from attendees.

## 2.1 Outcomes of the NSF CyberBridges Workshop

The attendees found the workshop to be useful, informative, and engaging with sufficient opportunities for networking and exploring collaborations. Based on survey results, participants agreed that the five thematic areas of the workshop included their areas of research and education, and that the disciplinary areas of workshop attendees were sufficiently broad to facilitate interdisciplinary engagements. Participants slightly less strongly agreed that the workshop was helpful to learn more about the NSF and available funding opportunities. Attendees identified a total of 46 unique opportunities for potential collaborations.

Two areas of discussion emerged from the keynote presentations and subsequent attendee discussion following each presentation. The first was focused on the ways in which new faculty work and thrive. The second was concentrated on advances in research areas related to the attendees' work. In the first area, similar to last year, concerns were focused on issues involved in interdisciplinary research and non-traditional forms of knowledge dissemination.  In addition to these continuing concerns, attendees discussed the need for more pervasive education in computational modeling, increasing the use of high performance computing in areas of science currently not well served by HPC, and the need for a comprehensive national strategy for HPC from the laboratory level through the campus level up to the national level. The second common area, advances in research, centered on several areas.  In the area of big data, attendees discussed the need for new approaches for uncertainly quantification, validation, and verification for simulating systems as well as addressing the reverse problem of inferring system models from observational data. Moreover, new approaches are needed for exploring the high-dimensional space present in big data.  In the area of computing platforms and architectures, attendees described the need for a range of computing architectures from campus level to national level to solve the types of problems they encounter.  Not all problems need national resources alone – campus and laboratory level resources are adequate for many of the problems.  Additionally, to better use existing high performance computing resources, attendees expressed the need for technologies (such as domain specific languages) to simplify the use of these systems, as well as the need for innovative research in new architectures to reach exascale.

**Figure 1.**  Map of the United States showing locations of the various attendee home institutions at the CyberBridges Workshop.

## 2.2 Workshop Attendees

Thirty-one attendees and five keynote speakers attended the workshop, from twenty-nine institutions (Fig. 1).  Most attendees were funded at least in part through the Division of Advanced Cyberinfrastructure, but several held awards partially funded through other divisions, including the Computing and Communication, Computer and Network Systems, Electrical, Communications and Cyber Systems, Information and Intelligent Systems, Materials Research, Atmospheric and Geospace Sciences, and Environmental Biology (Fig 2). Nine attendees were funded entirely from divisions outside ACI, and ten attendees were funded through multiple divisions. This year there was a greater breadth of interdisciplinary research, with thirteen divisions represented (compared to six last year).

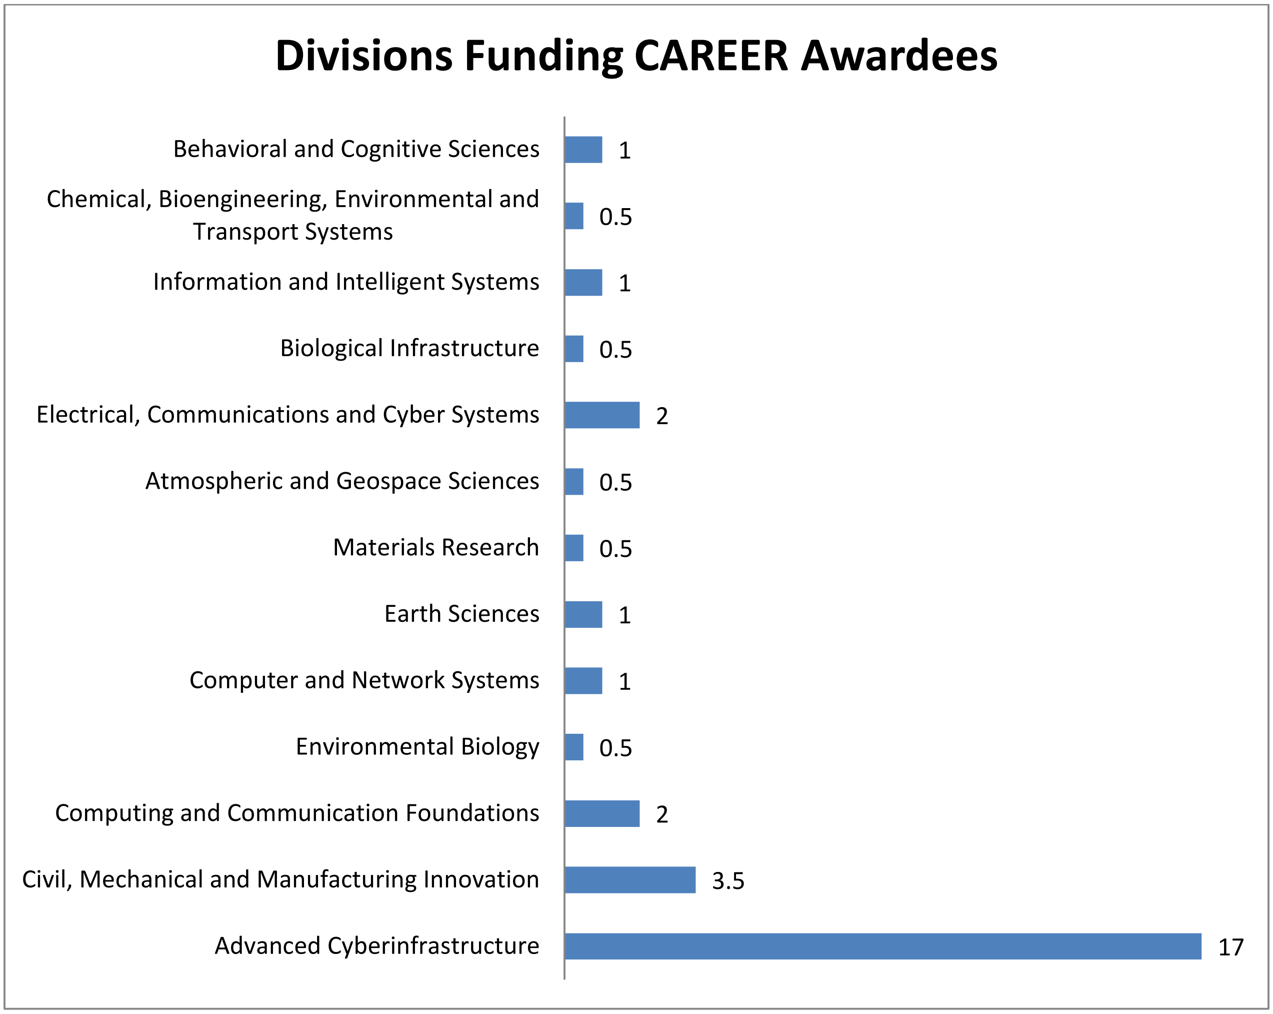


**Figure 2.** Distribution of the divisions funding the attendees of the CyberBridges Workshop. CAREER Awards funded through divisions are split evenly between the relevant offices.

### 2.3 Attendee Selection Process

Faculty who received NSF CAREER Awards from ACI (including through co-funding with another NSF directorate of division) were invited to attend the workshop. The invitation process (in priority order) was to invite new (over the past year) CAREER awardees before prior awardees, followed by ACI awardees before awardees from other NSF directorates.

## 3.0 Workshop Themes

The workshop focused on five thematic areas that reflect the spectrum of research and education activities in which the Division of Advanced Cyberinfrastructure (ACI) is focused, and they encompass the types of computational- and data-enabled science and engineering (CDS&E) in which ACI and many other NSF directorates are engaged.

1. <u>Computational- and Data-enabled Science and Engineering</u>

   The first thematic area of the workshop focused on *computational*- and *data-enabled science and engineering*, which involves the development of algorithms and cyberinfrastructure necessary to perform large-scale simulations or to process and interpret data generated from experiments, simulations, models, and observations in science and engineering. Dr. Omar Ghattas gave a keynote talk on solving inverse problems and led discussion among workshop participants. Fundamental issues that occur include developing efficient algorithms to solve inverse problems which: (1) handle the ill-posedness, non-causality, global behavior, and uncertainty exhibited by the Bayesian inverse problem; (2) explore and navigate high-dimensional parameter space coupled with expensive evaluations of the forward problem; (3) overcome the poor scalability of the Markov Chain Monte Carlo method; (4) reduce the effective problem dimension and exploit structure, and (5) map well onto extreme-scale systems and scale independently of parameter dimension, state dimension, data dimension, and the number of cores. An example of an application area where such techniques are needed is global seismic wave propagation in earthquake simulation.

2. <u>Scientific Visualization</u>

   The second thematic area of the workshop focused on *scientific visualization*, which focuses on the graphic illustration and rendering of scientific data for use by scientists and engineers to interpret their data. Scientific visualization is used in biology, medicine, meteorology, architecture, and fluid dynamics, for example. Dr. Chandrajit Bajaj gave a keynote talk on visualization for data-enabled modeling and uncertainty quantification for use in drug discovery. Fundamental challenges in scientific visualization include: (1) how to estimate the uncertainty present in data (e.g., from experiments or images) used as input to the visualization; (2) how best to reconstruct the data as a visualization (this is an ill-posed problem) in a stable and efficient manner; (3) how to perform scalable visualizations of a simulation with a projected runtime of $10^9$ hours; (4) data sets need to be curated and open-source to be of use to researchers; however there is not often funding for this, and (5) visualization software also needs to be open source to be of use.


**Figure 3**. Thirty one workshop participants attended talks and presented their work in a poster session.

**3.** <u>High Performance Computing</u>

The third thematic area of the workshop focused on *high performance* computing, the challenges affecting high performance computing today in petascale systems, and new challenges to be overcome to achieve exascale performance in the next decade. Dr. William Gropp gave a presentation on overcoming these challenges. The first challenge is in the area of big data. Dr. Gropp posited that big data requires “big compute” that is built on a low latency/high bandwidth communications interconnect.

Cloud technology can be useful for some services, but scales very poorly to the levels needed for analytics of petabytes of data. The second challenge is in the transition to exascale computing. Today, MPI is the primary programming paradigm used for applications that can effectively use petascale systems. Replacing MPI would be complex, and alternatives would likely be science domain specific in terms of providing a higher level of abstraction with some loss of generality as a consequence. Dr. Gropp expects that exascale systems in the next decade will suffer from extreme power constraints that will shape system architecture. In this environment, it will be expensive in terms of time and power to move data between processors. Moreover, the high frequency of faults in exascale systems will put more burden on the programmer to create resilient applications that can tolerate faults. Another problem will be performance variability across large systems, which will affect domain decomposition techniques that are not adaptive. To prepare for the next generation of systems, Dr. Gropp recommended focusing on improving the use of existing resources, adapting applications to better utilize the I/O capabilities available on systems today, and integrating fault tolerance and the ability to adapt to system conditions into applications.

**4.** Education

The fourth thematic area of the workshop focused on *education.* Dr. Steve Gordon described the growing unmet national need for trained scientist and engineers, and the problems attracting and retaining undergraduate and graduate students to science and engineering disciplines needed to fill the gap. One of the major problems he identified was the poor quality of instruction, and difficulties students encountered in learning the basic concepts needed for science and engineering. To address these problems, Dr. Gordon described efforts within the NSF XSEDE program focused on facilitating the integration of computational science within the curriculum that includes approaches based on inquiry-based learning. The program seeks to provide assistance with the initiation and enhancement of formal computational science and engineering undergraduate and graduate programs. Their approach is based on providing model curricula and course descriptions, promoting the development of virtual communities focused on education, and providing professional development for faculty. One of the key characteristics of this effort is the inherent interdisciplinary nature of the effort that involves material from Mathematics, Computer Science, Social Sciences, Physical Sciences, and Engineering. From his efforts, Dr. Gordon has found that the expertise necessary for this effort on campuses is often dispersed across multiple departments, colleges, and institutions. As a result of this dispersion of expertise, there are often difficulties negotiating the requirements, responsibilities, and arrangements needed to develop a computational science program.

**Figure 4**. Breakout groups discussed key questions offered by keynote speakers for each thematic area.

**5.** Grand Challenges and Interdisciplinary Research

The final thematic area of the workshop focused on grand challenges and interdisciplinary research. Dr. Brian Athey described lessons learned over almost 20 years of leading biomedical cyberinfrastructure projects. First, to be successful, cyberinfrastructure center projects need a driving set of science domain problems that help focus development and operations efforts. It can take several product cycles for the cyberinfrastructure to achieve relevance, and cannot be driven by education and training alone – meeting research needs helps to drive sustainability. Another lesson learned is the need to employ a qualified project managed from the first day of the project. Moreover, a contract between the grant agency for large projects requires a cooperative agreement that allows for "changes on the fly" in response to changing needs and conditions. Often, agencies will fund efforts up to the point of a successful demonstration, but sustainability over the long term frequently is not addressed. Dr. Athey also described the role of state level (rather than federal) funding to help position a research group or center for national impact when it is used as a "seed crystal" for a national center or an academic department. Finally, he mentioned the shifting paradigm for successful NSF cyberinfrastructure projects that rely on tight focus on deliverables in combination with a high level of participation from the community. This approach, in contrast to the extremes of classic "command and control" organizations and the "let a thousand flowers bloom" strategy for traditional academic research, is a new paradigm that is being successfully employed for NSF-sponsored cyberinfrastructure projects.

### 4.0 Invited Speakers and Panelists from the National Science Foundation

***Dr. Evelyn Goldfield*** discussed several major cyberinfrastructure challenges in chemistry which include: integration and inter-operability of existing software elements, codes, etc., for the benefit of the end user; positioning codes to quickly respond to disruptive changes in technology; development of efficient, effective algorithms for future HPC architectures; meeting big data challenges that arise from simulation or experiment, or both; review of cyberinfrastructure aspects of chemistry proposals and reward of PI's who develop high-quality cyberinfrastructure; modifying education to include more courses in advanced math, programming, and software engineering, and encourage more interactions with cyberinfrastructure experts.

***Dr. Daniel Katz*** presented on the challenges of software as infrastructure at NSF/CISE/ACI. The objective of the program is to "Create and maintain a software ecosystem providing new capabilities that advance and accelerate scientific inquiry at unprecedented complexity and scale". To support this objective, four areas of effort are needed: 1) support foundational research needed to advance scientific software; 2) the development of new policies for software that addresses the needs of the academic community for attribution while ensuring the sustainability of software; 3) enable transformative and collaborative science and engineering research through use of advanced software and services; and 4) develop a diverse workforce of scientists and engineers equipped with the necessary skills to develop and use software and services. Dr. Katz described several new and ongoing programs: Exploiting Parallelism and Scalability (XPS), Computational and Data-Enabled Science & Engineering (CDS&E), and Software Infrastructure for Sustained Innovation (SI^2). He then described some of the challenges that must be addressed, which involves finding a balance in funding existing vs. new infrastructure; understanding the true efficacy of the open source model for science; encouraging users to reuse existing software and to discourage duplication of efforts; and supporting career paths for software developers with universities and laboratories.

**Figure 5.** A panel of NSF program directors presented an overview of NSF programs and activities.

***Dr. Peter McCartney*** described some of the cyberinfrastructure-related activities in the BIO directorate, which included the Protein Data Bank (PDB) and the BIO Synthesis Center.

***Dr. Thomas Russell*** described several challenges in computational science (from the computational math perspective) in the area of modeling including: multiscale, multiphysics, and multimodels; verification, validation, and uncertainty quantification, and inverse problems and data assimilation. Other fundamental challenges include optimization under uncertainty, efficient algorithms at peta/exascale, and big data. He also described challenges in the area of big data including reproducible research and a desire for communitarian behavior with the goal of having a shared cyberinfrastructure. NSF data management plans describing plans for sharing and dissemination of research data and results are now required for all NSF proposals.

***Dr. Barry Schneider*** gave an overview of the NSF XSEDE programs and described some of the XSEDE resources available to the community. He described some recent developments in XSEDE, which includes the deployment of the Stampede system at TACC, the Keeneland system at Georgia Tech, and he described a potential new solicitation that is under review to expand the scope and range of NSF investments. Dr. Schneider described the growing use of XSEDE resources, and the broad impact the XSEDE program is making across the NSF and other federal agencies.

## 5.0 Attendee Feedback Survey

To collect attendee feedback, we conducted an anonymous Qualtrics survey that was emailed to workshop attendees after the workshop. The purpose of this survey was to collect feedback from attendees about what went well and what did not go well, and to ask for suggestions for improvement for the workshop. We received 16 responses to the survey (see the Appendix A for the complete text of the survey and responses).

**Table 1. Participant responses to the workshop survey. Responses are mean and standard deviation (S.D.). Survey participants answered questions on a Likert Scale with the following numeric assignment. Strongly Disagree (SD = 1), Disagree (D = 2), Neither Agree nor Disagree (AD = 3), Agree (A = 4), and Strongly Agree (SA = 5).**

| Question | Mean | S.D. |
|---|---|---|
| The five focus areas of the workshop included my area of research and education. | 4.0 | 1.0 |
| The disciplinary areas of workshop attendees were sufficiently broad to facilitate interdisciplinary engagement. | 4.0 | 1.0 |
| The workshop format (keynote talks followed by discussion) was useful and engaging. | 4.1 | 1.0 |
| The talks were relative, informative, and interesting | 4.1 | 1.0 |
| There were sufficient opportunities for networking and collaboration. | 4.1 | 1.1 |
| The poster session was useful and engaging. | 4.1 | 0.9 |
| The hotel accommodations, meeting space, and meals were adequate. | 4.3 | 1.1 |
| *The workshop was helpful in learning more about the NSF and available funding opportunities* | *3.8* | *1.2* |
| The workshop should include CAREER awardees beyond OCI. | 3.8 | 1.3 |
| The workshop should include attendees from outside the NSF CAREER program | 3.1 | 1.1 |

Based on survey results, participants agreed that the five thematic areas of the workshop included their areas of research and education, and that the disciplinary areas of workshop attendees were sufficiently broad to facilitate interdisciplinary engagements (Table 1). Participants also agreed that that workshop format (talks followed by discussion) was useful and engaging, the keynote talks were informative and interesting, and that the poster session was useful and engaging, and there were sufficient opportunities for networking and collaboration.. Participants slightly less strongly agreed (3.8) that the workshop was helpful to learn more about the NSF and available funding opportunities

In terms of participants' perceptions about future participants, results indicate some agreement (3.8) that the workshop should include CAREER awardees beyond ACI[1], but only low neutral agreement (3.1) that the workshop should include attendees outside the CAREER program.

Participants felt that the workshop length and the number of attendees were about right. Also, 60% of the survey responses indicated interest in attending the workshop even if full travel reimbursement were not provided. Half the respondents felt that the workshop should be held annually, while the other half would like to see it held every other year.

Although participants generally agreed that the workshop format was useful, several survey respondents suggested that the workshop could be improved by making the format more dynamic, with better organized small group discussions. One person suggested that the workshop could include a round robin discussion with NSF program officers to learn about upcoming funding calls and to help CAREER awardees begin thinking about the next steps in their careers.

Compared with results from last year, the mean response values were approximately 10% lower this year. Participants were less positive this year about the workshop focus areas, breadth of discipline, format, talks and poster sessions, and the hotel. Participants felt similarly this year and last year about the opportunities for networking and collaboration. However, the variance was higher this year and resulted in no statistical difference from last year. The largest mean difference from last year was for the question eight (*the workshop was helpful in learning more about the NSF and funding opportunities*) which was lower this year (3.8 ±1.2) compared with last year (4.7 ±0.48). This was also reflected in the comment suggesting a round robin discussion with NSF program officers.

[1] The survey instrument approved by the Purdue IRB used the term OCI (Office of Cyberinfrastructure), which should now be interpreted as the Division of Advanced Cyberinfrastructure (ACI).

## 6.0 Observations from the Workshop

*6.1 The ways in which new faculty work and thrive*

### 6.1.1 Interdisciplinary research

Similar to the workshop last year, there was concern among faculty about the difficulties of engaging in interdisciplinary research. At an institutional level, traditional promotion and tenure practices make it difficult to work collaboratively across disciplines. Moreover, funding agencies are not structured to support interdisciplinary research that does not clearly fit into one existing area. The perceived limited availability of interdisciplinary funding is an ongoing concern for faculty who work across traditional disciplinary boundaries.

**Figure 6.** NSF CISE Director Dr. Farnam Jahanian confers with workshop co-chairs Drs. Hacker and Shontz.

### 6.1.2 Education

Computing needs to be added more pervasively to K-12 and college curriculums and computational modeling needs to be introduced much earlier. Students in the sciences need more exposure to computation earlier in the curriculum through a holistic approach based on leveraging existing courses or programs.

A different approach to teaching that goes beyond traditional lectures in classrooms is needed to help improve student learning.

The barriers to some of these changes are due in part to the unfamiliarity of senior faculty to new techniques, technologies, and computational methods. The hesitancy to foster change within the faculty slows adoption of these new approaches by untenured faculty.

### 6.1.3 Non-traditional forms of knowledge dissemination are critical

Similar to the workshop last year, faculty described the need for recognition of new forms of knowledge dissemination in the form of community datasets, software, and by promotion and tenure committees as well as funding agencies

### 6.1.4 Increasing the use of computing and HPC in new areas of science

Attendees felt that to increase the use of computing and high performance computing in new areas of science (those traditionally underserved with computing) the technologies need to become easier to use.

Additionally, the access to these resources need to be made as real-time as possible. High-end HPC systems (such as Blue Waters) are highly oversubscribed. One area in particular is the need for science domain specific languages that can exploit the advanced capabilities available such as GP-GPUs, Phi processors and parallel computing resources. Attendees expressed a need for a coordinated national approach to work with the science community to identify emerging needs from the science community in which individually there is “subcritical” demand, but in aggregate have an adequate level of need to help drive new directions for architecture.

## *6.2 Advances in research areas*

### 6.2.1 Inverse problems

The emerging availability of vast amounts of high quality data collected from sensors and observations is driving new research into inverse problems, which seeks to ‘work backwards’ from the data to discover the underlying models, systems, or structures from which these data emerge. The high dimensionality of these data, along with the need for better methods for quantifying the uncertainty in these data, presents new challenges for the research community.

### 6.2.2 Future HPC architectures

Attendees described a perceived need for innovative research in new architectures to aid efforts to reach exascale. The momentum of the commodity hardware market from which HPC has been able to exploit is not sufficient to power the next level of advances in architecture needed to achieve exascale.

Attendees felt that there is a need for a diverse and balanced mix of computing resources that range from local small scale cluster and cloud computing systems up to national scale petascale and future exascale systems. There is an ecosystem of applications that can thrive, evolve, and grow in scaling and size over time within this balanced mix of capabilities.

### 6.2.3 Big data and data in general

Attendees felt that not all “big data” problems they encounter require very high end supercomputing resources. Applications involving the use and analysis of big data can fall into two groups: those that require very large scale systems (such as Blue Waters); and smaller scale machine learning and business analytics problems that can be solved with Hadoop: Many problems can be addressed on smaller clusters, perhaps on a campus level. Other than big data problems, there are many public databases and community datasets that have been created and supported with grant funds. However, the long-term sustainability of these data bases has not been adequately addressed, and there needs to be some mechanism by which researchers may request funds to help support these databases.

### 6.2.4 Verification and validation / uncertainty quantification

The growing body of high quality data collected from sensors combined with the tremendous computational power available today is stimulating work in computational modeling (the *forward* problem) to simulate the output of systems given a range of input parameters as well as work on the *inverse* problem that seeks to infer the input parameters of underlying models from observational data. Improving the accuracy of these analyses in shorter time is depends on better ways of quantifying uncertainty in the data as well as verifying and validating the correctness of computational

### 6.2.5 High dimensional space as the new frontier

High-dimensional space is a new frontier for which algorithms must be developed to explore high-dimensional data. There are several research avenues to be explored. First, model order reduction techniques can be developed and used to reduce the dimension of the high-dimensional data. This would allow for simulations involving the data to be performed more efficiently based on their structures. Second, there is a need for new optimization and sampling algorithms that allow for more efficient traversal of high-dimensional space. There is much work to be done in these areas, particularly in connection with the new wave of big data sets.

### 6.2.6 Visualization and high performance computing

There are many science problems that could benefit from both high performance computing and visualization of the numerical results from the algorithm. However, it is often not the case that researchers are generating visualizations of their scientific data. Attendees felt that in order for visualization to be more commonly used amongst high performance computing experts, visualization software needs to become easier to use. In addition, hybrid architectures that contain both GPUs and a sufficient amount of storage per node are needed so that both the high performance computing and visualization can be done on the same machine. This would help researchers avoid the endless transfer of data between machines with different architectures. In addition, based upon the visualization, dynamic steering of the high performance computing calculations could be performed so that only relevant parameters are used for calculations. This would help reduce the number of compute cycles, as well.

## 7.0 Summary of observations and suggested action items

We summarize a number of other observations from the workshop, and propose a series of action items based on our analysis of comments collected from attendees.

1. Interdisciplinary research is still an important area for new faculty. Unfortunately, institutional and disciplinary barriers remain that impede progress in advancing interdisciplinary research beyond computing. Funding for interdisciplinary research continues to be difficult to navigate and win.

**Action Item 1: Increase the number of interdisciplinary funding opportunities. Emphasize interdisciplinary aspects of existing funding opportunities that already include interdisciplinary work. Establish a communication process within NSF to assist researchers seeking funding for interdisciplinary projects than span divisions.**

2. Computation needs to be introduced earlier in the STEM curriculum to aid students who do not intend to become computer scientists to acquire computational skills needed for their discipline. It would be helpful if promotion and tenures policies and practices in place in institutions would evolve to put more value on new forms of teaching and value on computational science education.

3. New faculty continue to communicate the need for recognition from funding agencies and for promotion and tenure for non-traditional forms of knowledge dissemination such as community datasets and software.

**Action Item 2: Promote the teaching and use of computational models and tools across first and second year college courses, as well as high school science courses. Establish an award program to recognize the use of novel teaching methods involving computational methods for early career faculty.**

4. There are new areas of scientific discovery that are underserved in the use of computing and HPC. An overall national coordinated approach is needed to help identify and aid these communities to achieve the critical mass needed to make progress in the use of computing.

**Action Item 3: Encourage workshops to bring together members of research communities who traditionally have not reached a critical mass in their use of high performance computing to assess need and gather requirements from these communities. Encourage the diffusion and adoption of HPC technologies (through programs such as XSEDE) for these research communities.**

5. There is a clear need for a coordinated national approach in harmonizing allocation requests and high performance computing resources ranging from campus level clusters

up to high-end HPC resources (such as Blue Waters) to alleviate the oversubscription of national level HPC systems.

6. There is a clear need for a diverse mix of computing resources to support an ecosystem of applications, ranging from small, campus-level clusters and cloud computing systems to very large, nationally-owned petascale and future exascale systems

7. "Big data" problems need to be solved on high-end HPC resources, whereas smaller data problems can be assigned to smaller clusters, perhaps on the campus-level.

**Action Item 4: Establish a national program to provide access to federally funded computational, visualization, and storage resources (such as XSEDE, Blue Waters, and other federally funded (by NSF or other agencies)) through a single request, review, and award system (such as Research.gov) to facilitate the discovery of and access to these resources by the research community.**

8. There are emerging challenges arising from the availability of "big data" from sensors that lead to high dimensionality data from which the underlying models emerge and must be discovered. This leads to the formation and solution of inverse problems.

9. High-dimensional space is a new research frontier. There is a need for model order reduction, sampling, and optimization algorithms that allow for efficient representation of the data and efficient traversal of the space.

**Action Item 5: Include details in future NSF solicitations involving data (e.g., CDS&E, DIBBS, etc.) so that inverse problems, model order reduction, and sampling of higher-dimensional space are listed as potential proposal topics. Encourage the development of new algorithms and tools for high dimensional space for model order reduction, sampling, and optimization algorithms that allow for efficient representation of the data and efficient traversal of the space.**

10. There is a need for innovative architecture research in order to reach the exascale, as the traditional commodity hardware will not be able to reach exascale when improved.

**Action Item 6: Promote the development of novel computational approaches and architectures needed to achieve exascale capabilities in the coming decade.**

11. There are many community datasets and public databases that have been supported by grants. There is a need for support to address the long-term sustainability of these data.

**Action Item 7: Encourage and support the development of domain specific curated data repositories. An example of this for the civil engineering community is the NEES Project Warehouse and NEES Databases, available through the NEEShub. Another**

**example is IRIS (Incorporated Research Institutes for Seismology), which maintains a repository for geophysical time-series data.**

12. There is a growing need for better methods for validation, verification, and uncertainty quantification. This is motivated by the availability of large amounts of data from high resolution sensors along with emerging work on the inverse problem that seeks to model and quantify the unknown systems from which these data emerge.

**Action Item 8: Encourage the development of new approaches for validation, verification, and uncertainly quantification. To promote the dissemination of best practices for V&V and UQ, promote the development of educational course modules for the undergraduate curriculum in computational science.**

13. In order for HPC users to generate visualizations of their scientific data, visualization software needs to become easier to use.

**Action Item 9: Usability for visualization software should be emphasized through the addition of a solicitation-specific review question for ACI funded visualization projects which will be used by the scientific community.**

**Also see Action Item 5.**

14. There is a need for hybrid architectures with both GPUs and a sufficient amount of storage per node so that both HPC and visualization can be performed on the same machine. This would also permit dynamic steering of the computation and a reduction of compute cycles.

**Action Item 10: Encourage the development of opportunities such as the Exploiting Parallelism and Scalability (XPS) solicitation to promote the development of novel hybrid architectures and software environments that can be easily used by domain scientists. Also see Action Item 5.**

## 8.0 Lessons Learned from the Workshop

Keynote speakers overall provided excellent talks on scientific challenges in their respective fields. However, we should also ask keynote speakers to provide advice to CAREER awardees as to how to build their careers. It would also help to have at least one speaker discuss how his/her NSF CAREER Award (or a related award) helped launch his/her career. We are planning to do this for the 2014 workshop.

**Figure 12.** The poster session allowed attendees to share their work with colleagues and NSF program directors.

The invitation process we developed for inviting speakers worked well; however, we need to start earlier in order to attract the highest caliber speakers who are leaders in the international community. Starting earlier allows for more potential dates for the workshop.

The scheme we used for putting NSF CAREER Awardees in priority groups for receiving invitations worked well this year. However, given that we now need to send out several rounds of invitations, we need to start earlier on this (which we are doing for the 2014 workshop).

Assigning scribes to keep and submit notes was helpful in capturing the thoughts of attendees in breakout groups. However, attendees need to create their scribe notes and a summary during the workshop; otherwise, it can take a long time to receive notes from each scribe, which delays the report. Scribe notes and a detailed summary should either be written neatly on a form which has been formatted or typed during the last hour of the workshop and submitted to us at the conclusion of it.

**Figure 13.** Dr. Steve Gordon presented XSEDE project education efforts.

The format of the workshop, i.e., talks followed by group discussion, worked well.

Workshop attendees would like to have input into the topics to be discussed during the breakout sessions. Allowing attendees to provide input into questions discussed during the breakout groups would keep them more engaged in the discussion, as it would be a topic that they have the background to discuss. We should also work with keynote speakers to be sure that their questions are relevant to this audience.

The continuous poster session worked well.

In regards to the hotel venue, a hotel closer to NSF should be selected. Although this will cost somewhat more, more NSF program directors should be able to attend the workshop. In addition, a higher quality hotel venue will likely be better set-up to host the workshop (e.g., with working microphones and better quality food).

We should ask attendees to fill out the workshop survey at the end of the workshop. This would ensure that we have more respondents. Alternatively, we can distribute the survey via e-mail from one of the PIs. The survey should be made mandatory, as future workshop offerings are dependent upon the feedback we receive from attendees.

We should consider getting another NSF CAREER Awardee involved in the organization of the workshop given that 2014 is the third year that we are organizing it, and new ideas are sought.

**Acknowledgements**

Our efforts in planning and conducting the CyberBridges workshop and the development of this report were supported by Natasha Nikolaidis (Purdue), Linda White (Mississippi State), Debbie McBride (Mississippi State), Stephanie Schmidt (Purdue), Debbie Miethke (Purdue), Natalie Eaves (Mississippi State), Michelle Latham (Mississippi State), Cathy Chandler (Mississippi State), and Marilda O'Bryant (Mississippi State).

**Appendix A. Detailed Survey Results**

The complete text and responses of the survey sent out to attendees are described below.

In the first section, survey participants answered questions on a Likert Scale with the following numeric assignment. Strongly Disagree (SD = 1), Disagree (D = 2), Neither Agree nor Disagree (AD = 3), Agree (A = 4), and Strongly Agree (SA = 5).

1. The five focus areas of the workshop (Grand Challenges, Data, Visualization, Computational Science, and High Performance Computing) included my area of research and education).
   Results: Mean Value: 4.0. Responses: (1) SD, (2) AD, (8) A, (5) SA

2. The disciplinary areas of workshop attendees were sufficiently broad to facilitate interdisciplinary engagement.
   Results: Mean Value: 4.0. Responses (1) SD, (2) AD, (8) A, (5) SA

3. The workshop format (keynote talks followed by discussion) was useful and engaging.
   Results: Mean Value: 4.1. Responses (1) SD, (1) AD, (9) A, (5) SA

4. The talks were relative, informative, and interesting.
   Results: Mean Value: 4.4. Responses (1) SD, (10) A, (5) SA

5. The poster session was useful and engaging.
   Results: Mean Value: 4.1. Reponses (1) D, (2) AD, (6) A, (6) SA

6. There were sufficient opportunities for networking and collaboration.
   Results: Mean Value: 4.1. Responses (1) SD, (2) AD, (7) A, (6) SA

7. The hotel accommodations, meeting space, and meals were adequate.
   Results: Mean Value: 4.3. Responses (1) SD, (1) AD, (6) A, (8) SA

8. The workshop was helpful in learning more about the NSF and available funding opportunities.
   Results: Mean Value: 3.8. Responses (1) SD, (2) D, (1) AD, (8) A, (4) SA

9. The workshop should include CAREER awardees beyond OCI.
   Results: Mean Value: 3.8. Responses (1) SD, (2) D, (3) AD, (4) A, (6) SA

10. The workshop should include attendees from outside the NSF CAREER program.
    Results: Mean Value: 3.1. Responses (2) SD, (1) D, (8) AD, (3) A, (2) SA

Next, the survey asked participants to rate the following questions with a response from the options Too short/too few (1 = S); About right (2 = AR); and Too long/too many (3 = L).

11. Length of the workshop
    Results: Mean Value: 1.9. Responses (1) S, (14) AR

12. Number of attendees
    Results: Mean Value: 1.9. Responses (1) S, (15) AR

**The next questions asked about the frequency and cost of the workshop, and about the respondent's source of CAREER funding.**

13. Would you be interested in attending the workshop in the future without full travel reimbursement?
    Results: Yes 60%, No 40%

14. How frequently should the workshop be held?
    Twice a year: 0%, Annually 50%, Every other year 50%

15. Is a component of your CAREER award funded from OCI?
    Results: Yes 56%, No 44%

**The final questions provided an open form to allow participants to provide written feedback:**

1. Are there any new broad areas or topics that you would like to see covered in a follow on workshop?
   a. multiscale modeling, uncertainty quantification, fluid dynamics
   b. material informatics
   c. HPC in science and engineering
   d. more chemistry/materials science, alternative HPC infrastructures for shared memory applications
   e. software defined networks
   f. exascale computing, GPUs, hardware and software
   g. How to manage the tensions between Big Data, long-term archival semantics and reliable storage in light of limited budgets.
2. What changes or improvements could we make to the workshop in the future?
   Responses:
   a. I would have liked a round robin discussion where we could discuss our individual project ideas with various NSF program officers. It would be great for career awardees to start thinking about next steps and learn more about funding calls that are appropriate for their projects.
   b. Have longer talks

c. Format can be more flexible and dynamic
d. I would like to find more balance between science, software and infrastructure. What is new?
e. Do a better job of organizing small group discussion
f. Persuade the hotel to stop trying to imitate a cold autumn morning (68F in a locked up thermostat in summer was freezing cold...)

**Appendix B. Speaker and Attendee Biographies and Photos**
Biographies were current as of July 2013, at the time of the workshop

**Keynote Speakers and Invited Speakers from the National Science Foundation**

Dr. Farnam Jahanian serves as the National Science Foundation Assistant Director for the Computer and Information Science and Engineering (CISE) Directorate. He guides CISE in its mission to uphold the nation's leadership in scientific discovery and engineering innovation through its support of fundamental research in computer and information science and engineering and transformative advances in cyberinfrastructure. Dr. Jahanian oversees the CISE budget of over $850 million, directing programs and initiatives that support ambitious long-term research and innovation, foster broad interdisciplinary collaborations, and contribute to the development of a computing and information technology workforce with skills essential to success in the increasingly competitive, global market. He also serves as co-chair of the Networking and Information Technology Research and Development (NITRO) Subcommittee of the National Science and Technology Council Committee on Technology, providing overall coordination for the activities of 14 government agencies.

Alan Blatecky is the Director of Advanced Cyberinfrastructure at the National Science Foundation. Before coming to NSF, Alan was the Deputy Director of the Renaissance Computing Institute and has held executive leadership positions at the San Diego Supercomputing Center and the North Carolina Research and Education Network. Alan has focused on supporting research and development initiatives and programs in advanced high performance networking, computing, software, data, and visualization facilities, including early deployment and operations.

Omar Ghattas is the John A. and Katherine G. Jackson Chair in Computational Geosciences, Professor of Geological Sciences and Mechanical Engineering, and Director of the Center for Computational Geosciences in the Institute for Computational Engineering and Sciences at the University of Texas at Austin. He also is a member of the faculty in the Computational Science, Engineering, and Mathematics (CSEM) interdisciplinary PhD program in ICES, and serves as Director of the KAUST-UT Austin Academic Excellence Alliance. He has general research interests in simulation and modeling of complex geophysical, mechanical, and biological systems on supercomputers, with specific interest in inverse problems and associated uncertainty quantification for large-scale systems. His center's current research is aimed at large-scale forward and inverse modeling of whole-earth, plate-boundary-resolving mantle convection; global seismic wave propagation; dynamics of polar ice sheets and their land, atmosphere, and ocean interactions; and subsurface flows, as well as the underlying computational, mathematical, and statistical techniques for making tractable the solution and uncertainty quantification of such complex forward and inverse problems on parallel supercomputers.

William Gropp is the Thomas M. Siebel Chair in Computer Science, Computer Science Department; Director, Parallel Computing Institute; Deputy and Director for Research Institute for Advanced Computing Applications and Technologies at University of Illinois- Urbana-Champaign His research interests are in parallel computing, software for scientific computing, and numerical methods for partial differential equations. He has played a major role in the development of the MPI message-passing standard. He is co-author of the most widely used implementation of MPI, MPICH, and was involved in the MPI Forum as a chapter author for MPI-1, MPI-2, and MPI-3.He has written many books and papers on MPI including "Using MPI" and "Using MPI-2".He is also one of the designers of the PETSc parallel numerical library, and has developed efficient and scalable parallel algorithms for the solution of linear and nonlinear equations. Gropp is a Fellow of ACM, IEEE, and SIAM, and a member of the National Academy of Engineering. He received the Sidney Fernbach Award from the IEEE Computer Society in 2008 and the TCSC Award for Excellence in Scalable Computing in 2010.

Steven I. Gordon, Ph.D., is Interim Co-Executive Director for the Ohio Supercomputer Center (OSC) as well as director of the Center's Ralph Regula School of Computational Science. Gordon is the founding director of the Ralph Regula School of Computational Science. Gordon has also played a significant role in several programs in Science Technology Engineering and Mathematics education for high school and middle school students. Those include the Summer Engineering STEM Academy and the Young Women's Summer Institute.  Gordon's research applies models of storm water runoff and water quality to the analysis of watershed management and the applications of communications technology to distance education.

Chandrajit Bajaj is the director of the Center for Computational Visualization, in the Institute for Computational and Engineering Sciences (ICES) and a Professor of Computer Sciences at the University of Texas at Austin. Bajaj holds the Computational Applied Mathematics Chair in Visualization. He is also an affiliate faculty member of Mathematics, Electrical and Bio-medical Engineering, Neurosciences, and a fellow of the Institute of Cell and Molecular Biology. He is on the editorial boards for the International Journal of Computational Geometry and Applications, the ACM Computing Surveys, and the SIAM Journal on Imaging Sciences. He is a fellow of the American Association for the Advancement of Science (AAAS), and a fellow of the Association of Computing Machinery (ACM).

Brian Athey, Ph.D.is a Collegiate Professor and Inaugural Chair of the Department of Computational Medicine and Bioinformatics at the University of Michigan Medical School. He is also a Professor of Psychiatry and of Internal Medicine. He is the founding Principal Investigator of the NIH National Center for Integrative Biomedical Informatics (NCIBI), one of eight NIH National Biomedical Computing Centers, funded by the National Institute on Drug Abuse (NIDA) and the NIH Common Fund. Brian serves as US Academic lead and Co-CEO of the tranSMART Foundation, a non-profit company founded to coordinate the continued development of the open source tranSMART code base which supports an integrated open data sharing and analytics platform used world-wide to accelerate clinical and translational research. Brian has led the National library of Medicine (NLM) Next-Generation Internet (NGI) Visible Human Project and the DARPA Virtual Soldier Project. He has been a national leader in the NIH Clinical and Translational Scientists (CTSA) informatics Key Function Committee and U-M CTSA Informatics lead.

**Program Director Panelists**

Evelyn Goldfield is a Program Director of the Chemistry Division of the National Science Foundation and Professor in the Department of Chemistry, Wayne State University at Detroit. Her research interests include Quantum Dynamics: Methods and Applications, Dynamics of Chemical Reactions, Chemical Reactions in Confined Environments, High Performance Computing for Quantum Dynamics, Parallel Computing.

Daniel S. Katz is a Program Director of the Division of Advanced Cyberinfrastructure of the National Science Foundation. He is on leave from his position as a Senior Fellow in the Computation Institute, University of Chicago and Argonne National laboratory. He is also an affiliate faculty member at the Center for Computation and Technology (CCT) and an Adjunct Associate Professor in the Department of Electrical and Computer Engineering at Louisiana State University (LSU). He is interested in the development and use of advanced cyberinfrastructure to solve challenging problems at multiple scales. His technical interests include applications, algorithms, fault tolerance, and programming in parallel and distributed computing. His policy interests include citation and credit mechanisms and practices associated with software and data, organization and community practices for collaboration, and career paths for computing researchers.

Peter McCartney is a Program Director in the Division of Biological Infrastructure, National Science Foundation. He oversees research funded under several programs including Biological Databases and Informatics, National Ecological Observatory Network, Cyberinfrastructure for Environmental Observatories, Field Stations and Marine Laboratories, and Assembling the Tree of Life. Prior to NSF he was a Research Professor in the Global Institute of Sustainability at Arizona State University where he directed projects related to information systems for environmental and archaeological research use of metadata for designing automated internet access to data and applications; and workflow processing tools for incorporating multiple models into comprehensive analyses.

Thomas Russell is a Senior Staff Associate, Office of the Assistant Director, Mathematical and Physical Sciences, National Science Foundation. His primary responsibilities in the MPS OAD include leadership roles in Data Way and in the Cyberinfrastructure Framework for 21st Century Science and Engineering (CIF21) OneNSF investment, as well as continued major contributions to the Integrated NSF Support Promoting Interdisciplinary Research and Education (INSPIRE) One NSF investment. His research interests are in the numerical solution of partial differential equations, particularly with applications to subsurface flows in porous media, including groundwater flow and transport and petroleum reservoir simulation. His current major thrusts include control-volume mixed finite element methods, which compute accurate velocities/fluxes for flow equations in heterogeneous media on distorted meshes; Eulerian-Lagrangian localized adjoint methods, which compute accurate solutions for transport equations, even when advection-dominated; efficient algebraic equation solvers for these methods; and up scaling techniques based on stochastic models and the solution of moment equations.

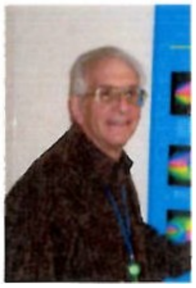

Dr. Barry I. Schneider is a Program Director for the National Science Foundation's Office of Cyberinfrastructure, specifically for the eXtreme Digital (XD) program. He received his Bachelors in Chemistry from Brooklyn College, his Masters in Chemistry from Yale University and his PhD in Theoretical Chemistry from the University of Chicago. Before coming to the NSF, he worked at Los Alamos National Laboratory (LANL) in the Theoretical Division, at GTE Laboratories as a member of the Technical Staff and since 1992 has held visiting appointments at LANL and at the National Institute of Standards and Testing (NIST).

**Workshop Attendees**

Lorena A. Barba is an Associate Professor of Mechanical and Aerospace Engineering at the George Washington University, in Washington DC. Her research interests include computational fluid dynamics, especially immersed boundary methods and particle methods for fluid simulation; fundamental and applied aspects of fluid dynamics, especially flows dominated by vorticity dynamics; fast algorithms, especially the fast multipole method and its applications; and scientific computing on GPU architecture. http://lorenabarba.com

Nico Cellinese is an Associate Curator at the Florida Museum of Natural History, University of Florida and a Joint Associate Professor in the Department of Biology. She is primarily interested in the systematics, evolution and biogeography of flowering plants. Additionally, part of her research revolves around tool development to facilitate biodiversity data synthesis and analysis. http://cellinese.blogspot.com

Alberto Cerpa was one of the three founding faculties of the Electrical Engineering and Computer Science program in the School of Engineering at UC Merced when he joined in 2005. His interests lie broadly in the computer networking and distributed systems areas, with recent focus in systems research in wireless sensor networks. Alberto is a recipient of the NSF CAREER Award (2013). http://www.andes.ucmerced.edu/~acerpa/

Diego Donzis is an Assistant Professor in the Department of Aerospace Engineering at Texas A&M University. He is interested in large scale computing, fluid mechanics, turbulence and turbulent mixing in incompressible and compressible flows. He obtained his PhD at Georgia Tech, and worked in the University of Maryland and the International Center for Theoretical Physics (Italy) before joining the faculty at Texas A&M.
http://aero.tamu.edu/people/faculty/?id=529

Benchun Duan is a faculty member in the Department of Geology and Geophysics at Texas A&M University. He is interested in earthquake physics and computational geophysics. He investigates factors and processes that control large earthquake rupture processes, near-field ground motion and deformation. Finite element method and parallel, high-performance computing are technical aspects in his research. His NSF Career project investigates controls on megathrust earthquakes along the Japan Trench subduction zone. http://geoweb.tamu.edu/profile/BDuan

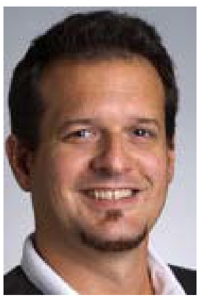

Adrian Feiguin joined Northeastern University as Assistant Professor in 2012, after spending 3 years as Assistant Professor at the University of Wyoming. His field of expertise is computational condensed matter, focusing on theoretical and computational aspects of low dimensional strongly interacting quantum systems. This physics is realized under extreme conditions, such as very low temperatures, high pressure, or high magnetic fields, and low spatial dimensions, and it is mostly governed by the collective behavior of the electrons inside a solid. http://www.northeastern.edu/afeiguin/

Antonino Ferrante is an Assistant Professor in the William E. Boeing Dept. of Aeronautics & Astronautics at the University of Washington, Seattle. He is recipient of the NSF CAREER (2011), U.S. National Academy of Sciences ICTAM Travel Award (2012), and Royalty Research Fund Award (2012) from the University of Washington. His research interests are in fluid mechanics, multiphase turbulent flows, high-speed turbulent flows, and chemically-reacting flows. His research tools are direct numerical simulation, large-eddy simulation, and high performance computing.
http://www.aa.washington.edu/faculty/ferrante/

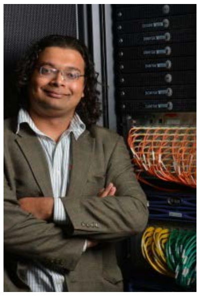

Baskar Ganapathysubramanian is an Assistant Professor of Mechanical Engineering and Electrical and Computer Engineering at Iowa State University. His research interests are in multi-scale multi-physics modeling, design of materials and processes using computational techniques, and stochastic analysis. The recent focus of his group is on advanced energy technologies including solar cells, and green buildings. Ganapathysubramanian completed his PhD and MS from Cornell University and holds a BS degree from the Indian Institute of Technology-Madras. http://www3.me.iastate.edu/bglab/

Kai Germaschewski is an Assistant Professor at the University of New Hampshire's Department of Physics and Space Science Center. His work focuses on in large-scale computer simulations of plasmas, with applications to the Earth's space environment and laboratory plasmas. He works with both fluid (extended MHD) and kinetic (using particle-in-cell) plasma models. He is interested in modern aspects of computing, like adaptive mesh refinement and implicit time integration, and heterogeneous architectures (GPUs, Intel MIC). http://www.eos.unh.edu/Faculty/kgermaschewski

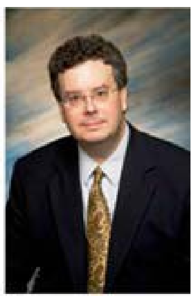

Thomas Hacker is an Associate Professor of Computer and Information Technology at Purdue University and Visiting Professor in the Department of Electrical Engineering and Computer Science at the University of Stavanger in Norway.  Dr. Hacker's research interests center around high- performance computing and networking on the operating system and middleware layers. Recently his research has focused on cloud computing, cyberinfrastructure, scientific workflows, and data-oriented infrastructure. https://tech.purdue.edu/profile/tjhacker

Dr. Ann Jeffers is an Assistant Professor in the Department of Civil and Environmental Engineering at the University of Michigan. Her research lies at the intersections between the fire sciences and structural engineering disciplines, and specifically seeks to establish novel computational methods that bridge the domains of fire science, heat transfer, and structural mechanics. She currently serves on the ASCE Fire Protection Committee and the SFPE Standards Making Committee on the Predicting the Thermal Performance of Fire Resistive Assemblies. http://www-personal.umich.edu/~jffrs

Shantenu Jha is an Assistant Professor at Rutgers University, and a Visiting Scientist at the School of Informatics (University of Edinburgh) and at University College London. His research interests lie at the triple point of Applied Computing, Cyberinfrastructure R&D and Computational Science. Shantenu is the lead investigator of the SAGA project (http://www.saga-project.org), which is a community standard and is used to support science and engineering applications on most major production distributed cyberinfrastructure -- such as US NSF's XSEDE and the European Grid Infrastructure. http://radical.rutgers.edu

Kapil Khandelwal is an Assistant Professor in the Department of Civil & Environmental & Earth Sciences at the University of Notre Dame. He received BS in Civil Engineering from IIT-Roorkee (India), MS in Structural Engineering from IIT-Delhi (India) and Ph.D. in Civil Engineering from the University of Michigan, Ann Arbor. His research interested includes: computational solid mechanics (FEM), gradient elasticity/plasticity, computational fracture mechanics, topology optimization, and progressive collapse of structural systems. http://www3.nd.edu/~kkhandel/

Xiaolin (Andy) Li is an Associate Professor in Department of Electrical and Computer Engineering at University of Florida. His research interests include Parallel and Distributed Systems, Cyber-Physical Systems, and Network Security & Privacy. He is directing the Scalable Software Systems Laboratory (S3Lab). He is in the executive committee of IEEE Technical Committee of Scalable Computing (TCSC) and the coordinator of BigData & MapReduce and Sensor Networks He received a Ph.D. in Computer Engineering from Rutgers University. He is a recipient of the National Science Foundation CAREER Award 2010 and a member of IEEE and ACM. http://www.andyli.ece.ufl.edu/

Laurence Loewe is an Assistant Professor at the University of Wisconsin-Madison. He investigates questions in the new field of evolutionary systems biology, which merges systems biology and population genetics. He is interested in bridging the gap between simple analytically understandable mathematical models and biological reality by building rigorous simulation models to answer various evolutionary questions. http://evolution.ws/people/loewe

Kamesh Madduri is an Assistant Professor in the Computer Science and Engineering Department at The Pennsylvania State University. He received his PhD in Computer Science from Georgia Institute of Technology's College of Computing in 2008, and was previously a Luis W. Alvarez postdoctoral fellow at Lawrence Berkeley National Laboratory. His research interests include high-performance computing, parallel graph algorithms, and massive scientific data analysis. He is a member of IEEE, ACM, and SIAM.
http://www.cse.psu.edu/~madduri/

Alison Marsden is an Associate Professor and Jacobs Faculty Fellow in the Mechanical and Aerospace Engineering Department at the University of California San Diego. Her work focuses on the development of numerical methods for simulation of cardiovascular blood flow problems, medical device design, application of optimization to fluid mechanics, and use of engineering tools to impact patient care in cardiovascular surgery and congenital heart disease.
http://maeresearch.ucsd.edu/marsden/AMarsden/Home.html

Dr. Pompili is an Associate Professor at Rutgers University-New Brunswick, where he is the director of the Cyber-Physical Systems Laboratory (CPS Lab), the site co-director of the NSF Center for Cloud and Autonomic Computing (CAC), and the associate director of application collaborations of the Rutgers Discovery Informatics Institute (RDI2). In 2011, Dr. Pompili was awarded an NSF CAREER Award to design efficient communication solutions for underwater multimedia applications. His research spans underwater acoustic communication and coordination of vehicles, ad hoc and sensor networks, thermal management of datacenters as well as mobile and green computing.
http://www.ece.rutgers.edu/~pompili/

Dr. Ioan Raicu is an Assistant Professor in the Department of Computer Science (CS) at Illinois Institute of Technology (IIT). He has received the prestigious NSF CAREER award (2011 - 2015) for his innovative work on distributed file systems for exascale computing. His research work and interests are in the general area of distributed systems. His work focuses on a relatively new paradigm of Many-Task Computing (MTC), which aims to bridge the gap between two predominant paradigms from distributed systems, High-Throughput Computing (HTC) and High- Performance Computing (HPC).
http://www.cs.iit.edu/~iraicu/

Pradeep Ravikumar is now an Assistant Professor in the Department of Computer Science, at the University of Texas at Austin. He is also affiliated with the Division of Statistics and Scientific Computation, and the Institute for Computational Engineering and Sciences at UT Austin. His thesis has received honorable mentions in the ACM SIGKDD Dissertation award and theCMU School of Computer Science Distinguished Dissertation award. He is also a recipient of the NSF CAREER Award. http://www.cs.utexas.edu/~pradeepr/

Suzanne Shontz is an Assistant Professor in the Department of Mathematics and Statistics at Mississippi State University. Suzanne's research is in parallel scientific computing and focuses on the development of meshing and numerical optimization algorithms and their applications to medicine, image processing, and electronic circuits. Suzanne is the recipient of a 2011 NSF CAREER Award and a 2011 NSF PECASE Award from President Obama for her research in computational- and data-enabled science and engineering. Along with Thomas Hacker of Purdue University, she is a Co-Chair of the 2012 and 2013 NSF CyberBridges Workshops. http://sshontz.math.msstate.edu

Dr. Melissa Smith was a research associate at the Oak Ridge National Laboratory (ORNL) for 12 years before her appointment to Clemson in 2006. Her research group is interested in the performance computer architectures for various application domains including scientific applications (modeling and simulation), high-performance or real-time embedded applications, and medical and image processing. Her group explores optimization techniques and performance analysis for emerging heterogeneous platforms, including many processors, Graphical Processing Units (GPUs) and Field Programmable Gate.
http://www.clemson.edu/ces/departments/ece/faculty_staff/faculty/msmith.html

Manoj Srinivasan is an Assistant Professor in the Department of Mechanical and Aerospace Engineering at The Ohio State University. His recent research focused on the understanding of human locomotion and biomechanics from the perspective of optimal control and dynamical systems theory. Srinivasan received an undergraduate degree in engineering from the Indian Institute of Technology, Madras, received a doctoral degree in Theoretical and Applied Mechanics at Cornell University, and was a post-doctoral researcher and lecturer at Princeton University. He is an NSF CAREER award recipient. http://movement.osu.edu

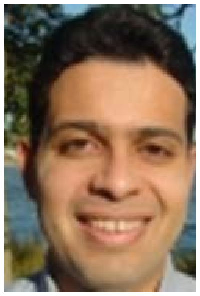

Andrés Tejada-Martínez is Associate Professor in Civil and Environmental Engineering at University of South Florida. Tejada-Martínez has received an NSF CAREER Award and various others NSF collaborative research awards for his work in large-eddy simulations of turbulent mixing in shallow shelf coastal regions and in the upper ocean mixed layer.
[www.eng.usf.edu/~aetejada](http://www.eng.usf.edu/~aetejada)

Richard (Rich) Vuduc is an Associate Professor at the Georgia Institute of Technology ("Georgia Tech") in the School of Computational Science and Engineering. His research lab, the HPC Garage (hpcgarage.org), is interested in all-things-high-performance-computing, with an emphasis on parallel algorithms, performance analysis, and performance tuning. He is a recipient of the NSF CAREER Award. His lab's work has received a number of best paper nominations and awards including most recently the 2012 Best Paper Award from the SIAM Conference on Data Mining.
[http://www.cse.gatech.edu/people/richard-vuduc](http://www.cse.gatech.edu/people/richard-vuduc)

Dr. Jun Wang is an Associate Professor in Department of Electrical Engineering and Computer Science at the University of Central Florida, Orlando, FL, USA. He has conducted extensive research in the areas of Computer Systems and High Performance Computing. His specific research interests include data-intensive high performance computing, massive storage and file system, I/O Architecture, and low-power computing.
[http://www.eecs.ucf.edu/~jwang](http://www.eecs.ucf.edu/~jwang)

Liqiang (Eric) Wang is a Castagne Associate Professor in the Department of Computer Science at the University of Wyoming. His research interest is the design and analysis of parallel systems for big-data computing, which includes two aspects: design and analysis. For design, he is currently working on optimizing performance, scalability, resilience, and load balancing of data-intensive computing, especially on Cloud, GPU, and multicore platforms. For the aspect of analysis, he focuses on using program analysis to detect programming errors and performance defects in large-scale parallel computing systems. He received an NSF CAREER Award in 2011.
[http://www.cs.uwyo.edu/~lwang7/](http://www.cs.uwyo.edu/~lwang7/)

Shaowen Wang is an Associate Professor of Geography and Geographic Information Science (Primary), Computer Science, and Urban and Regional Planning at the University of Illinois at Urbana-Champaign (UIUC). His research and teaching interests center on three interrelated themes: 1) computational theories and methods in geographic information science, 2) cyberinfrastructure and data-intensive computational science, and 3) multi-scale geospatial problem solving and spatiotemporal synthesis. He was a visiting scholar at Lund University sponsored by the National Science Foundation (NSF) in 2006 and NCSA Fellow in 2007, and received the NSF CAREER Award in 2009. http://www.cigi.illinois.edu/shaowen/

Dr. Lei Wu is an Assistant Professor of ECE Department at Clarkson University. He has experience working with NYISO, GE, and Siemens Energy Automation on various power system studies. He has extensive publications on power systems research and serves on the Research Committee of the Clarkson's Honors program and is a member of the Center for Sustainable Energy Systems. He is an Editor of IEEE Transactions on Sustainable Energy and Guest Editor of IEEE Transactions on Smart Grid. His educational and research activities are supported by grants from NSF, DOE, GE, and IBM. http://people.clarkson.edu/~lwu/

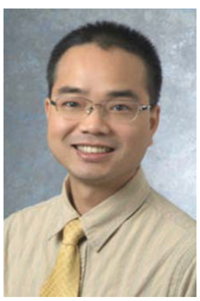

Dr. Xiong (Bill) Yu is an Associate Professor at the Department of Civil Engineering, Case Western Reserve University. His research interest is in the broad area of geoengineering related to infrastructure sustainability, environment and energy needs. His work embraces innovative sensors and materials to improve sustainability and intelligence of the civil infrastructure systems. He is the PI of over 25 research projects sponsored by the National Science Foundation , National Research Council, Ohio DOT, Federal Highway Administration, NCHRP-IDEA and other agencies such and private industry. http://filer.case.edu/xxy21/Index.html

Jessica Zhang is an Associate Professor in Mechanical Engineering at Carnegie Mellon University with a courtesy appointment in Biomedical Engineering. Her research interests include computational geometry, mesh generation, computer graphics, visualization, finite element method, isogeometric analysis and their application in computational biomedicine and engineering. She is the recipient of a 2012 NSF CAREER Award. http://www.andrew.cmu.edu/user/jessicaz/

## Appendix C. Poster Session

The poster session provided a forum for attendees to present work from their CAREER projects and facilitated connections and collaborations among researchers. All of the workshop participants presented posters. Poster topics included high-performance computing, parallel computing and distributed storage systems, fault tolerance and error detection, and scalability, to name just a few. Several researchers presented applied work where high-performance computing was being used to develop models and algorithms to address questions in various fields such as materials and structural engineering, fluid dynamics, biology, and medicine. The posters presented are listed in Table 1.

**Table 1**. Posters presented at the Cyberbridges 2013 conference.

| Name | Institution | Poster Title |
|---|---|---|
| Lorena Barba | George Washington University | Open software, GPU computing and flipped classroom |
| Nico Cellinese | Florida Museum of Natural History | TOLKIN-Workflow: a web application for conducting and managing complex research pipelines |
| Alberto Cerpa | University of California-Merced | No poster title given |
| Diego Donzis | Texas A&M University | Discoveries in compressible turbulence and turbulent mixing through petascale simulation and analysis |
| Benchun Duan | Texas A&M University | High performance computing and visualization in earthquake modeling research and education |
| Adrian Feiguin | Northeastern University | Non-equilibrium quantum dynamics in strongly correlated systems |
| Antonio Ferrante | University of Washington | Petascale DNS of evaporating droplet-laden homogenous turbulence. |
| Baskar Ganapathysubramanian | Iowa State University | A continuum-mechanics based, cyber-enabled approach to unraveling process-structure-property relationships in organic electronics |
| Kai Germaschewski | University of New Hampshire | Studies of 3D dynamics in the global magnetosphere using high-performance heterogeneous computing architectures |
| Thomas Hacker | Purdue University | Fault analysis and reliability guided scheduling for large-scale HPC systems |
| Ann Jeffers | University of Michigan | Traveling fires- do they really matter? |
| Shantenu Jha | Rutgers University | Developing middleware to support distributed dynamic data-intensive (D3) science on distributed cyberinfrastructure |
| Kapil Khandelwal | University of Notre Name | Topology optimization of linear and nonlinear multiscale systems |

| | | |
|---|---|---|
| Xiaolin Li | University of Florida | SMART: Scalable adaptive runtime management algorithms and toolkit |
| Laurence Loewe | University of Wisconsin-Madison | Modeling made easy |
| Kamesh Madduri | The Pennsylvania State University | Algorithmic and software foundations for large-scale graph analysis |
| Alison Marsden | University of California, San Diego | Optimization and parameterization for multiscale cardiovascular flow simulations using high performance computing |
| Dario Pompili | Rutgers University | CAREER: Investigating fundamental problems for underwater multimedia communication with application to ocean exploration |
| Ioan Raicu | Illinois Institute of Technology | Distributed storage systems for extreme-scale data-intensive computing |
| Pradeep Ravikumar | University of Texas at Austin | Statistical machine learning and big-p, big-n, complex data |
| Suzanne Shontz | Mississippi State University | Parallel dynamic meshing algorithms, theory, and software for patient-specific medical interventions |
| Melissa Smith | Clemson University | Multi-level performance modeling for heterogeneous petascale systems and beyond |
| Manoj Srinavasan | The Ohio State University | Towards an optimization-based and experimentally verified predictive theory of human locomotion |
| Andres Tejada-Martinez | University of South Florida | LES of full-depth Langmuir circulation in a tidal current |
| Richard Vuduc | Georgia Institute of Technology | How much time, energy, and power will my algorithm require? |
| Jun Wang | University of Central Florida | Enabling data-intensive HPC analytics for interdisciplinary community |
| Liqiang Wang | University of Wyoming | Scalable communication |
| Shaowen Wang | University of Illinois | CyberGIS for enabling data-intensive geospatial discovery and innovation |
| Lei Wu | Clarkson University | Stochastic multiple time scale co-optimized resource planning of future power systems with renewable generation, demand response, and energy storage |
| Xiong (Bill) Yu | Case Western Reserve University | Multiscale sensing and simulation for bridge scour |
| Jessica Zhang | Carnegie Mellon University | Multi-core CPU and GPU-accelerated multiscale modeling for biomolecular complexes |

## Appendix D. Collaboration Activities at the Workshop

Before the workshop, attendees were asked to provide a list of areas of interest in which they were seeking collaborators. These were included in the workshop program to help attendees identify potential collaborators. At the conclusion of the workshop, attendees were asked to complete a *Collaboration Card* identifying the areas of potential collaboration that they identified with colleagues at the workshop. A sample of the areas of interest and identified collaborations is summarized in Table 2.

**Table 2.** Sample of Areas of Interest and completed *Collaboration Cards*. Attendees listed both potential collaborators and the potential areas of collaboration.

| Researcher | Areas of Interest for Collaboration | Potential Collaborators | Areas of Collaboration Identified |
|---|---|---|---|
| Lorena Barba | Fast algorithms of the FMM family, developing benchmarks to help the community evaluate algorithmic innovations and new implementations, solvers and preconditioners for elliptic PDS where FMM may play a role, applications of structure-based energy methods in biomolecular physics where bioelectrostatics solvers use FMM as a numerical engine. | Antonio Ferrante, Diego Donzis, Rich Vuduc, Jessica Zhang | CFD modules that can be shared in courses, building a community for FMM algorithms, high performance computing for biomolecular modeling. |
| Kamesh Madduri | Big data analytics and mining applications, scalable data management and visualization. | Xiaolin Li, Rich Vuduc, Jessica Zhang, Baskar Ganapathysubramanian | Mesh partitioning and data mining, runtime systems, performance modeling and analysis, spatial data structures |
| Dario Pompili | Mobile grid computing, underwater communications, coordination of vehicles, cloud-assisted robotics, wearable computing | Andres Tejada-Martinez, Shantenu Jha | Distributed computing |
| Shaowen Wang | Advanced cyberinfrastructure, data-intensive geospatial sciences and technologies, scalable computing and information systems, sustainability science | Jun Wang, Xiaolin Li, Shantenu Jha | Advanced cyberinfrastructue, indoor location data analysis |
| Benchun Duan | Hybrid MPI/OpenMP parallelization of EQdyna, FEM mesh generation of complex geological models, including non-planar fault geometry, topography, and complex velocity substructure, visualization of 3D modeling results, visualization of earthquake generation processes in the 3D Earth for course modules. | Liqiang Wang, Jun Wang | Parallel I/O |

| | | | |
|---|---|---|---|
| Baskar Ganapathysubramanian | Parallel adaptive mesh generation, data-mining. | Alberto Cerpa, Adrian Feiguin | Multiscale materials modeling, green buildings |
| Adrian Feiguin | Quantum information, quantum chemistry. | Baskar Ganapathysubramanian | Fullerene-polymer based light harvesting devices. |
| Shantenu Jha | All domains of science and engineering that entail distributed and high performance computing. | Ann Jeffers, Shaowen Wang | Frameworks for multiphysics multicomponent simulation, GIS |
| Melissa Smith | Big data applications, heterogeneous HPC system users, and performance modeling and analysis | Andy Li, Rich Vuduc, Eric Wang. | Scalability, performance modeling and heterogeneous systems, performance optimization and load balancing |
| Jessica Zhang | Computational biology, finite element applications | Benchun Duan, Lorena Barba, Xiong Yu | Geology and meshing, bioapplications, bridge scouring and meshing CFD |
| Jun Wang | Data-intensive HPC applications, software defined networks, GPU computing, memory architecture | Xiaolin Li, Shaowen Wang, Benchun Duan, Liqiang Wang | Software defined networks, data intensive computing for GIS, parallel I/O, GPU |
| Andy Li | Not listed | Antonio Ferrante, Jun Wang, Shaowen Wang | Indoor GIS, LBS, CFD, SAMR, CDS&E, aeronautics, continuous detonation engine, SDN, storage/DFS |
| Liqiang Wang | Scalability of large-scale linear equation solvers, automatic load balancing of scientific computations, automatic cloud provisioning for HPC, storage and I/O optimization for big-data computing. | Benchun Duan, Xiong Yu, Melissa Smith | Parallel I/O, optimizing parallel computations, performance modeling |
| Rich Vuduc | Scalable applications and libraries, programming models, and computer architecture. | Diego Donzis, Lorena Barba, Antonio Ferrante, Melissa Smith | Performance tuning for GPUs, FMM community building, FFTs on GPUs, performance modeling minisymposium @ SIAM PP14 |

| | | | |
|---|---|---|---|
| Antonio Ferrante | Petascale elliptic solvers or fast Poisson solvers/parallel FFTs, optimized massive MPI communications, exascale computing: hybrid multicore GPUs, parallel I/O and visualizations of petascale datasets, code optimization on Blue Waters. | Xiaolin Li | Continuous detonation engine, AMR, CFD |
| Andres Tejada-Martinez | Finite difference, finite elements and spectral methods for fluid dynamics, stratified flows, parallel computing. | Dario Pompili, Antonio Ferrante | AUVs for fluid dynamics measurements in the ocean, volume of fluid methods for tracking interfaces in fluid dynamics |
| Ioan Raicu | Data-intensive computing applications, data provenance, large scale workflow based applications. | Xiaolin Li, Shantenu Jha | Cloud computing, targeting CSR programs, big data, data aware scheduling, job management systems. |
| Laurence Loewe | Bridging the gap between simple analytically understandable mathematical models and biological reality. | Alison Marsden, Baskar Ganapathysubramanian | How to model space in the Evolvix model, easy interface to PDEs |
| Suzanne Shontz | Mathematical modeling, model order reduction, scientific visualization, patient data sets. | Laurence Loewe , Alison Marsden, Kapil Khandelwal, Baskar Ganapathysubramanian | Mesh generation, adaptive meshing, vascular meshes, optimization of biomedical fluids |
| Alison Marsden | Image segmentation, machine learning algorithms, uncertainty quantification | Jessica Zhang, Baskar Ganapathysubramanian | Meshing, segmentation, UQ |

## Appendix E. Notes from Breakout Sessions

**William Gropp**

Discussion questions and breakout group comments

**What is the right balance between HPC and other computing infrastructure?**
**Scribe: Laurence Loewe**

The group discussed HPC vs. Cloud vs. Long-tail discussion, and concluded that a balance is needed. In their discussion, they found that it is important for research productivity and workforce development to have a balanced mix of computing resources available. The group identified the following tradeoffs:

1. Type of computing platform: high performance computing systems versus a cloud computing system.

In the use of high performance computing, some tasks are best run on high performance supercomputers, where most of the cost is not in the nodes but rather in the network (e.g. to facilitate MPI). Other tasks do not need such a fast network, and would be a good fit for a cloud computing system rather than a high performance computing system. These types of applications tend to belong to the pleasantly parallel universe, where high throughput approaches like Condor give the best performance / price ratio

2. Geographic location of computing resources: local versus remote.

HPC resources needed by the community are a mix of smaller local systems and larger remote systems. There will always be some need for local machines for testing runs, software development, etc. The quick availability and flexibility of local computing resources is much better compared with larger remote computational services. At the other end of the spectrum, larger computational, jobs should be “farmed out” to larger computer centers with larger systems. The economy of scales allow for much more efficient management (e.g. professionals managing computers rather than graduate students, who should be focused on their research).

3. High CPU vs High RAM density

Another continuum is a high computational core count versus high memory (RAM) density. Some computational tasks require many CPUs, but not much memory. These types of jobs are well serviced by the architectures available today. On the other hand, some computational tasks require a large amount of memory or large memory space (sometimes even without allowing for any parallel computing). Today it can be challenging to find such high -RAM machines in sufficient quantity to handle larger CPU loads that also require a large memory.

To enable students to collect relevant experiences with this broad range of computing approaches, we need a diverse array of compute resources that are readily accessible. To respond to these needs NSF has developed the following resources: 1) XSEDE caters to a broad range of large scale computing needs and provides a substantial number of compute hours to any funded NSF proposal at the mere effort of an online account registration; more can be requested as needed, however competition for CPU hours can be severe. 2) NSF welcomes PIs to put small compute resources (e.g. $15K server) on grant proposals, however, this is not designed to, for example, fund a $150K cluster (which should not be run by a PI anyway, but rather by a University - large equipment grant or centrally by the NSF).

In the discussion, several weaknesses surfaced in the current computational environment available to the research community. First, some PI's felt that the diversity of current compute resources could be increased. Second, due to a (perceived?) lack of other resources, sometimes the supercomputers are misused for jobs that could be run at much lower cost on massively parallel clusters of smaller machines with a slow (and cheap) network. Third, for computational groups there is a (perceived?) need to have larger computational resources locally (and readily) available (and have a funding model to support such clusters). Finally, some computational work is undertaken when the team is facing stringent deadlines (such as a conference submission deadline), and having additional resources that allow bypassing a national computing queue could be of interest

**How do we bring computing to new areas of science?**
**Scribe: Lorena Barba**

To promote the adoption of high performance computing in other areas of science, it is important to make the technology easy to use, and to provide easy access to resources as real-time as possible. Specifically, what is needed is domain-specific high level languages (such as Sundance and FEniCS) that can work on one core, a workstation, a cluster, or an entire HPC system that can exploit the capabilities (such as GP-GPUs or Phi processors) available in each platform. Moreover, computing needs to be added more pervasively to the K-12 and college curriculums, and computational modeling needs to be introduced much earlier in the curriculum.

**How can we support innovative computer architecture research?**
**Scribe: Rich Vuduc**

The conventional wisdom is that, with respect to hardware, the commodity market drives HPC system architectures and designs. That may be breaking because of the disruptive switch to parallel architectures: the commodity approaches are no longer sufficient to sustain continued performance improvements; as such, design may be at least partially up for grabs. There is high demand for HPC -- new systems like Blue Waters are highly oversubscribed. Thus, it may be possible that user communities and/or NSF could exert at least some control and influence over the directions for new architectures.

A big question is how to communicate requirements to architects in an "actionable" way. The breakout group discussed several potential means of doing so.

- Form a national advisory committee that could help identify multiple communities that, individually, may have "subcritical" demand, but in aggregate have enough computational or algorithmic commonality to constitute critical demand. This committee would help this super-community translate their common needs into a form that might inspire new directions for architecture.
- For critical communities that already exist, NSF could (continue to) support workshops and projects that specifically engage architects.
- Continue to expand support for research that bridges applications and architectures (and other levels of the stack) a la the Exploiting Parallelism and Scalability (XPS) program.
- Universities should revisit the way they approach computational science education. That is, is the computational science curriculum giving students the right "tools" to communication application characteristics in a meaningful way to architects? Conversely, can we encourage architects to not (in general) treat applications as "black-boxes?"

**What Big Data problems need HPC?**
**Scribe: Kamesh Madduri**

The group first identified computational characteristics of big data problems (> 10 PB) that would require supercomputers (e.g., Tianhe-2, Blue Waters). Of the three V's associated with big data, the group agreed that Velocity is an important descriptive feature for such problems: these problems have a real-time analysis requirement, are latency sensitive and just not throughput-oriented, and/or involve periodic, incremental updates of new data to larger data stores. Other characteristics of these problems may be that: 1) data is difficult to partition or split, or naive data decomposition schemes might involve significant load imbalances in parallel execution; 2) computations that involve irregular or random accesses to large data sets; and 3) analytics algorithms that are challenging to parallelize, i.e., might require sophisticated data structures, have irregular data-dependent concurrency in the computational routines, etc.
The group then identified several types of applications with the above characteristics:

- The process of building and validating weather forecasting models using big data - finer spatial and temporal resolutions along with incremental addition of data are challenges.
- Analysis requirements of petascale simulations (e.g., molecular dynamics) that already run on supercomputers. Performing analysis in-situ, or on the same supercomputer after execution, would minimize data movement.
- Simulations with established post-processing data analysis workflows, or analysis of massive experimental data (LHC, astrophysics) with complex pipelines. These would again benefit from tightly-coupled analyses on supercomputers.

- Smart Grid monitoring and Cybersecurity.
- Realtime monitoring of the structural health of critical infrastructure from sensor data.
- Monitoring high-frequency trading patterns (e.g., detecting the 2010 flash crash)
- Personalized medicine (e.g., genome sequencing and comparative analysis at a doctor's office), ICU monitoring, computer-assisted surgery

The group also identified several types of Big Data applications that do not require supercomputers:

- Pleasantly-parallel and throughput-oriented jobs that do not utilize/require the high-bandwidth, low-latency networks of supercomputers.
- Analysis of 10's of TB of data. This can be done on campus clusters or commercial clouds.
- Machine learning methods that can be expressed in the MapReduce paradigm and related languages and frameworks.
- Business analytics for which there are commercial solutions from SAS, Oracle, etc.

In considering the question: what are the main impediments to the use of supercomputers for Big Data? The group felt that cost was a big consideration. It may be cheaper and easier to build a Hadoop cluster or use a commercial cloud, than to acquire and program a supercomputer. The second reason brought up was that applications with realtime analysis requirements might need a dedicated parallel system, whereas jobs need to be submitted to shared supercomputers with possibly long turnaround times.

**Brian Athey**

Discussion questions and breakout group comments

**Data Life Cycle Management and maintaining data and information after the grant or contract end are critical to sustaining data and information intensive science? (which is now all science). How can the community work together to maintain this "big data" for the long term?**
**Scribe: Manoj Srinivasan**

(Data lifecycle management)

The group provided a brief report on their discussions. Participants included Nico Cellinese, Alison Marsden, Lorena Barba, Manoj Srinivasan, Baskar Ganapathysubramanian, Yongjie Jessica Zhang. In addressing the question, one basic issue seems to be that funding agencies, such as NIH and NSF, that invest millions of dollars in fundamental research that produce both data and software, do not seem to have simple mechanisms for supporting the long-term maintenance and management of such products -- long-term meaning 'long after the project funding period is over', perhaps tens of years. (The group learnt later that NIH does support R01 for long-term software maintenance, IF the software is used a lot by the community.)

There exist public databases such as Dryad, Figbank, Morphbank, genebank, etc., some supported by US agencies. Some such free public databases are supported by soft or temporary funding -- so it is unclear what the long-term sustainability of some such databases will be. On the other hand, it was mentioned that databases such as genebank are supported strongly by NIH, and the group wondered if it was because of the perceived value of such genomic data. While it seems reasonable to prioritize data management support based on perceived value of the data, one should ideally have a broader support for management of all data supported by public (or indeed private) funds.

The group discussed some mechanisms for supporting such databases. Proposals that propose to use data from a particular database might have a small part of their request allocated to support the database. (Perhaps the distributed maintenance of highly used datasets can subsidize some less or not at all used datasets in the same database?)

Perhaps the community could use user subscription models for public databases, either at the level of investigator, or at the level of institutions --- but perhaps this method will make the databases less affordable, for the smaller and/or less well-endowed organizations, other less affluent countries (say), and most importantly, the general public. (It becomes like expensive journal subscriptions.). Perhaps the for-profit journals that require the uploading of data into Dryad or other databases, which do benefit from the public support of science, could contribute somewhat to the support of such databases?

While one might argue that all the different stake-holders in the long-term maintenance of the data should be asked to contribute to its maintenance -- but perhaps one does not know future stake-holders derived from as yet unknown uses of the data. Also, we might want to save some data for posterity, even though there are no current stake-holders interested in paying for its maintenance -- after all, we have been maintaining journal articles and books for hundreds of years now, even though we use only a very small fraction of this collected information.

Also, in the long-term, it may be necessary to decide what to store and what to delete. This seems a difficult question to answer objectively, and might have to be decided on a case by case basis.

**<u>Steve Gordon</u>**

Discussion questions and breakout group comments

**What strategies might work best for integrating computational science into the curriculum in the face of limitations on total credit hours for a degree?**
**Scribe: Shaowen Wang**

The group discussed strategies for developing education and training programs of data and computational sciences. Three specific strategies were identified: 1) add new data and computational sciences courses to existing education programs; 2) replace existing courses of other areas with data and computational sciences courses; and 3) adapt or modify existing courses to include modules for data and computational sciences.

For campuses that currently do not have any formal program for education of data and computational sciences, a holistic approach is desirable for implementing curricula by leveraging existing resources (e.g., from http://www.hpcuniversity.org/). In the case that a campus has already had existing courses or programs, it is important to recognize the multidisciplinary nature of this area, and promote a community-driven approach to flexibly applying the aforementioned three strategies.

For establishing a sustainable education program for data and computational sciences, it is crucial for a campus to have grass-roots-level engagement of faculty interests across disciplines while high-level support from university administration is necessary.

**How can we revise the instruction in large, interdisciplinary lecture classes to integrate inquiry-based learning?**
**Scribe: Kai Germaschewski**

The group discussed approached to improve large lecture classes, and distilled a list of a number of active learning techniques:

- Research in physics education, which has by now some decent literature on similar issues shows that large lecture classes are rather ineffective in helping students learn.
- "Peer instruction" (Erik Mazur) where students are encouraged to discuss questions with their neighbor in class. Clickers or colored cards can be used to quickly scan the students' answers.

- "Studio" classes are typically smaller classes, a large lecture may be broken up into 2 or more studio sections (up to 90 students per section). Students work on activities to figure out new concepts and problems themselves.
- "Reverse classroom". Students watch the lecture online at home (or elsewhere), while actual class time is used for problem solving, group work

The group identified several barriers to switching to such techniques:

- Lack of familiarity with the new methods
- "Change is hard", also attitude of "the lecture classes worked fine for me, so they can't be that bad" (but future college professors are not an adequate sample of the general college population), and resistance to actually believe that those methods work (though there is research out there!). "It's easier to do a lousy job."
- P&T attitude is generally, "you just need to do ok in teaching", which is not much of an incentive to try to do better, if one really better worries about one's research standing, because that's what makes or breaks tenure.
- Lack of resources (studio rooms, large computer/physics labs, sufficient faculty/TA time).

The group identified several possible approaches that would help:

- The availability of pre-made materials that could be used in an instructors class (including in-class exercises, videos).
- Evaluations that actually evaluate learning, rather than how happy students are with their grades and how entertaining a given class was.
- Class release time to prepare a new active-learning course.
- Incentives from funding agencies (the ways in which they could do this is an open question)

**What kind of campus activities from projects like XSEDE can advance the computational engineering program in your campus?**
**Scribe: Jessica Zhang**

The group discussed the need for several types of campus activities. The first is courses on modeling and simulation. The second was the need for an interdisciplinary institute and program that could integrate multiple departments with a CSE program or minor that would encourage collaborative research. Finally, open courses or workshops that could provide modules on with applications and parallel libraries (such as OpenMP or MPI) would be helpful.

**Are there infrastructure barriers that inhibit the integration of computational modeling into instruction?**
**Scribe: Chandrajit Bajaj**

There are two fundamental barriers: a “topic” barrier, and barriers for instructors. In terms of a “topic” barrier, since the area of computational modeling encompasses foundational mathematics, computer science, physical sciences, and engineering, there are several inhibiting factors. First, in terms of the level of preparation of students, the requirements and prerequisites required are asking a lot of the current student population. Second, multiple course sequences which delve into theoretical and practical (lab based) are difficult at times as our individual course curriculum is already very full. Finally, students often prefer “easier” courses when computational modeling is not part of the core but is an elective. In terms of barriers for instructors, some disciplines find it easier while others find it extremely difficult to absorb material outside their discipline, due in part to the lack of training for the instructors. Moreover, there are inadequate recognition and rewards for instructors indulging in “new” course development for computational modeling. Finally, the lack of resources leads to a minimum required number of students needed to sign up for a specialized track sequence of courses.

**Summary of Breakout Session for Chandrajit Bajaj's Talk:**

**Question 1: What are some challenges and approaches to data sciences and visualization?**
**Scribe: Andy Li**

*No scribe notes submitted.*

**Question 2: What are some challenges and approaches to high performance computing and visualization?**
**Scribe: Melissa Smith**

There are several challenges in the arena of using high performance computing in conjunction with visualization. First, the resources for HPC and visualization are very different. Thus, data sets must be migrated from the HPC to storage to the visualization application, which is very inefficient. Second, the use of visualization tools is intimidating for some HPC users. In particular, the use of different visualization tools requires a significant effort to learn how to use them. In addition, some HPC users lack a basic understanding of visualization. Third, scalability issues prohibit the effective use of HPC and/or visualization to solve large science-enabling problems. Fourth, storage systems on HPC resources are not always adequate to support computation and visualization needs. Fifth, some applications could benefit from more support in regards to schedulers and runtime systems to support dynamic steering in conjunction with visualization. This could reduce the number of wasted cycles, potentially this could be reduced drastically. Finally, in situ computation and visualization can be challenging if the HPC system does not have a hybrid architecture, e.g., containing both GPUS and adequate per node storage.

**Question 3: What are some challenges and approaches in teaching computational visualization?**
**Scribe: Chandrajit Bajaj**

*No scribe notes submitted.*

**Question 4: What are some challenges and approaches in open source visualization software?**

**Scribe: Ann Jeffers**

Participants of this group studied the differences between the molecular community and other science and engineering fields that leads to a large repository of open source software and data. They noted that the molecular community has a culture of sharing that is not present in other scientific communities. There is also the incentive to share data since the molecular community requires large data sets for their research. In addition, sometimes sharing of data is mandatory.

The group identified several challenges for open source visualization software in this context. First, the quality of open source software in many fields is poor since the "best" research groups often do not share their software in order to main competitiveness. Second, there are IP issues that groups face when deciding whether or not to share their software. Third, whenever there is the possibility of commercialization of software, there is often no incentive to make the software open source. Some software is used solely for research and the advancement of science. There is a tradeoff between commercializing software when the research is federally funded and taxpayers paid for it, as commercialization of software can lead to job creation, which is also often part of the funding agency's mission. Programs such as NSF I-Corps incentivize tech transfer. Fifth, visualization is a multidisciplinary field, and it is often the case that scientists and engineers are unaware of general visualization tools that are available for their work. Visualization tools require a lot of time and effort to learn how to use and adapt to a particular application. Sixth, there is a tradeoff between making software open source. A PI may lose his/her competitive edge; however, there is the benefit of having a large user base who can contribute to further development of the code and citations of it.

## Omar Ghattas

Discussion questions and breakout group comments

**Summary of Breakout Session for Omar Ghattas's Talk:**

**Question 1: How can big data and big models be integrated to produce better predictive models?**
**Scribe: Pradeep Ravikumar**

*No scribe notes submitted.*

**Question 2: What are promising new ideas for exploring high dimensional space?**
**Scribe: Suzanne Shontz**

There are two major issues which must be addressed in order to efficiently explore high dimensional space. The first issue is that the high-dimensional models must be small enough to be to simulate them efficiently. The challenges lie in determining an appropriate parameter space in which to reduce the multiscale models; the physical space cannot be used for this purpose, as all scales must be preserved and interactions between various parts of the model must be maintained. Developing parameter space model order reduction (MOR) techniques for multiscale models is further complicated by the fact that various application domains, such as ice sheets, earthquakes, electronic circuits, airplanes, and buildings are very different. Thus, application-specific MOR techniques need to be developed. In addition, global (as opposed to local) MOR techniques are needed for use in solving many inverse problems. The second issue for exploring high-dimensional parameter space which we touched upon briefly is the need to move about the space efficiently. Structure-exploiting, scalable algorithms that map well onto extreme-scale systems are needed. For example, conventional Markov Chain Monte Carlo is too expensive and does not scale well when used to sample high-dimensional parameter space.

**Question 3: What are promising new ideas for quantifying uncertainties in modeling and simulation?**
**Scribe: Alison Marsden**

Uncertainty quantification (UQ) for models and simulations is not currently being practiced widely enough. This is often due to the large computational cost required to perform uncertainty quantification and a lack of checkpoints (or more advanced constructs) in simulations; reduced order models are needed to address the simulation time. It is also due to a fear of demonstrating inadequacies in models and simulations. For some modelers and computational scientists, it is due to a lack of knowledge as to how to properly sample the space or quantify the uncertainty. There is also a lack of understanding the separate but complementary roles of verification, validation, and UQ. Historically, there was a large safety factor which sufficed. On the flip slide, there are those who are using validation to over claim that "one test means my model or simulation is valid", and this leads to a false sense of security in the model or simulation. Community standards are needed in this area. There should also be a shift in focus from developing new methods to solve partial differential equations (PDEs) to developing UQ techniques and theory for PDEs.

**Question 4: How can we adapt/reinvent the important algorithms of CS&E so they better map onto high-throughput accelerators? Onto systems with massive numbers of cores?**
**Scribe: Ioan Raicu**

Reinventing important CS&E algorithms so that they better map onto high-throughput accelerators is a challenging task. One major question is whether the hardware should adapt to the software or the other way around? Or should they be co-designed? The Department of Energy, for example, has ongoing initiatives in co-design. In addition, it seems as though hardware can be expected to converge to many-core heterogeneous architectures. Thus, it should be the software that adapts to the hardware. One challenge with this is that FPGAs are difficult to program, and only the algorithm developer understands the algorithm in a deep way. Significant communication is required in order to determine how best to program the algorithm on FPGAs. Programming languages and programming models are the key to future adaptability to radically different computing architectures. Collaborations should be promoted across disciplines, centered on grand challenges in CS, Math, and Engineering, and with solid funding. One trend in the community is to start new Data Science degrees; this is a great first step in regards to realization of this collaboration and towards producing the next generation interdisciplinary workforce.

**Question 5: How can we transform our universities and federal agencies to become more hospitable to cross-cutting research/education at the interfaces of science/engineering, mathematics, statistics, and computing?**

**Scribe: Xiong (Bill) Yu**

Federal agencies, universities, and researchers can all make changes to become more hospitable to cross-cutting research and education at the interface of the above-mentioned areas. First, federal agencies can ensure the continuity of funding in these areas, can support junior faculty to growth the research community in these area, and can institute additional review criteria (i.e., require clarification as to how the work can be utilized by the community). Within universities, administrators can be educated to be advocates for interdisciplinary research, evaluation criteria for promotion and tenure can be modified to be made more appropriate for work in these areas, and resource allocation and management can be handled through an interdisciplinary research center/institute as opposed to a department. Researchers can realize that change happens most effectively from the bottom-up and can engage in interdisciplinary dialogue and raise the standard of research. They can also work to close the gap between computing and science/engineering.