TeraGrid Planning Process Report:
August 2007 User Workshops

Katherine A. Lawrence
<kathla@umich.edu>
Ann Zimmerman
<asz@umich.edu>
Collaboratory for Research on Electronic Work
School of Information
University of Michigan
Ann Arbor, MI 48109-2112
http://www.crew.umich.edu

December 3, 2007

Table of Contents

Executive Summary

TeraGrid is a national, distributed infrastructure integrating multiple resources at nine resource provider facilities. In late spring 2007, the NSF awarded a grant to the University of Michigan's School of Information (UM-SI) to facilitate a participatory planning process to help guide the future evolution of TeraGrid. This process is intended to anticipate changes that are already occurring and will take place in the areas of high-performance computing (HPC) and computational science over the next five to seven years. To gain an understanding of the requirements of current and potential future users of TeraGrid, UM-SI conducted a series of three workshops. Two of these workshops focused on the needs of TeraGrid and HPC users in ten research domains (astronomy, computer science, engineering, mathematics, and physics for Workshop A; biology and other biosciences, chemistry, earth and environmental science, materials science, the social sciences, and humanities for Workshop B). Thirty-five people participated in the two events, which were held on August 21–22 and August 23–24, 2007, near Chicago, Illinois. This report summarizes the results of the workshops.

The guiding question for the workshops was "Where could HPC take you and others in your field over the next 5 to 7 years?" Answering this question requires a better understanding of the specific role that high-end computing resources and services (such as those provided by TeraGrid) play in meeting researchers' needs. The workshops were organized to encourage small group discussion of these topics within and across disciplines. The first day of each focused on generating issues pertinent to individual disciplinary areas, and the second day was directed toward user needs based on the type or level of use.

Workshop participants identified issues related to HPC use, including the development and recruitment of people capable of using HPC, architecture and software, and the logistics of gaining access and administering their accounts. Some individuals feel that TeraGrid has yet to fully exploit its potential as a distributed computing resource that can aggregate data from many sources as well as parcel out computational jobs to multiple supercomputers. Participants in Workshop A identified seven key parameters to consider when planning the future of TeraGrid: architectural diversity/variety, resource scheduling, compute nodes, help/human resources/training, data access, data management/transfer, and data analysis and post-processing. There is consensus on the treatment of some of these parameters, but in other cases needs vary. Topics of particular interest to individuals in Workshop B were funding a future TeraGrid, meeting diverse user needs, improving the allocations process, enhancing TeraGrid's interaction environment, and including users in ongoing planning.

In sum, participants in both workshops recognized the reality and potential of TeraGrid to advance their research. Although they are a diverse group, they shared some common opinions about ways to better meet their needs. For one, they would like more consideration given to software in the future. Data management and schedulers appear to be weak links as HPC moves toward petascale. Barriers to HPC use include the allocation process, the steep learning curve, and a lack of courses and programs in computational science. Broader promotion of the research benefits of HPC and additional funding could help alleviate concerns related to education and training.

Introduction

TeraGrid is a national, distributed infrastructure integrating multiple resources at nine resource provider facilities.[1] The project is funded by the National Science Foundation (NSF) and provides access to computational resources, primarily in the form of supercomputers, large amounts of storage space, visualization services, fast networks, and software. Following a 5-year construction phase, TeraGrid became operational in late 2004. TeraGrid's resources currently include more than 250 teraflops of computing capability and more than 30 petabytes of online and archival data storage. High-performance networks provide rapid access and retrieval to data. TeraGrid supports a variety of use cases ranging from exploiting a single TeraGrid resource to combining resources across sites. The latter capability opens up new possibilities for conducting scientific work. Since TeraGrid became operational, many changes have occurred. For example, TeraGrid resources have been opened up to new user communities through the TeraGrid Science Gateways program, the number of resources providers has grown, and a user portal has been developed to make it easier to obtain information about accounts and resources and to simplify access to TeraGrid services. Further changes are certain to come. For instance, by the time the current grants expire in 2010, a petascale resource will be on the horizon, the user community will evolve, and new policies and services are likely to be needed to meet the needs of users and the expanding pool of high-performance computing (HPC) resources.

In late spring 2007, the NSF awarded a grant to the University of Michigan's School of Information (UM-SI) to facilitate a planning process. The process itself is led by a steering committee consisting of individuals representative of key stakeholder communities and various areas of expertise. The outcome of this process will be a report to NSF and to stakeholders that discusses options for the delivery of TeraGrid services and resources based on the diverse needs of science and engineering communities over the time frame 2010-2015. The report will be written by the steering committee utilizing input gathered from current and emerging users and from other stakeholders. The planning process is intended to anticipate the changes that are already occurring in the use of HPC services and resources. By relying on a community-driven, participatory process, the steering committee and facilitation team aim to develop well-informed options for evolving TeraGrid and wide acceptance of the planning outcomes.

One objective of the first phase of the process is to gain an understanding of the requirements of current and potential future users of TeraGrid. During summer 2007, as the steering committee was being established, researchers and staff at the UM-SI organized and conducted a series of three workshops to elicit preliminary information on user requirements that could be used to inform subsequent phases of the planning process. This report describes the results of the two workshops that focused on the needs of current and potential users of TeraGrid and HPC in ten research domains. The workshops were held on August 21–22 and August 23–24, 2007, near Chicago, Illinois. The first workshop, held in June 2007, focused on TeraGrid Science Gateways (Lawrence & Zimmerman, 2007).

The remainder of this report is organized as follows. We begin with a description of the workshop's purpose and participants. Next, we summarize the activities that were used to

[1] Since these workshops were held, two more resources providers became part of the TeraGrid, bringing the total number of sites to eleven.

address the guiding question and to acquaint attendees with each other and their projects. The main part of the report analyzes the results of the workshop discussions and activities. We conclude with a summary of the key findings from the workshop and the results of the meeting evaluation survey.

Workshop Purpose and Participation

The workshop was designed to assess how TeraGrid could meet the needs of current TeraGrid and HPC users in the future. The guiding question for the workshop was "Where could high-performance computing (HPC) take you and others in your field over the next 5 to 7 years?"[2] Specifically, we wanted to identify how TeraGrid could improve the capabilities available to those in research disciplines that currently use or that anticipate needs for HPC in the near future. The workshop activities were organized with the intention of identifying commonalities and differences in the requirements and priorities of the participants. Secondary goals of the workshop were to solicit ideas for engaging others in the planning process and to provide an opportunity for attendees to meet and interact with their colleagues.

Because the community of HPC users is so diverse, we organized the two workshops around affinity groups based on fields of research. Most of the invitees were current TeraGrid users drawn from the disciplinary areas that are the heaviest users of TeraGrid. Researchers from the social science and humanities, two areas that anticipate future growth in the use of HPC, were also invited. Workshop A invited individuals from astronomy, computer science, engineering, mathematics, and physics. Workshop B included researchers from the biosciences, chemistry, earth and environmental sciences, materials science, the social sciences, and humanities.

In addition to disciplinary diversity, we sought to include individuals with various levels of experience and knowledge regarding HPC and TeraGrid. Thus, we employed multiple methods to identify potential attendees. We asked participants in the June 2007 workshop for TeraGrid Science Gateways to recommend people, especially users of their gateways, who had a basic understanding of high-end resources and services and who could offer thoughtful and constructive input regarding the use of HPC to support their research. TeraGrid personnel and members of TeraGrid's Cyberinfrastructure User Advisory Committee (CUAC) also provided recommendations. We emphasized that potential participants did not need lots of HPC experience or knowledge and could be "emerging users" at any level—PIs, graduate students, post docs, and research assistants. We also consulted with experts who helped us to identify people from minority-serving institutions and underrepresented groups such as humanities researchers who use (or would like to use) HPC.

We contacted more than ninety individuals, inviting them to attend the workshop or send a qualified team member in their place. (See Appendix A for a copy of the invitation and Appendix B for the final list of participants.) In total, thirty-five people participated in the

[2] We purposely chose to use the word HPC in addition to TeraGrid because not all workshop attendees were TeraGrid users.

workshops.[3] Approximately forty additional people replied who could not attend but who sent a colleague in their place and/or expressed interest in future events.

Results of Pre-Workshop Survey

Prior to the workshop, we asked participants to answer questions in a brief online survey. The goal of this survey was to gather general information about the participants and to assist us in forming workshop breakout groups. The survey results showed that almost everyone in Workshop A had used HPC for more than 5 years, but about one-third of participants in Workshop B had either no experience or less than one or two years of experience with HPC. As might be expected, almost everyone in Workshop A considers HPC critical for answering research questions that interest them, whereas in Workshop B only 70% of participants consider HPC to be necessary to their work. Although all the participants had heard of TeraGrid before and had at least some degree of familiarity with its purpose, about one-third had never used TeraGrid's services (including both experienced HPC users and those new to HPC).

Workshop Structure and Activities

The workshop was organized to elicit user requirements based on fields of research, level of experience with HPC, and the nature of HPC or TeraGrid use. Appendix C contains a copy of the agenda for each workshop, which began with lunch on the first day and ended at noon on the second day. The first day was focused on generating issues pertinent to individual disciplinary areas, and the second day was directed toward user needs based on the type or level of use. In response to the timing of day one activities and participant comments, we adjusted the agenda, so the large group discussions and the plans for the second day were different for each workshop.

The workshop began with an overview of the TeraGrid planning process, an introduction to the planning process website (www.teragridfuture.org), and a presentation about TeraGrid by Dane Skow, lead of TeraGrid's Grid Infrastructure Group. Slides from all the presentations are available on the planning process web site.

The first group activity was a short physical exercise that was designed to help participants get to know each other, to prepare them for their small group discussions, and to illustrate the similarities and differences between their disciplines. To begin, we asked attendees to gather in the open space at the center of the room and to arrange themselves according to their geographic location in the U.S. We also asked them to introduce themselves and their field of research. Following this, they organized themselves in order based on the prevalence of people in their field using HPC. Finally, the participants stood in a two-dimensional space that represented the number of processors they use for a typical job versus their perceived level of experience with using HPC. An interesting difference existed between the two workshop groups. In Workshop B, there was a strong correlation between the number of processors used for a job and a participant's level of experience with HPC. In contrast, the participants of Workshop A use a wide range of number of processors even though all were very experienced with HPC. This

[3] Fourteen people participated in Workshop A and twenty-one attended Workshop B. An NSF program manager was also present at Workshop A. Due to weather problems that caused flight cancellations, the program manager scheduled to attend Workshop B was unable to travel to Chicago.

difference spoke to the varying usage models used to address particular types of research questions, both within and across fields.

After the warm-up exercise, we asked participants to meet in groups based on their field of research. We made these assignments prior to each workshop based on responses to the online survey. For Workshop A, these four groups consisted of astronomy/astrophysics, computer science, fluid dynamics/mechanics, and physics (such as particle physics). The five groups in Workshop B were computational chemistry, physical chemistry and biology, earth and environmental sciences, materials science, and social science and humanities.

Exercise/Day 1: SWOT and Storyboarding

On the first day, we employed two techniques to help achieve the workshop goals. The initial activity utilized a Strengths, Weaknesses, Opportunities, and Threats (SWOT) analysis, which is a strategic planning tool commonly used in business. In a SWOT analysis, an individual or team examines the helpful and harmful attributes (strengths and weaknesses) and helpful and harmful external conditions (opportunities and threats) related to the topic at hand. We began by asking participants to conduct an individual SWOT analysis of HPC in their field and to then spend time sharing their thoughts with the others in their group.

Storyboarding, the second technique we employed is a method for generating many different issues while grouping them into common categories. A storyboard begins by identifying the purposes of the topic under consideration, and then branches into generating key categories and specific issues within each category. The objective is to generate as many ideas as possible. For this part of the exercise, we asked the participants to examine *the future of HPC in their field over the next 5–7 years,* considering any technical, logistical, organizational, education-related, social, and other issues that may be relevant, but using their SWOT discussion as a starting point for generating ideas. They used large sticky notes to capture their ideas and post them on large pieces of paper on the wall.

Each group's final task was to prepare a presentation to the other groups that addressed the following questions:

- How is HPC used in your field?
- What are the key headings in your group's storyboard?
- What do people in your field *have in common*?
- What do people in your field *not* have in common?

During the small group presentations, we asked the participants to listen for and take notes on:

- What does that field have in common with your field?
- What is different/unique about that field?
- What don't you understand?

As each group reported back, the participants had the opportunity to ask questions. The presentations were followed by a large group discussion. The results of the small group presentations and the large group discussion are summarized in the Findings section (p. 9).

Exercise/Day 2: Supporting User Needs

With the goal of bringing to life some of the challenges and issues that users face as they work with HPC systems, we planned to conduct a series of activities on the second day that would help participants identify nuances of the process that they might not otherwise notice. As preparation for this activity, we asked the participants to identify the steps for using a distributed or grid HPC system. However, the first portion of the second day involved a large group discussion of issues that had been raised by the previous day's activities and any thoughts or concerns that had emerged overnight. In response to this discussion, the remainder of the workshop was different for each session.

The participants of Workshop A were particularly eager to explicitly capture technical needs for HPC going forward. To this end, we asked the participants to work in small groups to identify key axes that could define the parameters which should be considered in designing the next generation of TeraGrid. Each small group took turns describing important axes until the entire group felt that all axes had been captured. These axes, and any appropriate subtopics that were part of them, became the titles for flip charts that were posted around the room. The workshop participants then circulated to all of the flip charts and marked their needs on each axis with an "x." (These marks were essentially a "vote" by each person.) If an individual felt comments were necessary, he added them on the charts. Figure 1 shows an example of one of the flip charts under the topic *Architectural Diversity/Variety*. (All the flipcharts are compiled in Appendix D.)

# of processors sharing memory

```
LARGE
| xxxxx
| x
| x
|
| xx
|
| x
|
|
| xx
SMALL
```

Figure 1: Example of technical axes and votes from Workshop A

Participants in Workshop B followed their initial morning discussion with an activity based on the steps of high performance computing. They broke into pre-assigned groups based on their level of experience or type of HPC computing. These groups included:

- Non-TeraGrid users (some new to HPC, others experienced with HPC but new to TeraGrid)
- People with fewer than 2 years of experience
- People with more than 3 years of experience, divided into three groups based on the number of processors used for a typical job
  - small (<64)
  - medium (64-128)
  - large (256 or more)

The groups drew analogies between each of the steps of HPC and various products and services unrelated to HPC, identifying systems that could be used as models for improving the delivery of TeraGrid services.

The remainder of this document reports the key findings from each workshop based on the products from Day 1 and Day 2 activities and our own notes and observations.

Findings

Issues Generated by Field-Based Groups on Day 1

Working in small, discipline-focused groups, the participants explored the topic "HPC in [our field] in 5-7 years," drawing on the strengths, weaknesses, opportunities, and threats they had identified individually. We observed that even within groups, people from nominally similar fields often had very different computational needs for their research. Even so, the groups did their best to identify key concerns and what they did and did not have in common with each other. In this section, we identify themes shared across disciplines as well as discipline-specific considerations for HPC.

Common Themes

The participants explored a range of subtopics in their small group discussions such as purposes for HPC in their field, barriers to entry and other challenges, diverse technical issues, success stories, educational needs, hardware requirements, future applications, possibilities, and concerns, and ways to attract others to computational research. The common themes that emerged from these discussions are described below.

Use

- For most scientists (other than computer scientists), HPC is a tool and not an end in itself.
- Several fields have two primary research approaches: modeling and experiments. Although these two communities historically have not overlapped, they are beginning to work together on problems that combine their expertise. As a result, research discoveries are increasingly the result of collaborations between modelers and experimentalists.
- The volume of data in many fields is tremendous, and it is growing at a rate that exceeds researchers' abilities to analyze it.

People: Development and Recruitment

- Researchers are concerned about insufficient funding for expert personnel, both graduate students and professionals capable of using HPC or supporting its use. Hardware infrastructure requires significant technical support beyond installation and basic maintenance, both at the computing centers and from local staff members. Researchers suggested that these issues will only magnify with the next generation of machines. For example, users will need help from experts in order to develop, port, and debug codes if they are to utilize a petascale resource.
- Students often do not begin their programs with the computational knowledge or experience necessary for working on state of the art HPC, so they have to gain knowledge in this area in addition to learning about their domain. Thus, they face a significant learning curve. Further, extended waiting times for gaining access to HPC resources (e.g., for debugging) is very

discouraging for newcomers. One solution would be to provide new users with access to smaller, possibly local systems that have limited wait times.

Technology

- Some participants advocated for flexible or heterogeneous architectures for different kinds of use. For example, some research areas cannot take advantage of parallel processing because their algorithms are not inherently parallelizable.
- Participants are under the impression that the vendors do not try to find out what the users need, although NSF is aware that vendors would like to do so. The implication is someone (not necessarily TeraGrid) needs to identify incentives or deals that would make this communication exchange happen more readily.
- Generally, participants are concerned to make efficient use of the resources. They would like more memory, data, and I/O bandwidth (memory-to-disk, disk-to-archive, and site-to-site) for computational purposes as well as for moving data around quickly because they have difficulty with the logistics of transferring many thousands of files from the data sources. Data archiving is also a concern. Automation through more sophisticated schedulers would be a desirable way of addressing the challenges of data transfer and archiving. However, there was less consensus regarding the quality of the processors (i.e., expensive, high-speed); some want petaflops while others are content with "low quality."
- Many fields could use different levels of machines (e.g., local, lower-end supercomputers in addition to TeraGrid resources). This requires institutional buy-in and recognition that computers are a resource for research as well as for learning and enterprise management. This also introduces issues of compatibility across local and TeraGrid software environments.
- A unified environment that allows standardization both across the grid and with local environments would make it easier to move data and codes back and forth.
- Many disciplines will need algorithms to be redesigned/rewritten for use on petascale machines (or simply to use more processors than they are using now) in order to use them efficiently and effectively. These codes are often not open source, and they are not something that is appropriate for a graduate student to write a thesis about (whether in a science domain or in computer science), so there are few opportunities for getting the code rewritten within the university context. NSF-funded software development and the promotion of scientific computing as an interdisciplinary field of study would be very valuable.

Logistics and Administration

- People who use TeraGrid would like a longer allocation cycle (e.g., 3 years instead of 1 year) to better match the pace of their research.
- In general, people would like a faster turnaround on jobs.
- As the computational meshes that compose grids grow and larger numbers of processors are employed, the probability of failure for one node (which will kill the job) is higher. In addition, scheduled maintenance as well as unscheduled machine down times cause jobs to be checkpointed. Consequently, several restart runs may be needed in order to make a complete computation, but the computation time for each submission is limited and may be insufficient for finishing the calculations. Even if the scheduling policy allows very long jobs, checkpointing will be necessary to safeguard against failure. As scientific problems

prompt computing times to grow, even the petascale systems running large problems on many processors will face a high likelihood of hardware failure during every run.
- Researchers find it difficult to run visualizations because they cannot easily transfer the data back to their home base—perhaps just the summary statistics. Thus, they would like to be able to view visualizations directly from the remote supercomputers on the TeraGrid.

Discipline-Specific Issues

- **Computer science** focuses more on development and how machines are used, so a sandbox for trying out more exotic architectures is desirable to those working in this area. Researchers lament the lack of knowledge among administrators, which increases the difficulty of explaining the value of investing in these systems to them.
- Although they are ostensibly in the same field, **physics** researchers have diverse needs that are influenced by the different scales with which they work. Their HPC needs vary in terms of numbers of processors, length of job, number of collaborators sharing an instrument, and the data that are produced.
- **Fluid mechanics** integrates the science of physics with applications in engineering, geophysics, the medical sciences, and other areas. Thus, people in this field pursue research questions that are quite diverse, and they apply a wide variety of algorithms to their computations. Some fluid dynamics problems are highly non-linear, which can make it difficult to achieve perfect parallel scalability. With massively parallel architecture, knowledge of the underlying hardware is often necessary for writing efficient codes, and the optimal algorithm for an application is hardware-dependent.
- Like physics, researchers in **astronomy and space physics** do different degrees of physics. Their big challenges are schedulers as well as hardware and software for archiving because their jobs are so large and consumptive of processors. They are concerned that if they can ramp up to using more processors than they do now, the fault tolerance may not exist to support the demand.
- **Computational materials science** does not have large datasets for the most part, but they do use a lot of temporary disk space for the calculations. They have an "unbelievable profusion of minor variance in codes" due to the variety of problems they study. They do visualization locally on their own machines, so data transfer capabilities are important.
- **Computational chemistry** requires a lot of complex mathematics, but researchers have differences in their needs for large memory and long compute times (months). The needs for visualization also vary, depending on the size of molecules and types of systems they study. Data transfer and long term storage is an issue.
- **Earth sciences** researchers study diverse phenomena using very large data sets for numerical modeling, data mining, visualization, and many other uses. They spend a lot of time improving the models, even though they would rather put their attention toward the science. Some of what they do requires strong communication with the disk, so it may not be readily scaled. Strong I/O and communication is necessary for the code to run well on large machines with many processors.
- The group representing **computational biology, chemistry, and bioinformatics (chemistry with a biology slant or vice versa)** included people who maintain infrastructure at universities, and they suggested that infrastructure experts look at TeraGrid with a different perspective than researchers because they are more mindful of the costs and logistics associated with setting up and maintaining HPC systems. Bioinformatics researchers are not

interested in dealing with the complex details of HPC, so finding people who are capable of writing and adapting code that can handle many tiny text files is going to be an issue.

- **Social sciences and humanities** scholars have little in common in terms of the research they do, but they share the use of multi-modal data (e.g., text, audio, spatial, visual, census) and the lack of models with which to begin their analysis. They use HPC to generate hypotheses with data that have never before been gathered together in one place. These data are often non-standard and are not in digital form. Once data are digitized, moving them around can be problematic because of copyright issues. Some researchers are challenged by the interface to HPC, which is aimed at more mathematically inclined people (and people used to command-line interfaces) and is generally unfamiliar to those with humanities background.

Discussion Activities on Day 2: Workshop A

Large-Group Debrief of Key Issues

When the workshop reconvened on the second day, we began with a large group discussion of issues and concerns that had surfaced upon reflection or over dinner. The participants expressed confusion as to whether the planning process was directed toward the future of HPC generally or TeraGrid, specifically, as it relates to their requirements. In particular, they noted that what is unique and exciting about the TeraGrid is distributed computing—taking isolated resources and making them available in a globally distributed way. In contrast, they noted that our treatment of HPC tends to reflect the use of a specific resource. Additionally, they stated that the planning process needs to consider using distributed systems as an additional set of processes.

The group spent some time discussing their expectations about what they would like to communicate to the Steering Committee that is leading the planning process. Along the way they raised some interesting issues about the opportunities and challenges for using TeraGrid.

One potential that has yet to be fully realized is the aggregation of data from multiple sources. In particular, the physics and astrophysics communities have enormous data sets that would be very powerful if aggregated. However, to accommodate this volume of data the bandwidth must increase substantially. In addition, scientists need to identify useful sources of data that would be linked into the TeraGrid.

Significant environmental roadblocks impede the broadening of computational science (as a means to using HPC). At the level of researchers adopting new technologies, faculty who are located at teaching-oriented colleges and universities are challenged to find time to conduct research. The same is true at many minority serving institutions. This trickles down to students, who are not very often taught about HPC. Further, in computer science departments there is a general need for curriculum development in this area. Thus, across current faculty and students in the pipeline, there is scant awareness of what people could accomplish with HPC. Participants suggested that one avenue is to have “PR” in the form of scientists showing what they’ve accomplished using HPC. Nevertheless, they recognized that it is not necessarily up to TeraGrid to spread the word; institutional and professional involvement is also essential.

Overall, participants in the Workshop A felt that identifying technical areas of computational improvement was more important to them than figuring out how to get more people to use the system. Specifically, they wanted the opportunity to identify and convey their technical requirements to the planning process Steering Committee. One participant noted that TeraGrid is like a library—the users don't interact with all the people in the library, but they all have to share the resource. The goal is not to have giant, homogenous resources because TeraGrid is not one thing to all people, but identifying the boundaries of what is needed is still important. The next section describes the important technical topics identified by the participants.

Important Parameters for Technology

The individuals in Workshop A identified seven key technical issues. Within each of these topics, the group identified and voted on multiple parameter axes. (Histogram-style details of the voting are in Appendix D.) The nature of these issues and the voting results are described below.

- **Architectural Diversity/Variety:** Participants tried to parameterize the computational needs of the disciplines represented in the workshop, but it became clear that no single architecture would be considered ideal by all present. For example, some fields (such as fluid dynamics) solve problems that cannot be parallelized or scaled for computational efficiency and are not well served by symmetrical multiprocessing (SMP) or highly parallel (shared or distributed memory) computer architectures used by many disciplines. Also, researchers in computer science favor the opportunity to develop and study multiple types of systems.

  Appendix D displays the range of preferences for the following characteristics of machine architecture: number of processors sharing memory; fast communication software; low-latency, globally-shared memory; exotic architectures, memory bandwidth; interconnect speed; and number of processors. The majority of workshop participants were interested in architectures with large numbers of processors, shared memory, and efficient processor interconnection, as well as mainstream architectures for central processing units (CPUs), graphics processing units (GPUs), Field Programmable Gate Arrays (FPGAs), etc. Fast communication software was also very important to most participants. However, there was little consensus as to the importance of low-latency, globally-shared memory. Discussion prior to voting raised the idea of forming an expert advisory group to recommend particular types of compilers, schedulers, processors, and memory that would be best for certain types of research problems.

- **Resource Scheduling:** Concerns in this area focused on the ability to have "game-proof," realistic job scheduling that would be easier to use across multiple resources. Voting indicated that participants would like greater consistency of resource scheduling, although complete consistency is not necessary. Also, the group was fairly evenly divided among those who would like metascheduling capabilities across resources and those who did not consider it to be important. Not surprisingly, metascheduling is particularly important for those who use more than one resource for storage and computing.

- **Compute Nodes:** The key parameters for compute nodes were memory bandwidth, interconnect speed, and size. Almost everyone considered high memory bandwidth and a fast interconnect speed to be important. However, they indicated little consensus about the size of nodes; those who considered it most important would like to see something like 10^6 cores in 5 to 7 years, but many participants did not have strong opinions either way as to the size of the compute nodes.

- **Help/Human Resources/Training:** This topic includes the extent to which help centers and other training, curriculum, and HPC education activities should be centralized versus regionally distributed. A second facet of this issue is the nature of the help resources provided. For example, should the bulk of effort be focused on low-level users or on high-end users? The voting (and the comments written in the margins of the flipcharts) showed that participants feel that some aspects of help—face-to-face consultations, resource-specific questions, routine or straightforward use—are best addressed with regionally distributed help centers whereas training and high-end users are effectively supported by centralized resources.
- **Data Access:** This issue focuses on the consistency of the interface, particularly how transparent or standardized access is across resources. Most participants preferred a transparent, uniform interface to data storage, but a small number would rather allow unique interfaces if they offered better performance.
- **Data Management/Transfer:** This issue focuses on the transfer of data for three primary relationships: disk to archives, site to site, and main memory to disk. In terms of valuing transfer speeds between the processor and memory versus the memory and disk, more people voted for the processor to memory speed, and such speed is particularly necessary for certain algorithms (such as those used by LIGO). In terms of the importance of disk-to-archive and site-to-site transfers, votes ranged across the spectrum.
- **Data Analysis & Visualization (Post-Processing):** Data analysis and post-processing require both special resources as well as people to help TeraGrid users with those resources. These resources also have implications in terms of hardware, software, and bandwidth. The primary questions are how these resources should be distributed. Most people were inclined toward regionalized help as well as regionalized resources (such as small, specialized facilities). However, there are some types of research which require absolutely no post-processing once the results have been calculated.

Discussion Activities on Day 2: Workshop B

Large-Group Debrief of Key Issues

Before the end of the first day, we asked participants to think about issues that had been raised in their presentations and discussions and to identify additional topics of concern. We began the second day with about ninety minutes of general discussion. Five broad areas were discussed: funding a future TeraGrid, meeting diverse user needs, applying for allocations, improving TeraGrid's interaction environment, and including users in ongoing planning.

Funding a future TeraGrid

Participants were concerned about how NSF is allocating its money given the needs expressed by the participants in the workshops. They questioned if funds are most effectively spent through the procurement of "glamorous" machines. They would prefer to see half the money spent on additional Track 2 machines, noting that researchers need heterogeneous resources, not just petascale "big iron." Their needs range from no coupling to widely coupled, and many are concerned about large data services and storage, particularly the role of new research centers and instruments that will generate a lot of data (e.g., new observatories). In addition, many would like to see NSF fund more software development, which is a major stumbling block to using

HPC more efficiently. These questions raise issues about whether government spending for HPC should be based on the needs of a limited number of researchers or should consider the requirements of researchers more broadly. If a balance between the two is most appropriate, how is that best achieved?

Participants identified additional sources of funding for HPC that they would like to pursue or encourage. Some suggested that in the future the National Institutes of Health (NIH) should play a role in funding HPC. Politicians and taxpayers understand the value of biomedical research, and they should be made aware that as bioinformatics research grows, the demand for computing resources may increase substantially and that funding is needed to support it. Meanwhile, the participants would like to consider how scientists can convey the importance and value of their research to taxpayers. Can there be better support from NSF to help scientists with public relations (or to teach them how to do this while they are still graduate students), just as they do with education and outreach?

Meeting diverse user needs

Participants noticed that some fields or subspecialties have a well-defined class of models and software and can be readily supported with known technologies. In some cases, their technical needs are unique to their fields; however, a concern is whether NSF should be supporting specialized equipment on the TeraGrid. Consequently, to serve the needs of a broader audience, participants thought that it could be valuable to consider the needs of less well-defined disciplines for moving their science forward.

TeraGrid should also consider how to support low-end users, especially new users. For example, initially, many users need guidance to help them select the resource that best meets their needs. Participants suggested that "starter grants," might be employed to encourage beginning users to utilize local resources (university-based systems, some of which are partially or completely funded by NSF) before graduating to TeraGrid. One person suggested a mentoring program that would match less-experienced users with other researchers who could provide guidance along the way. Some suggested maintaining a playpen with retired machines (although retired machines would incur expenses for maintenance, energy, and cooling) or locally managed machines (that would save costs because need for fast communication is not necessary). Others recommended providing emulators on inexpensive machines, with plenty of user support, perhaps on a local machine or partition of a TeraGrid machine, so that learning how to code and debug is faster and more easily scheduled.

Applying for allocations

Many participants find the allocation application process to be disheartening. For example, the reviews are often harsh—written by someone not in their field—and provide little guidance. Further, people would like examples of good proposals. Participants noted a special need for resources to help them estimate their allocation needs. In addition, like those in Workshop A, participants in Workshop B perceived the one-year grant cycle as being too short (and poorly timed in January as the holiday season ends and a new teaching semester begins). Meanwhile, the off-cycle allocation option is not well publicized, but it could benefit those who cannot fit their research into the regular cycle. Some suggested that TeraGrid consider block grant funding (like NIH or NCAR) to help researchers more flexibly manage their budgets over a longer period of time.

Improving TeraGrid's interaction environment

Further to the issues with using allocation time for remote debugging, participants find that the current systems are too slow or cumbersome to work efficiently. Some feel that the basic interface works well enough, but they noted that it has few features. On the other hand, the more developed interface is too slow. One person suggested creating a more distributed system for gaining access to debugging: TeraGrid could provide support for developing and debugging code on a local machine (e.g., kits that match local environment to TeraGrid environment). The problem is that people who use many processors cannot get sufficient access for debugging; some participants recommended that it would be valuable to direct funding toward identifying a more efficient and inexpensive way of debugging large jobs.

Including users in ongoing planning

Many felt that the planning process should be ongoing and that it would be profitable for users and for TeraGrid if users have continual input. They suggested a user advisory board as one way to accomplish this. The User Advisory Committee of the Pacific Northwest National Laboratory's Environmental Molecular Sciences Laboratory was cited as an example.[4] Further, participants perceived past user surveys as being focused on the needs of the top twenty users, and they wished to see them cover the user population more broadly. They also suggested that "alienated" users—those who stopped using the resources—be targeted in future surveys in order to gain a more comprehensive picture of user needs. In addition, they supported methods to provide anonymous feedback. Finally, they would appreciate a response after they provide input, so they know if or how their ideas were considered.

Analogies: Suggestions for Improving the Delivery of TeraGrid Services

As mentioned earlier in this report, participants in Workshop B were asked to identify the steps of using HPC, and they came up with the following.[5]

1. Request allocation based on necessary machine (need to know how to use the resource somewhat prior to request)
2. Learn about the machine environment: login, password, compiling, applications, storage, scratch space, how to submit jobs, monitoring use
3. Transfer and manage data (scratch space, etc.)
4. Run test and monitor it
5. Administer account (ongoing), such as checking hours used, adding users (e.g., students)
6. Estimate timing and choose queue based on nodes available
7. Run job
8. Get results/analyze data

The participants then met in small groups to generate analogies as a way to consider means to to make these steps easier to accomplish. Some groups generated similar analogies, underscoring

---

[4] See http://www.emsl.pnl.gov/homes/uac.shtml.
[5] Workshop A also generated a list of the steps of HPC, but because their activities on Day 2 were changed in response to their feedback, the discussion here only represents Workshop B.

the broad potential for reconsidering the delivery of services in new ways. In the following subsections, we describe the analogies generated by the small groups for each of the steps of HPC, identifying the different ways that these ideas are important to users with different levels of experience or with different computational needs.

1. Request allocation based on necessary machine (need to know how to use the resource somewhat prior to request)

For those unfamiliar with HPC, requesting an allocation—enough Service Units (SUs) to get familiar with the resource—should be as easy as getting a Hotmail account.[6] With greater levels of experience, the process could be more like the process of applying for grant money at NSF. For example, applicants could talk to someone like a program officer for guidance (particularly how to estimate SUs for a first-time process) and have their applications read by more like-minded reviewers. In general, the process should be simplified—more like NSF's FastLane (and *not* like grants.gov)—with some sort of test or series of questions (e.g., regarding the processors and commercial software applications required) that would steer applicants to the right computing resource and help them know where and how they should apply. If applicants are unable to determine which resource is appropriate for their needs, then perhaps they should be advised to consider using a Science Gateway to simplify their access.

2. Learn about the machine environment: login, password, compiling, applications, storage, scratch space, how to submit jobs, monitoring use

Users at all levels would appreciate greater detail and support for learning new machine environments. Newer users would benefit from "wizards" to help provide an initial configuration (like those provided with Microsoft products), Flash animations to explain how to do things, and step-by-step guidelines for using a resource (with more details available as needed). Users would like websites with specific information about using each system, including an explanation of acronyms and how to add software to a machine (identifying dependencies). A significant hurdle is figuring out what can be done on each machine, so users suggested having an information website organized by *capability*, not by machine. For example, researchers need to find the right machine and the right software together, and some prefer to search by software first with the machine selection being a secondary consideration.

3. Transfer and manage data (scratch space, etc.)

Users would appreciate simple guidelines on data storage for each resource—for example, storage options, time limitations, what kinds of data should be stored on what parts of the system, effective methods of transfer, how to work with data that are separate from the processors, etc.

This should be as straightforward as keeping two ordinary desktop computers synchronized, which would require some sort of distributed global file system that could be accessed from a desktop machine.

[6] TeraGrid provides start-up allocations up to 30,000 service units which are called Development Allocation Accounts (DACs). These allocations require only minimal information from the principal investigators.

4. Run test and monitor

Users invoked quality control (QC) and closely-tracked delivery systems such as FedEx, UPS, and Amazon.com as models for improving the process of running test jobs. Like an assembly line or a shipping service, the resource software should let people know how much of their allocation/processor capacity is really being used so they know if they are achieving execution at the expected pace (or with a predictable delivery time). Other monitoring features could verify that the code is fully functional and debugged and that users will get their desired output when they run the job. (See also item 7 for issues associated with debugging queues.)

5. Administer account (ongoing), such as checking hours used, adding users (e.g., students)

Multiple groups suggested that account administration should be displayed like online banking (and praised the improvements of the current portal, which is no longer like monitoring multiple, separate accounts, but rather one bank account). Users like to see how much they are spending and what their balance is so they can make better or more efficient use of their remaining SUs.

6. Estimate timing and choose queue based on nodes available

Ideally, application software (commercial packages like Gaussian 03, Amber 10, ANSYS, or open source projects like Siesta) would make it easier to estimate how long a job will take; this is not really a TeraGrid problem, but rather something that application developers could support. However, one group had other alternatives for allocating and using nodes. For example, TeraGrid could consider partitioning large Track 1 and Track 2 machines so that some partitions could be used by researchers for a continuous, dedicated period of time (such as 2 weeks), just as an equipment resource like the SuperCollider would be used. (Note: This opportunity to schedule a large portion of a machine for a significant amount of time is currently possible, but some participants were unaware of it.) Meanwhile, for the shared portion of the resource, the maximum running time of a job should be 48 hours. Because some people request more time than they need or put in placeholders for their anticipated use (queue stacking), the time that people have to wait before their job starts is longer than it should be. One group recommended a better monitoring system to see who is "gaming" the system and penalizing people accordingly. The less experienced participants thought that newbies should be eligible for some priority treatment to help them get access to using the system—much like the federally-funded EPSCoR (Experimental Program to Stimulate Competitive Research) grants or a handicap in golf (perhaps based on the inverse of their allocation?).

7. Run job

Relevant to all aspects of maintaining an account and running jobs on TeraGrid is the ease of navigating the environment. One group suggested having a customizable environment like a Google home page or a My Yahoo page. Another interface suggestion addressed the need to change paths and keep track of where you are when working on HPC systems. This is a challenge for new users in particular (but probably applicable to all users). One suggestion was to color-code long path names to make them easier to read and track.

Each supercomputing center has to perform a balancing act to achieve maximum, and fair, machine usage. Typically the balancing includes considerations of the length of jobs, wait time in the queue, and machine availability for small jobs or debug jobs. Some of these, like length and wait time, are dynamic and are also dependent on the load, but for short jobs and debug jobs a specific partition, such as a debug partition, is very useful. It was suggested that the debug partition should be active at least from 8 am to 8 pm every day for jobs up to 1 hour, but optimally it would be available 24/7 for those who may be traveling to far-away time zones or collaborating internationally. Some suggested that no person should use more than 1/20 of debug queue hours per day in order to prevent people from "chaining" one job to the next. Such scheduling issues will become particularly important as researchers develop new algorithms for petascale machines which cannot be tested on a local system.

8. Get results/analyze data

Some suggested that the TeraGrid-provided applications should offer an interface for analysis. They thought it could be like Netflix; you say, "This is what I want," and it comes in the mail a few days later. The applications could also provide an interface for choosing what formats you would like.

Discussion and Conclusions

The disciplinary areas represented in the two workshops are diverse both within and across them, which shows the range of needs and uses for HPC. In spite of this variety, we found that certain issues were prominent for most, if not all, the participants.

One refrain we heard repeatedly is that petascale computing is far from what will adequately address this *diversity of HPC use*. Architectures need to be varied to support multiple user groups, though standard, consistent interfaces are desirable. The greatest consensus was for the importance of fast communication software/interconnect speeds (particularly for processor to memory transfer), large numbers of processors sharing memory, and mainstream architecture for the bulk of the processors, but the caveat is that with every specification, there were others who required something completely different.

Another dominant concern is that *software is neglected* as a minor issue compared to the "glamorous" hardware. For most scientists to take advantage of the upcoming Track 1 and Track 2 computers (or even simply scaling up to use larger numbers of processors), funding for upgrading software algorithms and codes is essential. Likewise, researchers want software that is mindful of the end-user experience, with well-designed interfaces for account maintenance and debugging, plus documentation organized from the perspective of users, not the system builders.

As more digital observatories come online and computing power becomes more advanced, the *volume of data* being transferred, generated, and stored is a huge concern. Handling so much data has implications for architecture, connectivity/transfer, archiving, and fault tolerance. Some participants believe that the potential of data aggregation has yet to be fully realized, which will create further demands on the system.

Another consistently stated requirement was the desire for *more sophisticated and flexible scheduling* to take advantage of grid infrastructure (distributed computing), to partition large data analyses over time, and to reduce wait time for debugging (emulators could help with this). Schedulers that prevent people from "gaming" the system, allow greater consistency of scheduler interfaces, and let people personally monitor the length of job runs are also desirable.

The *allocation process* itself needs reconsideration in terms of length and timing of the cycle, size of allocation blocks, guidance for writing applications, and simplicity for new users. Also, users would like better real-time feedback about their use of their allocation while they are running test jobs (and in between jobs) so that they use their time most efficiently and effectively.

Whether or not more disciplines begin to use HPC for research, NSF and TeraGrid still need to address ways of *simplifying the learning curve* for newcomers and low-level entrants. For instance, graduate students and postdocs do much of the hands on work, and they do not always come to a project with the requisite HPC skills, knowledge, or expertise. Participants suggested that NSF fund local institution-based systems outfitted with emulation environments or TeraGrid-compatible software, offer support systems for preliminary staging of new code and calculations as well as local debugging, and maintain lower-end machines (perhaps retired ones). To make learning HPC easier, TeraGrid could develop unique funding/allocation processes, friendlier interfaces and more navigable environments, documentation and "wizards" from the perspective of new users, and offer some priority treatment for accessing the system.

Participants identified a *huge need for education* about scientific computation, both as a field itself and within non-computer science disciplines. Because computational science layers high-level coursework on top of already rigorous programs, educators also need to figure out how students learning HPC within these programs can finish their graduate degrees in reasonable time frames. Better training and education should help reduce the deficiency in the number of expert personnel for infrastructure support. Users are eager for access to human technical support for consultations, resource-specific questions, training, better documentation and guidelines for use, and regional visualization support.

Finally, there are ways that TeraGrid can *raise its profile* within its user community and beyond. Participants suggested that HPC needs better PR to convey its value and potential to state funding agencies, government officials, taxpayers, university administrators, future students, and computer science educators both to attract more resources and to develop new users. Among those already using HPC, particularly TeraGrid, people would like to be a continual part of the conversation. Participants suggested an ongoing planning process taking user input into consideration through advisory boards, broader surveys, anonymous feedback, and regular reports to constituents.

In conclusion, users recognize the potential of TeraGrid to help them conduct their research effectively and efficiently, but they face significant hurdles to making the process feel smooth and straightforward. They are a diverse group, but they have strong, coherent opinions about what would benefit their use. In particular, they would like more consideration given to the potential of software. Data management and schedulers appear to be weak links as HPC moves toward petascale. Barriers to use include the allocation process, the steep learning curve, and a lack of courses and programs in computational science. Broader promotion of the research benefits of HPC and additional funding could help alleviate concerns related to education and training. The diverse needs of stakeholders will pull limited resources in many directions, necessitating careful consideration of what services and improvements will best use available funds while adequately responding to the needs of future TeraGrid users.

Workshop Evaluation

At the end of the workshop, we left time for participants to complete a survey that asked them to evaluate the information received prior to the meeting, the clarity of the workshop goals, the quality of the presentations, instructions, and workshop activities, and their overall impressions of the event. Based on this feedback, participants in Workshop B were more satisfied than those in Workshop A. The evaluations of the workshop activities for the Workshop A averaged between 3 and 4 on a scale of 1 to 5, suggesting that the activities offered a better than average benefit and the workshop had a comparable level of organization. The evaluations for Workshop B averaged above 4 (good to very good), with particularly high marks for the organization and length of the workshop and the workshop as a whole. Participants in Workshop A generally preferred the large group discussions and the opportunity to discuss and vote on the technical axes. Participants in Workshop B tended to express a preference for one activity over another, but each exercise was singled out as helpful and effective by some participants. For future workshops, some useful, but often conflicting, suggestions were made:

- In Workshop A, some would have preferred a more directed process with less ambiguity, the inclusion of pre-digested material, and more focus on current needs, although others were comfortable with an emergent format.
- In Workshop A, some would have liked to focus more specifically on technical issues (and even more specifically on TeraGrid rather than HPC), but others were cautious about such a technical focus.
- In Workshop B, several participants suggested inviting more people to obtain even greater diversity but others advocated for more focused meetings oriented around particular research areas.
- A few individuals suggested that greater involvement from TeraGrid and/or supercomputing center personnel would be desirable.

Participants also look forward to seeing their suggestions implemented with the result that wait times will be reduced, the proposal process will be improved, and future computing developments will continue to add value to their research. They also expressed the desire for ongoing opportunities to provide feedback as the planning process moves forward and the hope that future NSF solicitations will incorporate suggestions that they generated in these workshops.

Acknowledgements

We are grateful to the participants for their active participation and interest in the workshop, traveling amidst summer storms and airline chaos to join us and for their comments on a draft of this report. We also thank Dane Skow, lead of TeraGrid's Grid Infrastructure Group, for his attendance and presentation at the start of each workshop. The assistance of Rebecca O'Brien was crucial to the success and smooth running of the workshops. The TeraGrid Planning Process is supported by the National Science Foundation under Grant No. OCI-0724300 to the University of Michigan. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

Lawrence, K. A. and Zimmerman, A. (2007). *TeraGrid Planning Process Report: June 2007 Workshop for Science Gateways*. Collaboratory for Research on Electronic Work, School of Information, University of Michigan, Ann Arbor, MI.
Available: http://teragridfuture.org/events/workshop/formatofworkshop/report

Appendix A: Workshop Invitation

Hello,

We would like to invite you or a member of your research team to attend an invitational workshop whose purpose is to solicit information about the needs of current and future users of high-performance computing (HPC). Specifically, the input gathered at the workshop will be used to help guide future plans for TeraGrid. If you have questions after reading the information below, please contact Katherine Lawrence (kathla@umich.edu; 734-994-7904).

When and where is the workshop?
The workshop will be held near Chicago's O'Hare Airport, beginning on [August starting date for each workshop] at 1 p.m. and running through to noon the following day. Further details regarding the specific location will be available soon.

When do I need to reply?
Please reply to Katherine Lawrence (kathla@umich.edu; 734-994-7904) **by Tuesday, July 31** as to whether you or a colleague will be able to attend. We would appreciate hearing from you even if you cannot attend.

Do I have to pay for anything?
We will provide support for travel-related expenses for you or one individual from your project to attend. We will send travel information once you reply to the invitation.

What is the context of the workshop?
This workshop is part of a series of activities associated with a planning process funded by the National Science Foundation (NSF). This process to help guide the future evolution of TeraGrid will be led by a Steering Committee representing key stakeholder communities. The committee Chair is Tim Killeen, and the Associate Chair is Roberta Balstad. The process is being facilitated by the University of Michigan's School of Information (UM-SI). The upcoming workshop will be conducted by Katherine Lawrence of the UM-SI. Please see the end of this message for further information about the planning process.

Why am I receiving this invitation?
Invitees were identified in several ways, including recommendations from colleagues familiar with your work and suggestions from people associated with TeraGrid. Basically, you are receiving this invitation because you are recognized as someone who would offer insights and expertise helpful to the planning process. We hope **you** can attend. If not, we would welcome a graduate student, post-doc, or other researcher from your team. Participants need not have lots of HPC experience or knowledge; however, it would be most useful if they have a basic understanding of high-end resources and services and how HPC might be used in support of their research.

I'm not a TeraGrid user. Can I still come to the workshop?
You do not need to be a current user of TeraGrid or of HPC to attend the workshop. In addition, no prior knowledge of TeraGrid is required.

What will be done with the workshop results?
The information collected at this and other workshops will be used to better understand user needs and priorities and to help develop options for delivering HPC services when the current TeraGrid grants expire in 2010.

Who else will be at the workshop?
The workshop will include researchers from fields including [astronomy, computer science, engineering, mathematics, and physics for Workshop A; biology (and other biosciences), chemistry, earth and

environmental science, materials science, the social sciences, and humanities for Workshop B]. Some participants will have lots of experience with HPC while others may just be getting started.

What else do I need to do?
To make this short meeting as productive and engaging as possible, we will ask attendees to answer a brief survey in advance of the meeting. This survey will help us prepare activities in advance and will allow participants to efficiently share information with colleagues at the workshop.

Thank you for considering this invitation. We hope to see you in Chicago!

Sincerely,

Katherine Lawrence and Ann Zimmerman, Workshop Co-Organizers
School of Information, University of Michigan

Timothy L. Killeen, Chair, Steering Committee
Director, National Center for Atmospheric Research

Roberta Balstad, Associate Chair, Steering Committee
Senior Fellow, Center for International Earth Science Information Network

****

ABOUT THE TERAGRID PLANNING PROCESS

The NSF is providing support for a community-driven, participatory planning process whose goal is to provide information that will help guide the future evolution of TeraGrid. Current awards for the operation, user support, and enhancement of the TeraGrid facility will expire in 2010. By this date, a petascale computing resource will be on the horizon, the user community will have grown and diversified, and new policies and services are likely to be needed to meet the needs of users and the expanding pool of high-performance computing resources. In anticipation of these changes, the planning process is focusing on the needs of current and emerging user communities as a critical aspect in the development of a path forward for TeraGrid in 2010 and beyond.

Planning activities will be conducted over the space of approximately one year and will include a combination of face-to-face and online engagement designed to:

- gather information on user needs and priorities;
- compare user requirements; and
- develop options for the delivery of high-performance resources and services

The results of the planning process will be a report to the stakeholders that outlines options for the design of the next generation TeraGrid and for the delivery of high-end resources and services based on user requirements. The report will be written by the Steering Committee using the information and input gathered from stakeholders throughout the planning activities. The final version of the report is targeted for February 28, 2008.

The Steering Committee is currently being formed and a web site is in development. For further information on the planning process please contact Ann Zimmerman, School of Information, University of Michigan, at asz@umich.edu.

Appendix B: List of Workshop Participants, Disciplines, and Affiliations

Workshop A

Richard Aló	Computer and Mathematical Sciences, University of Houston-Downtown
James G. Brasseur	Mechanical Engineering, Bioengineering, and Math, Pennsylvania State University
Gene Cooperman	Computer Science, Northeastern University
Rupert Croft	Physics, Carnegie Mellon University
Diego Donzis	Aerospace Engineering, Georgia Institute of Technology
Charles Goodrich	Astronomy, Boston University
Steven Gottlieb	Physics, Indiana University
Mark Hagen	Spallation Neutron Source, Oak Ridge National Laboratory
Anthony Harkin	School of Mathematical Sciences, Rochester Institute of Technology
David Hebert	Mechanical Engineering, University of Massachusetts-Amherst
Scott Koranda	LIGO Scientific Collaboration (LSC), University of Wisconsin-Milwaukee
Robert J. Moorhead II	Electrical and Computer Engineering, Mississippi State University
Tim Olson	Science and Engineering, Salish Kootenai College
David Porter	Astronomy, University of Minnesota
Steve Meacham	Office of Cyberinfrastructure, National Science Foundation

Workshop B

Jerry Bernholc	Physics, North Carolina State University
Andrew Beveridge	Sociology, Queens College; Graduate School and University Center of the City University of New York
Sen Chiao	Engineering, Marine and Environmental Systems, Florida Institute of Technology
Erik Deumens	Chemistry and Physics, University of Florida
Olga Dmytrenko	Chemistry, University of Delaware
J. Stephen Downie	Library and Information Science, University of Illinois, Urbana-Champaign

Wayne Dyksen	Computer Science and Engineering, Michigan State University; MATRIX Center for Humane Arts, Letters & Social Sciences Online
David J. Earl	Chemistry (Materials Science), University of Pittsburgh
Lev Gelb	Chemistry (Materials Science), Washington University in St. Louis
Gautam Ghosh	Materials Science and Engineering, Northwestern University
Hassan Karimi	Information Sciences, University of Pittsburgh
Keith Laidig	Biochemistry, University of Washington
Gillian Lynch	Chemistry, University of Houston
Matt Mazloff	Oceanography, Massachusetts Institute of Technology
Heather Netzloff	Chemistry, Iowa State University
Scott Perrin	Chemistry, Georgetown University
Jim Rattling Leaf	Sicangu Policy Institute, Sinte Gleska University
Patricia (Pat) Seed	History, University of California, Irvine
Angela Wilson	Chemistry, University of North Texas
Don Wuebbles	Atmospheric Sciences, University of Illinois, Urbana-Champaign
David A. Yuen	Geology and Geophysics, University of Minnesota

Appendix C: Workshop Agenda

Time: Tuesday, August 21 [and Thursday, August 23], 2007, from 1:00 pm to 6:00 pm
Wednesday, August 22 [and Friday, August 24], 2007, from 8:30 am to 12:00 noon

Location: Lindbergh A Room, Ground Floor, Hyatt Rosemont

Hosts: Katherine Lawrence, Ann Zimmerman, and Becky O'Brien
(kathla@umich.edu, asz@umich.edu, and beckyobr@umich.edu)

Workshop Guiding Question

Where could high performance computing (HPC) take you and others in your field over the next 5 to 7 years? Answering this question requires a better understanding of the specific role that high-end computing resources and services (such as those provided by TeraGrid) play in meeting your research needs. Thus, the workshop activities are organized with these topics in mind.

Workshop Goals

1. Assess the requirements of HPC users. Specifically, how can TeraGrid improve the capabilities available to you and those in your research discipline? To do this, we will identify the common and different needs and priorities among the participants, with attention to the different levels of experience represented by the group.
2. Solicit your ideas for engaging members of your disciplinary community in the use of HPC and the TeraGrid planning process.
3. Provide you with an opportunity to interact with each other around topics of shared interest and give us ideas about how to support your continued participation and interaction throughout the planning process...and beyond.

Workshop Outcomes

We will produce a report that will be posted on the planning web site after the workshop. We will ask you to comment on the report before it is posted. The information collected during the workshop will carry forward in the planning process. It will be used in creating an overall picture of user needs and requirements. It will also be used to gain stakeholder input on the development of options for the future delivery of high-performance resources and services.

Day 1

12:30 pm: Check-in and get lunch

1:05 pm: Introduction (and lunch continues)

- Welcome from workshop organizers, explanation of the TeraGrid planning process, and brief demo of the TeraGrid Planning Process website
- Dane Skow, Director of TeraGrid's Grid Infrastructure Group, talks about TeraGrid's plans for the near future.
- Review of the goals and agenda of the workshop

2:10 pm: Brief bio-break

2:20 pm: Warm-up exercise to familiarize participants with each other's HPC work

2:35 pm: Idea generation activity in small groups from perspective of your discipline
Participants will identify current and future strengths and weaknesses of HPC for their research, sharing their insights in small groups composed of others from their same disciplinary area. In these groups, participants will explore the future of HPC in their field, generating ideas of how it could be used by themselves and others over the next 5-7 years.

4:05 pm: 15 minute break

4:20 pm: Presentations by small groups and assembly of key themes
Presentations by small groups will answer the questions:

- How is HPC used in your field?
- What do you, as a group, have in common and *not* in common in terms of how you use HPC?
- What is unique about your uses?

The larger group will then identify key themes across the presentations.

5:10 pm: Discussion of outcomes and preparation for second day
Participants will discuss the key themes to consider for future HPC and TeraGrid planning. We will briefly prepare for the second day's activities.

6:00 pm: Adjourn for the evening

Day 2

8:00 am: Breakfast served

8:30 am: Debrief from first day

8:45 am: Exploration of the HPC process
Participants will identify and explore key aspects of using HPC, focusing on each of the stages necessary to use a system such as TeraGrid.

9:30 am: Small group discussion of HPC use
Divided into groups based on experience with using HPC, participants will discuss what they observed in the prior exercise and how it applies to their own experience with HPC. They will focus on the following questions to generate solutions for improving their use of HPC:

- How can HPC providers better meet your needs in the future? What would make it easier for you to use HPC?
- What do you have in common with other users of your level of experience? How would you advocate that HPC service providers support your group (especially based on different levels of experience)?

10:15 am: 15 minute break

10:30 am: Small group presentations of HPC use and prioritization of solutions
Participants will share their favorite solutions in brief presentations, followed by individual prioritization of solutions.

11:15 am: Reactions, feedback, and workshop evaluation survey
Participants will identify key issues and opportunities to consider going forward. We will discuss how, in future workshops, events, and initiatives, we can most effectively engage the end users and developers of TeraGrid and HPC.

12:00 noon: Workshop ends

Appendix D: Voting on Important Technical Axes (Workshop A)

Marginal comments are noted near the associated "x" marked in voting. Participants voted on whatever dimensions were important to them, so not all axes have the full number of votes.

(1) ***Architectural Diversity/Variety***

# of processors sharing memory

LARGE
| xxxxx
| x
| x
|
| xx
|
| x
|
|
| xx
SMALL

Fast communication software

VERY IMPORTANT
| xxx
| xxxx
| x
|
| x
|
|
|
|
| x
NOT IMPORTANT

Low-latency, globally-shared memory

VERY IMPORTANT
|
|
| x
|
| x
| x
|
| x
| x
| x
NOT IMPORTANT

Immature architecture (cell processors, CPUs, etc.)

EXOTIC ARCHITECTURE
| x
|
|
| x
|
|
| x
| xx
| xxxx
| xx
MAINSTREAM ARCHITECTURE

(2) ***Resource Scheduling***

CONSISTENT
| xx
| x
| x
| x
| x*x* (would like to be able to estimate job
| xxxx submission)
|
|
|
|
VARIABLE/VARIETY

Multiple resource scheduling

IMPORTANT
| x
| xxx
|
|
|
|
| xx
| xx
| x
|
NOT IMPORTANT

(3) ***Compute Nodes***

Memory bandwidth

IMPORTANT
| xxxx
| xxxxxxx
| xx
| x
|
|
|
|
|
|
NOT IMPORTANT

Fast interconnect speed

IMPORTANT
| xxxxx
| xxxxxx
|
|
|
|
|
|
| x
| x (LIGO: serial jobs)
NOT IMPORTANT

Size

IMPORTANT
| xx*x* (10^6 cores in 5-7 yrs)
| x
|
| x
| x
| x
| xxxx
| xxx
|
|
NOT IMPORTANT

(4) ***Help/Human Resources/Training***

DISTRIBUTED
| x*x* (*for some things*)
| x*x* (face-to-face consultations, help)
| x
| *x* (help line for each resource)
| xxxxx*x* (centralized for intense users,
| xx distributed for routine, simpler assistance)
|
| *x* (training)
| x
| *x* (*for other things*)
CENTRALIZED

(5) ***Data Access/Archive Access***

TRANSPARENT
| x*x* (desire standard interfaces to all storage)
| xxx
|
| x
| x
| x
|
| x*x* (performance)
|
|
UNIQUE TO HOSTS

(6) ***Data Management/Transfer***

Which needs to get "more" faster

PROCESSOR ←→ MEMORY
| xxx (LIGO: memory bandwidth more important
| xx*x* since fundamental to basic algorithm)
| *x* (both need maximum)
| x
| xx
| x
|
|
| x
| x
MEMORY→ DISK

Importance of disk-to-archival transfer

IMPORTANT
| x
| x
| xx
|
| xx
|
| xx
| xx
|
| x
NOT IMPORTANT

Importance of site-to-site transfers

IMPORTANT
| xx
| x
| xx
|
| xx
|
| xx
| xx
|
| x
NOT IMPORTANT

(7) ***Data Analysis & Visualization (Post-Processing)***

REGIONALIZED HELP
| xxx
| xxxx
| x
| x
| x (tier 1=local; higher=centralized)
|
| x (focus on support through
| a gateway mechanism)
|
|
CENTRALIZED HELP

REGIONAL RESOURCES
| xxx*xx* (but little; many small & specialized)
| xxxx
| xxxx
|
|
|
|
|
|
| *x* (visualization not an issue)
"ONE" RESOURCE