HPC University

Brad Armosky, Texas Advanced Computing Center, Shawn Brown, Pittsburgh Supercomputing Center, Tony Drummond, Lawrence Berkeley National Laboratory, Jim Ferguson, National Institute for Computational Sciences, Richard Gerber, Lawrence Berkeley National Laboratory, Thomas J. Hacker, Purdue University, Scott Lathrop, University of Chicago/Argonne National Laboratory, Mary Ann Leung, Krell Institute, Laura McGinnis, Pittsburgh Supercomputing Center, James A. Rome, Oak Ridge National Laboratory, Kathy Traxler, Louisiana Optical Network Initiative, Julia White, Oak Ridge National Laboratory, John Ziebarth, Krell Institute

1 INTRODUCTION

HPC University is a virtual organization launched in 2007 focused on high-quality, high-performance computing (HPC) learning and workforce development activities and resources. HPC University priorities are driven by community needs and requirements.

HPC University has a broad mandate to address the needs of a large and diverse community that includes K-20 educators and students; graduate, post-doc and senior researchers; administrators; system administrators; and practitioners in all fields of study related to HPC. The primary emphasis to date has been to address training requirements of the research community on topics ranging from introductory to petascale level performance of scientific codes. The virtual organization will extend the effort to include resources and activities to address graduate, undergraduate and high school education.

Participation in this virtual organization is open to all interested organizations that want to help expand the breadth and depth of the range of resources and services as well as to help broaden community involvement. Participation to date includes representatives from Argonne National Laboratory, Computational Science Education Reference Desk, Indiana University, Krell Institute, Lawrence Berkeley National Laboratory, National Center for Atmospheric Research, National Center for Supercomputing Applications, National Energy Research Scientific Computing Center, National Institute for Computational Sciences, Oak Ridge National Laboratory, Ohio Supercomputer Center, Open Science Grid, Pittsburgh Supercomputing Center, Purdue University, San Diego Supercomputer Center, Shodor Education Foundation, Inc., the Texas Advanced Computing Center, and the University of Chicago.

During the TeraGrid'08 Conference, the HPC University virtual organization will provide an update of the HPC University requirements analysis and implementation plans. The team will identify current and potential actions to address community needs. There will be a question and answer period to solicit additional community input and foster increased collaboration and participation among the HPC community. We invite all interested organizations to join in developing effective strategies for expanding and scaling-up the opportunities to best serve the education and training needs of the HPC community.

2 REQUIREMENTS

The HPC University team organized a Requirements Analysis Team (RAT) that identified several promising paths for creating qualified, effective HPC professionals capable of exploiting current terascale and upcoming petascale technologies for the advancement of scientific research. Gaps in training materials and delivery methods were identified, with recommendations for filling these gaps prepared. In addition to developing materials for addressing the evolving community needs for high performance computing, the RAT identified specific concerns with respect to the following issues:

- Persistence - Are up-to-date materials available?
- Quality assurance – Do the materials provide a validated, verifiable experience for the users?
- Delivery methods – Are the materials available to the users independent of geography or temporality?
- Scaling the training – Are good training practices identified and deployed for the development of the instructor pool?

2.1 Current catalog of materials, mapped onto technology continuum

The HPC University Requirements Analysis Team (HPCU RAT) conducted a survey of existing education and training materials, focusing on materials available through multiple TeraGrid Resource Providers, the Department of Energy (DOE) national laboratories, and other organizations with HPC resources we could identify through web searches and contacts with people in the HPC community. The HPCU RAT also conducted surveys of users and HPC allocations committees. Analysis of these materials resulted in the identification of the HPC training topics shown in Figure 1.

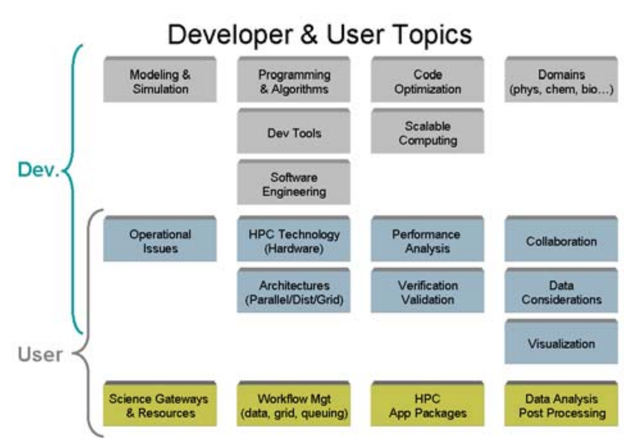

Figure 1: HPC Training Topics

In this figure, topics shown in gray are topics considered necessary for scientific application developers and topics shown in yellow for application users. Topics shown in blue are applicable for both HPC developers and application users.

Table 1 shows the training topics applicable for people running user application codes and writing their own HPC programs. Table 2 shows the training topics applicable for various audience types: Novice, Apprentice, Journeyman, or Master. These audience types roughly correspond to the academic levels: Undergraduate, Master, Ph.D., and Post-doc/faculty and reflect levels of expertise in the high performance computing skill set, not necessarily scientific competence or academic achievement.

Table 1: Training Topics Vs. User/Developer

Topic	Application User	Developer
Modeling & Simulation		✓
HPC Technology (hardware)	✓	✓
Architectures (Parallel, Dist, Grid)	✓	✓
Workflow management	✓	
Programming & Algorithms		✓
Development Tools		✓
Software Engineering		✓
Operational Issues	✓	✓
Performance Analysis	✓	
Verification and Validation	✓	✓
Code Optimization	✓	✓
Data Considerations		✓
Data analysis/post processing	✓	✓
Visualization	✓	
Scalable Computing	✓	✓
Domains (physics, chem, etc.)		✓
HPC Application packages		✓
Science gateways & resources	✓	
Collaboration	✓	

The training materials identified in the survey are categorized into the topics shown in Figure 1 and listed by topic below. The training materials range from short self-paced web tutorials to instructor-led courses, slide presentations of seminar materials, and video recordings.

The categories with the largest number of training resources available are:

- Operational Issues
- Programming & Algorithms

- Development Tools

While Operational Issues have the largest number of entries, many of the materials are site and hardware specific. Programming & Algorithms entries are predominantly focused on MPI, OpenMP, and parallel programming. The next group of topics with the most entries are:

- Architectures (parallel/distributed/grid)
- Science gateways & resources
- Performance analysis
- Visualization

Topics with no entries include:

- Workflow management
- Software engineering
- Verification and Validation
- Data Analysis & Post processing

2.2 Petascale Requirements

The primary focus of HPC University is high-quality, high-performance computing (HPC) learning and workforce development activities and resources, and that includes preparing people across a continuum that leads them to a level of knowledge and proficiency to deal with the largest computing systems available (petascale in 2008 and even larger scaling in the not-so-distant future). Providing and supporting training at this level of science is interesting and challenging, as the requirements upon developers and users are being hashed out in the modern arena of research. Achieving the scale of performance and reliability required for petascale computing is a significant challenge. Nevertheless, it is necessary to forge ahead as large NSF investments in the so-called "Track 2" architectures, such as Ranger at TACC, and the planned installation of Kraken at NICS, move the field rapidly toward the petascale.

The RAT sought to understand the gaps that exist in current parallel computer training, especially with an eye toward what will be needed to get to petascale computing. We consulted the Extreme Scalability RAT (XS-RAT), a team charged with determining what hurdles TeraGrid users would have to leap in order to reach petascale computing [1], in an effort to map our work to the experts view of the field. The XS-RAT targeted key areas of

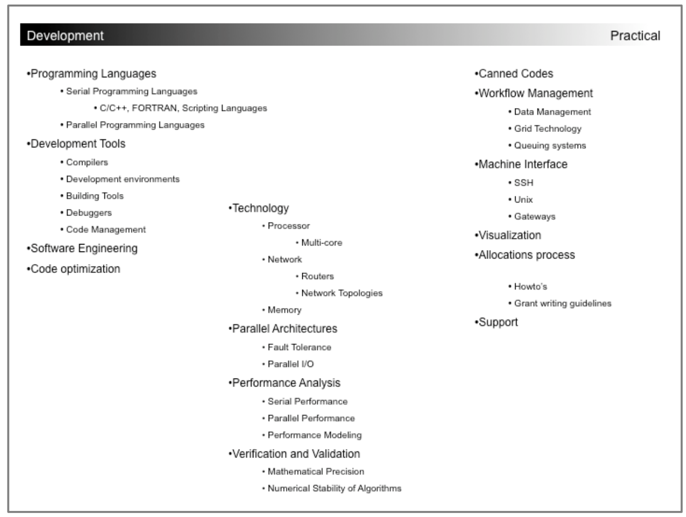

Figure 2: Petascale training map

development: multi-core and parallel architectures, debugging and profiling technologies, and visualization. Other areas included reliability and fault tolerance, modern programming models, parallel I/O, and workflows. One clear, permeating theme: the boundaries between user and developer would need to blur, with each adopting characteristics of the other, in order to create a meaningful

Table 2: Training Topics Vs. Audience

Topic	Novice	Apprentice	Journeyman	Master
	Undergrad	Master	Ph.D.	Post-doc/faculty
Modeling & Simulation	✓	✓	✓	✓
HPC Technology (hardware)	✓	✓	✓	✓
Architectures (Parallel, Dist, Grid)	✓	✓	✓	✓
Workflow management	✓	✓		
Programming & Algorithms		✓	✓	✓
Development Tools		✓	✓	✓
Software Engineering		✓	✓	✓
Operational Issues	✓	✓	✓	✓
Performance Analysis		✓	✓	✓
Verification and Validation	✓	✓	✓	✓
Code Optimization		✓	✓	✓
Data Considerations	✓	✓	✓	✓
Data analysis/post processing	✓			
Visualization	✓	✓	✓	✓
Scalable Computing		✓	✓	✓
Domains (physics, chem, etc.)		✓	✓	✓
HPC Application packages	✓			
Science gateways & resources	✓			
Collaboration	✓	✓	✓	✓

computational environment at the petascale. Figure 2 displays our mapping of various training topics to a continuum based on this merging. At the poles of the figure are topics central to one group or the other, with the key list lying at the midpoint of development and practical concerns. This map, while not complete, provides a starting point for analyzing what requirements are necessary for petascale training.

2.3 Quality Assurance (VVA)

High-end computing systems are becoming more readily accessible for scientists, with faster and more powerful systems coming on-line every year. For efficient and effective use of high-end computational resources, the community needs timely materials that allow them to learn about HPC in a rapidly changing environment. It is important that quality assurance for all resources be provided by HPC experts in conjunction with community perspectives.

The Shodor Education Foundation developed the Computational Science Education Reference Desk (CSERD) to provide access to computational science resources. CSERD (http://cserd.nsdl.org) is a Pathways effort of the National Science Digital Library (http://www.nsdl.org) to provide a digital repository of materials for teaching and learning across all domains. CSERD utilizes a comprehensive review mechanism, called Verification, Validation and Accreditation (VV&A).

The VV&A review process provides an on-line review mechanism similar to a journal review process. An editor assigns an entry in CSERD to reviewers. The reviewers conduct a VV&A review. The editor uses these reviews to then publish the resource for the community, or to withhold publication until the resources can be improved to meet quality standards. In addition, the community may also submit on-line reviews for all resources in CSERD.

The Verification, Validation and Accreditation process builds on V&V proven strategies for reviewing materials used by many organizations and agencies where:

- Verification provides assurance that the resource works as advertised on the computing platforms as advertised. This helps to answer the question, does this resource "solve the problem correctly".
- Validation provides assurance that the resource is based on current, valid scientific methods. This helps to answer the question, does this resource "solve the correct problem".
- Accreditation provides assurance that the resource is appropriate for the advertised audience. This helps to answer the question, does the resource "match the learner's skill level".

All of the HPC University resources will be subject to the VV&A review process to provide the community with high-quality HPC education and training resources.

2.4 Methodologies for Delivery

The RAT identified potential paths to creating qualified, effective HPC professionals, capable of exploiting current terascale and emerging petascale technologies for the advancement of scientific research. Gaps in training materials and delivery methods were identified and recommendations for filling these gaps were prepared.

2.4.1 Face to Face

Face to face (FTF) training is usually considered the gold standard of training in both professional and academic settings despite the progress that has been made in computer based training methods over the last twenty years. More recently there has been growing evidence of increased success and potential advantages to on-line learning systems. There is a large body of literature investigating the effectiveness and innovative approaches to on-line learning systems, however FTF remains the predominant delivery method. But in the growing global economy and with widespread use of the Internet and technological advances, the need to train more diverse and geographically disperse groups of people has increased.

Twigg [2] describes key factors to developing innovative on-line learning systems, including descriptions of work completed by a variety of academic institutions. Twigg asserts that designing on-line systems that utilize learner-centered training is essential to innovation. Twigg suggests starting with an initial assessment so that students' skill levels and learning styles can be determined, then an array of high-quality interactive learning materials and activities can then be used to build an individualized study plan that includes continuous assessment and instant feedback. Twigg also points out that successful on-line learning systems cannot just put FTF courses on-line, they must redesign courses to adapt to the differing learning methods available on-line; they must "move beyond merely reading text".

Twigg concluded that the cost of developing effective on-line learning systems is far greater than the cost to develop FTF training. The American Federation of Teachers [3] suggests that preparation time for distance education courses may be as much as 66-500 percent longer. A strategy to increase the cost effectiveness of on-line systems is to redesign course development and delivery teams for a more efficient distribution of labor, which would include the use of instructional technologists, tutors, faculty, and information technologists.

Key advantages for FTF training include lower cost for development and widespread acceptance while important disadvantages include higher costs for geographically disbursed participants and more difficult delivering training customized to individual needs.

2.4.1 eLearning

In a recent review of computer mediated communication (CMC) research for education, Luppicini [4] reported that students utilizing CMC education performed at least as well as students in FTF classes. Evidence was presented supporting the idea that students using CMC education experienced less normative pressure, engaged more equally in discussions, and contributed more ideas than students using the FTF methodology. Additionally, CMC education outperformed FTF in critical thinking, personal perspective sharing, and task-oriented interaction. However, the review described indications that FTF students rated group cohesion and group effectiveness higher than CMC students [4].

The use of synchronous instruction techniques is being explored to improve on-line learning. Lee studied the use of synchronous electronic discussions and task-based instruction to improve communication [5]. Hrastinski [6] and Boulos, Taylor, and Breton [7] provided evidence that students' sense of group and sense of participation was higher when using synchronous medium. Delfino and Persico [8] performed a 5-year case study comparing the delivery of a course on educational technology using FTF, on-line, and a combined on-line and FTF delivery, as a blended learning approach. Some of the problems they described include: participants felt the additional workload was a drawback; given their lack of familiarity with the technology, managing a big on-line community was not a simple task; on-line tutoring required time, competence, and commitment; and difficulty recruiting enough tutors to maintain a reasonable tutor/trainee ratio. A blended approach provided a way for the instructional designers and tutors to integrate the best FTF and on-line techniques. Lectures were best for introducing and providing a general framework, while on-line media were best for student-centered activities such as problem-based learning, case studies, and inquiry learning.

Another study compared the use of four learning methods on students' scientific inquiry skills: asynchronous learning networks (ALN) and FTF interaction with and without instruction on how to plan an approach to learning, monitoring comprehension, and evaluating one's learning progress, also known as metacognitive instruction [9]. They found students utilizing ANL with meta-cognitive guided inquiry significantly outperformed all other groups and the FTF without metacognitive instruction group had the lowest mean scores. They posit that the use of metacognitive training within an ALN learning environment demonstrates the advantages of enhancing the effects of ALN on students' achievements in science. An important quote they mention is ″we learn by doing and by thinking about what we are doing″.

Key advantages of on-line learning systems include the ability to create customized training plans, less normative pressure, more equal engagement in discussions, and stronger performance in critical thinking, personal perspective sharing, and task-oriented interaction. Disadvantages include a higher cost of development, less widespread acceptance, technical access and usage issues, and the need for trainers and students to learn how to effectively use on-line learning technologies. For more information and research on distance learning methods and technologies, see references [13-17].

2.5 Scaling Training and Education Efforts

Although most survey respondents claimed to teach themselves what they need, the RAT members felt that often face-to-face (possibly remote) training is useful. There are two modes of doing this teaching (classroom, workshop, or seminar), or collaboration/apprenticeship. Both of these modes have their challenges.

2.5.1 Training Courses

Hanson [10] has set up training courses for many years utilizing computers at national supercomputer centers. He points out several issues that make this a challenge:

- It is often difficult to get user guides for particular computers, compilers, and other necessary software.
- It takes a "super amount of effort" on the instructor's part because of frequent changes in the systems and software.
- There are difficulties coordinating the supercomputer sites from the diverse training locales. Often the instructor is supposed to use Windows computers with the wrong software and no administrator password.

The key to success according to Hanson is to do real "hands on" computing with real problems. To understand how things scale, it is essential that you solve super problems on supercomputers; otherwise the computational overhead gets in the way. A lesser problem can actually show a slowdown when run on a supercomputer. Such a course takes real work on the part of students, with usually a week of sustained effort, to be effective.

2.5.2 Collaboration

Collaboration and apprenticeship can occur over a much longer period, and be directed at the particular scientific field of the user. Catalina Danis [11] is in the Social Computing Group at IBM and did an interesting study of collaboration and learning. She probed the social side of learning how to use a supercomputer. Most users are application scientists who need to solve a problem. She quoted a scientist: "you need to develop a body of knowledge about hardware, memory, protocols. Basically, you need to dig deeper, not in Computer Science, but in Computer Engineering. ...I am no longer a practicing scientist." This points to the issue: optimizing a code can take time away from doing science. Danis proposed solution to this dilemma is to team up a domain scientist with computational experts, preferably with some domain expertise. She divides this into long-term collaboration (being a team) and short-term consultancy.

Short-term consultants are generally located at the computer centers and are assigned to "...help the scientist fix any problems that prevent the code from achieving a production run." The degree of help is a function of the user's skills. Contact is initiated via an e-mail to the center's help desk after the user's proposal to use the center's resources has been approved. Although it is not always necessary, it helps if the consultant has domain expertise for the scientist's problem so that the scientist's intent is not violated. Without domain expertise on the part of the consultant, the scientist must detect whether or not the results seem right. Because the consultant is usually not co-located with the scientist, this introduces inefficiencies of communication and dilutes the opportunities for mutual learning. Again social issues arise. There is often disagreement over the division of labor, and who should have the final say over techniques. One consultant noted "that many users are only interested in getting their code to run, and are unwilling to work to get it to run well. This takes a time commitment." Consultants also try to shape the scientists coding behavior. One consultant refused to look at a scientists code unless he put all the variable declarations in one place. "Users are stuck in the old practices [11]"

Apprenticeship might be somewhat different than collaboration in a team. A student could be assigned a mentor to work with to learn HPC without a strong tie-in to a scientific discipline. An example might be a graduate student running benchmarking tests across multiple computer systems. The student would be taught how to optimize code (e.g., LINPACK) without needing to know where such routines might be used. Another approach to mentoring might enlist something like MentorNet (http://www.mentornet.net). Jim Rome has been a MentorNet mentor for about five years now. This year, MentorNet has been expanded to allow any ACM student member to join. However, thus far the mentoring that is encouraged in MentorNet is more oriented towards career advice rather than actual instruction. Every few weeks MentorNet sends e-mail to its mentors suggesting topics for discussion between the mentor and student. TeraGrid could facilitate the social aspects of mentoring using this technique.

Mentoring works effectively, but more research on the social aspects of such collaboration (e.g., the Danis paper [11]) should be encouraged. We need to factor in the time spent mentoring that keeps a scientist from their science.

2.5.3 Training the Trainers

At least one study [12] has concentrated on "training the trainers." The authors have been teaching HPC to faculty in science, technology, engineering, and mathematics (STEM) keeping the educators current with modern HPC methodologies. They have presented 16 workshops to about 400 mostly undergraduate STEM faculty. To overcome the problems cited in [10], bootable CDs have been created to turn a MS Windows or Machintosh computer lab into a computational cluster in under five minutes. They also built an 8-node cluster for under $3000 that can travel in airplanes to deliver HPC education to any place with an ac outlet. To make this training succeed, participants are required to submit daily pre- and post-workshop surveys, which have allowed on-the-fly course modification over night. Most of these sessions last a week. As was mentioned above, it is necessary to use supercomputers to solve super problems. For example, the study authors simulated galaxy formation, with thousands of point masses, by solving a giant n-body. It is important to solve problems that actually speed up with more nodes, to see the results in real time, and to visualize them. It was found that when participants actually developed computer code they absorbed the material better than passive observers. It motivated the teachers to include more examples of solving realistic problems in their expertise domain. The authors summarize their work as follows: "these introductory examples and our educational computing environment are far from the computing experience that students will have as professionals--the limitations of CPU power, the constraints of shared resources, and the realities of day-to-day management of running jobs will demand that HPC professionals adapt to the hardware at hand, minimize and unnecessary calculations, ... and submit to the thumb-twiddling monotony that is the queue. It is not yet clear what will be the best transition to help students move from introductory to advanced parallel computing."

2.5.4 Facilitate Mentoring

In TeraGrid we have users spanning the whole range of HPC capabilities from novice to expert, each with characteristic skill sets and knowledge gaps. Exciting user's demand for training must be, however, matched by equal efforts to identify and train instructors and mentors. A study of collaboration in the fusion program showed that trust and security is an essential element for success [18]. As discussed in the RAT, there is little incentive for people to share codes compared with the disincentives. Most numerical codes work in limited parameter ranges, yet poor performance or incorrect results are often blamed on the code developer, rather than the misapplication of the code. Argonne National Laboratory (ANL) was the DOE code-sharing repository, but never achieved its full potential due to the absence of a relationship between the user and developer.

A webpage should be created to establish these and other training or mentoring relationships, where HPC experts and novices can each offer and ask for help. A good way to increase the effectiveness of asynchronous training would be to assign mentors to answer questions for specific courses. A mechanism to encourage and provide appropriate motivations for skilled practitioners to provide this mentoring must be devised. One approach that we believe may reward mentors is to compensate mentors by increasing their TeraGrid resource allocation by an amount in proportion to the mentoring effort – similar to the model used by the Amazon.com Mechanical Turk [http://www.mturk.com/mturk/welcome].

2.6 Getting to Petascale

As mentioned in section 2.2, one of the critical drivers for HPC University efforts is building a user population that is functional in petascale environments. It is expected that the transition from tera- to petascale applications will be more challenging than the transition to terascale computing was, since many HPC users are still not scaling even to terascale (e.g. beyond 256 processors) levels.

Issues we have identified include:

- How do we know that the jobs are running efficiently?
- Is there a framework to easily test code scalability?
- Will anticipation of post-petascale architectures dramatically shift programming and science methodologies?

Our recommendations, which complement the other actions we feel are needed for HPC University in general are:

- Find and engage the experts, leveraging groups like the NSF PetaApps winners and the TeraGrid Extreme Scalability Working Group, to suggest directions for these issues.
- Collaborate with the vendors of these systems to provide tools that will make it easier to use petascale computers.
- Include petascale applications in the case study libraries being developed as TG EOT initiatives.

3 Implementation Strategies

The HPC University Virtual Organization (HPCU-VO) is in now in the throes of turning the RAT recommendations into action plans. The team is addressing the following key aspects to ensure that HPC University is responsive to the community's needs.

- HPC and petascale competencies
- HPC Roadmap
- Identifying gaps to be filled
- On-line instruction methodologies
- Quality assurance
- Evaluation
- Community engagement strategies
- Dissemination

The HPCU-VO is building on the computational science competencies developed by the Ralph Regula School for Computational Science, to develop HPC and petascale competencies. Building on these competencies, the team will create a recommended roadmap people can follow to develop HPC skills and knowledge.

To date, over 200 HPC training resources have been identified within NSF, DOE and other national and state HPC centers. We know the list is not complete. The HPCU-VO team will continue to poll HPC centers to identify all of the currently available HPC training resources.

With the competencies and roadmap in hand, an assessment of user needs, and a list of known training resources that are available, the HPCU-VO team will then prioritize new materials and content that should be developed to address the gaps in the available training. We already know from the HPCU RAT that there are a number of gaps to be filled.

The HPCU-VO places a high priority on adapting as many materials to an on-line tutorial format as possible. The objective is to reach more people than can be reached through live events, and to provide just-in-time training when it's needed. The CI-Tutor environment (http://ci-tutor.ncsa.uiuc.edu) is being used for many of the on-line HPC tutorials. Effective on-line methodologies will be followed to ensure that the tutorials are engaging, effective, and well received by the community.

All of the HPC resources will be subject to a thorough VV&A review process to provide the community with assurance that the materials have been reviewed to meet quality standards for HPC training.

Formative and summative evaluation methods will be utilized to ensure that the training resources are meeting the needs of the community, that the training materials are improved based on community feedback, and that additional gaps in community training needs are prioritized for further development. Further, participants in the training will be surveyed and interviewed at least 3 to 6 months after having used the resources to assess the impact of the training on their practices.

The HPCU-VO will pro-actively and continuously engage the community to understand their training needs and requirements, to seek their assistance and reviewers of the materials, and to provide advice on how the HPC training resources can be improved.

Materials that have been developed, reviewed, and evaluated for positive impact by the community will be broadly disseminated for use.

We invite the community to join us in one or more of these areas to enrich the plans and broaden the breadth and depth of the content available to the community.

Acknowledgment

The authors wish to acknowledge the support of the National Science Foundation, the TeraGrid Resource Providers, the HPC RAT leaders Laura McGinnis and Brad Armosky, and the numerous external organizations that have contributed to this work, which are listed in the introduction to this paper. This work was supported by the National Science Foundation Office of Cyberinfrastructure, grant number 0503697 "ETF Grid Infrastructure Group: Providing System Management and Integration for the TeraGrid."

References

[1] The recommendations of the Extreme Scalability RAT and the current endeavors of the working group are available at the webage: http://www.teragridforum.org/mediawiki/index.php?title=Extreme_Scalability_Working_Group

[2] C. A. Twigg, "Innovations in On-line Learning: Moving beyond no significant differences", Tech. rep., The Pew Learning and Technology Program, Center for Academic Transformation, Rensselaer Polytechnic Institute, 2001.

[3] H. E. Program, P. Council, "Distance education guidelines for good practice", Tech. rep., American Federation of Teachers, 1998-2000.

[4] R. Luppicini, "Review of computer mediated communication research for education", Instructional Science vol. 35, pp. 141–185, 2007.

[5] L. Lee, "Enhancing learners' communication skills through synchronous electronic interaction and task-based instruction", Foreign Language Annals vol. 35, pp. 16–24, 2002.

[6] S. Hrastinski, "The relationship between adopting a synchronous medium and participation in on-line group work: An explorative study", Interactive Learning Environments vol. 14, pp. 137–152, 2006.

[7] M. Boulos, A. Taylor, A. Breton, "A synchronous communication experiment within an on-line distance learning program: A case study", Telemedicine Journal of e-Health vol. 11, pp. 583–593, 2005.

[8] M. Delfino, D. Persico, "On-line or face-to-face? Experimenting with different techniques in teacher training", Journal of Computer Assisted Learning vol. 23, pp. 351–365, 2007.

[9] M. Zion, T. Michalsky, Z. R. Mevarech, "The effects of metacognitive

instruction embedded within an asynchronous learning network on scientific inquiry skills", International Journal of Science Education vol. 27, pp. 957-983, 2005.
[10] Floyd B. Hanson, "Local Supercomputing Training in the Computational Sciences Using National Centers", Future Generation Computer Systems, vol. 19, pp. 1335 – 1347, 2003.
[11] Catalina Danis, "Forms of Collaboration in High-Performance Computing: Exploring Implications for Learning", CSCW'06, November 4-8, 2006, Banff, Alberta, Canada.
[12] David Joiner, Paul Gray, Thomas Murphy, Charles Peck, "Teaching Parallel Computing to Science Faculty: Best Practices and Pitfalls", PPoPP'06 March 29-31, 2006, New York, NY, USA.
[13] O. L. N. T. F. on Quality Distance Learning, Quality learning in Ohio and at a distance, Tech. rep., Ohio Learning Network, December, 2002.
[14] "American distance education consortium guiding principles for distance teaching and learning", webpage: http://www.adec.edu/admin/papers/distance-teaching_principles.html.
[15] "Illinois on-line network, pedagogy & learning, webpage", webpage: http://www.ion.uillinois.edu/resources/tutorials/pedagogy/index.asp.
[16] I. Adam Newman, "Eduventures, Measuring success in web-based distance learning, Research Bulletin 4", EDUCAUSE Center for Applied Research, February, 2003.
[17] "Best practices for electronically offered degree and certificate programs, Commission on Institutions of Higher Education", webpage: http://www.neasc.org/cihe/best_practices_electronically_offered_degree.htm.
[18] David P. Schissel, "Grid Computing and Collaboration Technology in Support of Fusion Enegy Sciences", Physics of Plasmas, vol. 12, pp. 058104, 2005.
[19] "The Ralph Regula School for Computational Science is an initiative of the Ohio Supercomputing Center and the Ohio Board of Regents", website: http://www.rrscs.org/.