

The Ranger Virtual Workshop

1 INTRODUCTION

The Ranger Virtual Workshop (VW) is a set of online modules covering topics that help TeraGrid users learn how to effectively use the 504 teraflop supercomputer at the Texas Advanced Computing Center. The Virtual Workshop is not meant to replace live workshops, but rather to complement them. The Virtual Workshop is available at all times and requires no travel budget. Participants may need just a portion of the material or may simply learn better from written materials they use at their own pace. It also serves as a place where attendees at live workshops can go to review material that they did not completely understand during the necessarily short period of the live workshop. The Ranger Virtual Workshop is available to all users via the TeraGrid User Portal.

2 AVAILABILITY

The Virtual Workshop is available to everyone, but registration is required to use the materials. TeraGrid portal users are automatically registered, and can go directly to the materials after login to the portal. Those without a portal account can register at the Virtual Workshop site.

Registration and login is required for two reasons: it allows us to identify TeraGrid users, because we are allocating consulting help only to that group, and it allows us to add functionality for all users, such as note taking and progress tracking. It also gives developers more accurate usage data to base decisions on content and features. For example, data on incorrectly answered quiz questions allow identification of content topics that need to be presented more clearly, and page logs aid understanding of how people use the materials, and which topics are most important.

The primary target audience is academic researchers on TeraGrid who need information on high-performance computing on Ranger. However, given that the content ranges from introductory to advanced topics, a much broader audience could use the materials.

3 TOPICS

There are seven Virtual Workshop modules completed and published with six additional topics to be added over the next quarter. The completed modules are intended to be living documents that grow and change as the implications of running applications on Ranger are more completely

understood. Usage data is recorded for published modules so that under-utilized modules can be identified and to ensure that training effort and expertise is focused on the most relevant areas. The goal of all of the modules is to provide the user with a solid footing for using any HPC system, while ensuring that the specific issues associated with Ranger are addressed.

The screenshot shows the 'Topics' page of the Ranger Virtual Workshop. At the top, there is a navigation menu with links for Home, Topics, Reference, Glossary, Help, and Notebook. Below the menu, the page title 'Ranger Virtual Workshop' is displayed. The main content area is titled 'Topics' and lists several modules with brief descriptions:

- Parallel Programming Concepts and High-Performance Computing**: Concepts concerning parallel processing and its efficient realization within different environments.
- Ranger Environment**: This module provides the information you need in order to start working on Ranger, logging on to Ranger, setting up your environment, moving files to Ranger, and so on.
- Message Passing Interface (MPI)**: MPI is a de facto standard specifying the interface and functionality of a message passing interface (MPI) library. It is a standard for facilitating communication (exchange of data and synchronization) between processes in a distributed memory parallel program. MPI is the first standard and portable message passing interface that provides good performance.
- MPI Point to Point Communications**: This module details and differentiates the various types of point to point communication. Point to point communication involves transmission of a message between a pair of processes. Collective communication, which involves a group of processes.
- OpenMP**: In the shared-memory environment that Ranger has on each node, it is much easier to integrate OpenMP into your code with OpenMP than to do pthread programming from scratch or to use MPI. This module introduces OpenMP and describes how to use it.

Fig 1. The Topics page gives a short description for each module.

The System Information modules include the Parallel Programming Concepts module, the Ranger Environment module, and a planned module on GridFTP. The Concepts module is provided to ensure that all VW users have a common set of terms, concepts and definitions for moving to more advanced modules. This module is code-free and covers general concepts such as hardware terminology (e.g. clusters vs. grids), inter-process communication, and distributed memory systems. The Ranger Environment module is designed to provide the user with all of the information needed to begin using Ranger. The module is highly Ranger-specific and walks the user through accessing Ranger, configuring the environment, moving files to Ranger, compiling applications, and using the Sun Grid Engine (SGE) batch system. These topics are covered by easy-to-follow step-by-step examples as well as multimedia clips to walk users through things like

submitting and monitoring jobs on SGE. All examples were written and tested on Ranger to ensure that the examples apply directly to the system. Source code is provided for several simple applications and scripts that will allow users to verify that their account is fully operational. The GridFTP module addresses many of the issues users might encounter while moving large-scale data around the TeraGrid. The mechanics of this tool are discussed and explained, and a thorough discussion of when and where GridFTP can be used to accelerate the research process. Together the system information modules provide users with all of the information they need to begin understanding and utilizing Ranger.

There are currently two introductory modules on MPI and OpenMP with additional modules planned on Advanced MPI Topics, Hybrid MPI/OpenMP, and OpenMP. The MPI module covers the basics of writing and compiling C/C++ and Fortran MPI programs, and focuses on ensuring that the user understands the general working of MPI. The various implementations of MPI and the functions provided by the MPI standard are explained. This module introduces the user to basic MPI programming and uses a very simple MPI program to demonstrate the key concepts. It focuses on ensuring that the user understands the content, style, and scope of MPI messages and spends the most time on discussing MPI messages, and how these messages function in MPI communicators. Specific information about compiling and running MPI on Ranger is also included. The OpenMP module addresses efficient use of the 16-core Ranger blades using OpenMP. The various pragmas and compiler directives that drive OpenMP are covered along with a complete discussion of their advantages, disadvantages, and pitfalls. This module also shows how to control memory management across individual processes and the environment variables needed to use OpenMP. A Hybrid MPI/OpenMP module will address the issues specific to combining OpenMP for cross-core parallelization and MPI for cross-node parallelization on Ranger.

The Profiling and Optimization section contains the Profiling and Debugging module on the use of common code profiling and debugging tools, as well as Optimization and Scalability. The Profiling and Debugging module covers the basic components of profiling and debugging. This includes how to generate call graphs, listing files, what timers are available, and what data to acquire for profiling. This module provides an in-depth example on using gprof and gmon to identify optimization targets. The Performance Application Programming Interface (PAPI) is also discussed as an effective strategy for analyzing parallel programs. Ranger's sheer size and high core density provide unique opportunities as well as new pitfalls as applications are scaled up. The Optimization and Scalability module covers

some of these issues and highlights some common programming concerns that users should understand as they try to get the most out of the system. The module provides guidance on what scalability is, how to measure scalability, and how to detect and correct problems that may occur as applications move to petascale systems like Ranger.

The five advanced MPI modules which will cover a number of topics applicable to everyday researchers writing parallel code are planned for the future. The topics include point-to-point communication, collective communication, derived datatypes, groups and communicator management, and one-sided communication. In each of the advanced sections, particular care will be taken to ensure that the information is pertinent to Ranger and MVAPICH2, the MPI-2 implementation for InfiniBand used on Ranger, as there are important optimizations provided by the Ranger InfiniBand interconnect. The point-to-point module will cover all aspects of this type of communication, including blocking and non-blocking methods, communication modes, and the Wait/Test functions. Discussions on how to pick a communication mode, comparisons of system and synchronization overhead as a function of communication choices, dealing with deadlock, and using persistent communications will also be provided. The collective communication module will include a discussion of broadcast, scatter/gather methods, all-to-all, and global computing. This module will also focus on performance issues related to scaling of the collective communication model to petascale-size problems, with an emphasis placed on scaling examples performed on Ranger. All MPI-2 additions to collective communication will be discussed in this module. Additional modules will cover the details needed to create custom datatypes and communication groups. These modules will include detailed examples focusing on practical situations where using these features are reasonable and beneficial. The one-sided communication module will cover this MPI-2 topic with details on implementation provided by MVAPICH2. In each of these sections, functional example code will be available to demonstrate functionality as well as pitfalls. Each of these modules will also provide in-depth programming examples and exercises for the reader. The advanced MPI topics will focus on issues relating to scaling the various communication methods and protocols to petascale problems and will make heavy reference and use of topics discussed in the Scalability module.

An additional module on visualization topics is also planned. This module will discuss both the tools and methods which can be used for both interactive and non-interactive visualizations. The focus will be on configuration and use of tools available or practical for use on Ranger or other petascale systems.

4 FEATURES

The VW includes a number of features to support the learning environment.

The navigation is designed to be self-explanatory and shallow. After login, 99% of the pages can be reached in three clicks or fewer. General reference pages are available from all pages via the top navigation bar. The topics are arranged in suggested order, but sequential progression is not enforced. Within a module, the participant can either progress page by page or jump to any page. The current page is marked in the table of contents.

The Reference page (see Fig. 2) gathers bibliographic and Web references for all topics in one place, sorted by topic. All content pages link to the references as appropriate. The Reference page is complete, but continues to grow as additional relevant information is found and as topics are added to the Virtual Workshop.

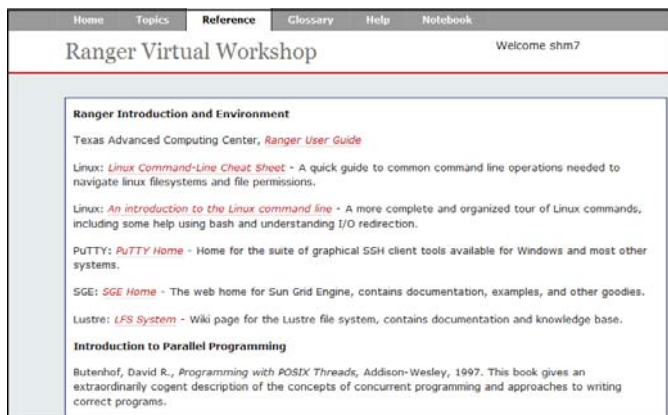


Fig 2. The Reference page lists references for all topics.

The Virtual Workshop includes an HPC Glossary. The Glossary covers terms specific to Ranger as well as general HPC terms for completeness. The Glossary has about 50 terms and continues to grow. The alphabetical Glossary is available from all pages via the top navigation bar (see Fig. 3). Glossary terms are also linked as appropriate in the text. The links pull the definition into the top of the page for an in-line reference (see Fig. 4).

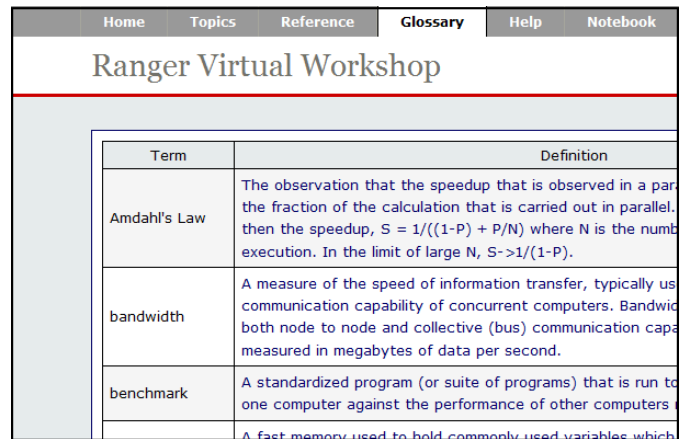


Fig 3. The full Glossary is available via the top navigation bar.




Fig 4. Using a link to a glossary page pulls the definition in to the page.

Participants have two tools to reinforce and check what they have learned: exercises and quizzes. Exercises are practical labs, where the participant is given instructions to try out what they have learned. All exercises are written to be run on Ranger (see Fig. 5). Consultants are available for questions on the exercises as well as on the main content by either using links inside the Virtual Workshop or by submitting a ticket to the TeraGrid User Portal help desk system.



Fig 5. Example of a lab exercise.

Interactive quizzes presented as multiple-choice questions will be added in the next few months. Quizzes are automatically graded for the users. Incorrect answers are returned with the correct answer, a short explanation of why it is the best answer, and a link to the corresponding material. The results are also written to a database, enabling the content developers to improve materials for concepts that are commonly misunderstood, allowing effort to be focused appropriately.

To address differences in the way people learn and the occasional difficulty in following written instructions, some procedures are also available as flash movies with audio narration. In the sections where a demonstration is appropriate, screen recordings with audio narration are included (see Fig. 6).

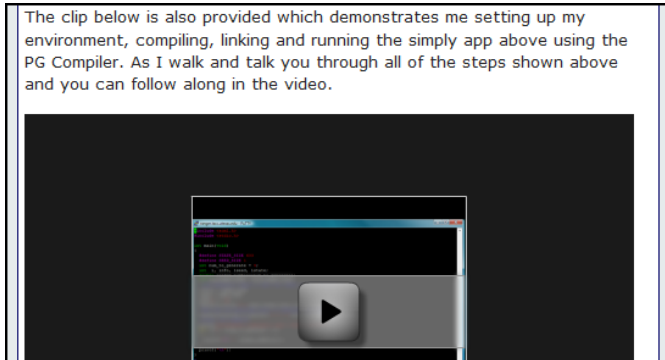


Fig 6. Demonstrations are included as short video clips.

Participants have the opportunity to customize their topic pages. The participant can mark their progress through each topic with a Virtual Workshop-specific bookmark. When they return to the Virtual Workshop, they can easily find their “marked page” for each module. They can also take notes, or “write in the margins” on as many pages as they choose. When a note is added to a page, that note will appear on the page on subsequent visits (see Fig. 7), but only for the person who created the note. Notes can also be updated or erased at any time. If feedback shows sufficient interest, the page note capability will be expanded to enable note sharing with other specified participants.

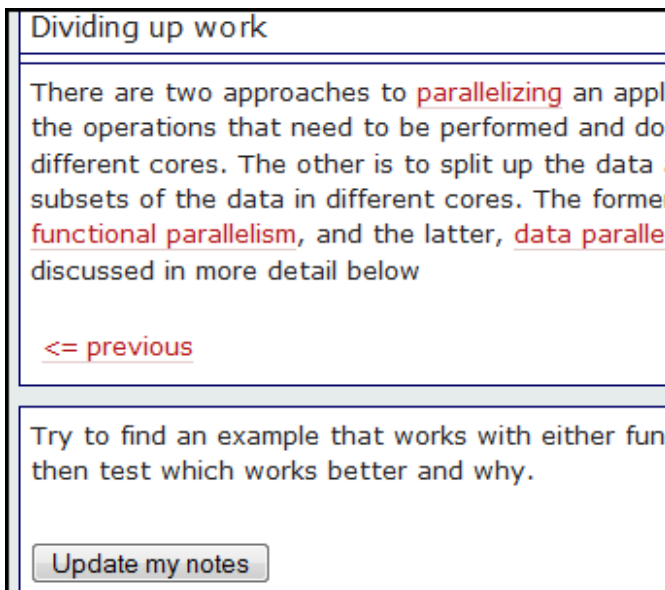


Fig 7. This example shows a note that has been added to a page.

In past Virtual Workshops, participants have asked for a way to view information on how much of the course they have completed. In response, a Notebook page has been provided, which includes a summary of their notes, bookmarks, questions asked, and quizzes submitted, so they can easily review their progress and jump to points of interest (see Fig. 8).

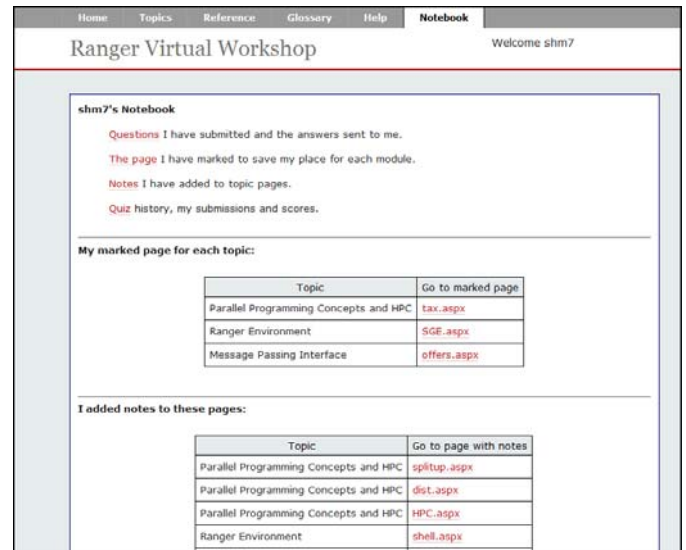


Fig 8. The Notebook page includes information on notes, bookmarks, and completed exercises and quizzes.

5 ASSESSMENT PLANS

The Virtual Workshop was designed based on experience developing and delivering Web-based training in High Performance Computing topics since 1995. The development model has been to incorporate new features each year, evaluate, and retain the features that enhance the workshop. Experimentation in past workshops included virtual office hours, module walk-throughs, scheduled progression through the topics, and using video and audio clips. Usage of all features in the Ranger Virtual Workshop will be tracked. Special attention will be paid to understanding which content and features are most heavily used, how participants progress through the materials, and determining whether participants find the personal notebook valuable. Later in the program experiments will be added to determine whether participants are interested in adding public content and comment via a Wiki. Evaluation will be based on usage data and annual surveys.

6 FUTURE WORK

The Ranger Virtual Workshop has been in production since March, 2008. At that time it included a core set of introductory modules and all of the basic feedback and note-taking features. Content and features continue to be added and improved. Existing topics will be expanded and improved based on feedback from questions,

comments, experience running performance tests, and analysis of tools such as quiz results.

As mentioned earlier, graded quizzes will be added for interactivity and immediate feedback, and a Wiki will be added for user contribution and discussion. Audio and video clips will be added to the content modules. The clips will be authored by content experts and will be used to enhance a portion of a topic rather than provide a full presentation, e.g. the majority of a module is text, with a video included to introduce the topic or explain an important or difficult concept.

7 CONCLUSIONS

Existing topics will be augmented and new topics will be added to the Virtual Workshop. Usability will be improved by adding a Wiki and print options. Usage tracking began in March 2008, and future development will be guided by the results.

ACKNOWLEDGMENTS

The authors wish to thank everyone at TACC for sharing their materials, feedback, time, and expertise with us in the development of the materials.

REFERENCES

Allen, Michael W. (2003). *Michael Allen's Guide to e-Learning*. Hoboken, NJ: John Wiley & Sons, Inc.

Mehringer, Susan H. (2002). Enhancing Web-Based Materials with Video Clips. *E-Learn 2002 World Conference on E-Learning in Corporate, Government, Healthcare, & Higher Education. Proceedings (7th, Montreal, Quebec, Canada, October 15-19, 2002)*.

Weston, T. & Barker, L. (2001). Designing, Implementing, and Evaluating Web-Based Learning Modules for University Students. *Educational Technology*; v41 n4, 15-22.

Williams, R. (2002). *Web Design Workshop*. Berkeley, CA: Peachpit Press.

Willie, S., Pui, C., & Palmer, A. (1999). Remote Lecture Presentation Preferences for Internet Delivered Continuing Medical Education. *Proceedings of ED-MEDIA 99 World Conference on Educational Multimedia, Hypermedia & Telecommunications, 1999*, Association for the Advancement of Computing in Education, Charlottesville, VA. 190-195.