



University of Puerto Rico, Río Piedras
 New Natural Sciences Building Phase II
 XSEDE Workshop Agenda

Thursday, January 29, 2015

8am	Sign-In, Courtyard adjacent to the Computer Science Department, Natural Sciences Building Phase II, A-150		
9:00am	Welcome, Room A-211 Amphitheatre UPR High Performance Computing - Humberto Ortiz-Zuazaga XSEDE Overview – Scott Lathrop		
Noon	Lunch & Consulting		
1:00pm – 5:00pm	<i>XSEDE Campus Bridging</i> Instructor: Barbara Hallock Room: A-140	<i>Getting Started with HPC: XSEDE New User Training</i> Instructor: Jay Alameda Room: A-141	<i>Linux/Unix Basics for High Performance Computing</i> Instructor: Daniel Lucio Room: A-143
5:00pm	Consulting		

Friday, January 30, 2015

7:30am	Sign-In, Courtyard adjacent to the Computer Science Department, Natural Sciences Building Phase II, A-150		
8:00am – Noon	<i>Introduction to Parallel Computing with MPI and Open MP</i> Instructor: Daniel Lucio Room: A-141	<i>Getting Started with HPC: XSEDE New User Training</i> Instructor: Jay Alameda Room: A-143	
Noon	Lunch*		
1:00pm – 4:30pm	<i>CUDA™: GPU Programming</i> Instructor: Vince Betro Room: A-141	<i>Bioinformatics Topics</i> Instructor: Marcela Madrid Room: A-143	

Session Descriptions

XSEDE New User Training

This tutorial will help new users become comfortable with the steps needed to access XSEDE resources, and to successfully use those resources to accomplish their research or educational goals. Topics will include how to request an allocation, find documentation, ask for help, login and run on an XSEDE supercomputer, submit a job and troubleshoot a job that has not run. The practice section of the tutorial will consist of hands-on activities including submitting a job, figuring out why it has not run and transferring files between supercomputers. Participants who are unfamiliar with Unix will benefit from attending the Linux/Unix Basics tutorial beforehand, but if unable to do so, we will provide personalized help during this session.

Linux/Unix Basics for High Performance Computing

An introduction to the Linux and Unix basics will be provided in this session. An overview of interacting with clusters and HPC resources will also be provided through interactive teaching. While there are no labs associated with this session, participants will be able to follow along with the instructor to gain experience and familiarity with the Linux/Unix environment.

Introduction to Parallel Computing with MPI and OpenMP

An overview of a typical XSEDE high performance computing (HPC) resource and the user-environment will be provided in this session. There are no labs associated with this session. The topics that will be discussed are:

- HPC System Overview
- Theoretical Background
- Parallel computing systems
- Parallel programming models
- MPI/OpenMP examples
- HPC User-Environment

CUDA™: GPU Programming

This training module is a beginners/intermediate course on programming NVIDIA GPUs with CUDA. After a short segment on why we are using accelerators in High Performance Computing and on accelerator hardware I will describe all pieces necessary to write a program in C and FORTRAN. The example will be a stencil update, which is simple enough to be written in a few lines of code. The specific code design will be guided by the hardware with emphasis common design principles for writing fast code for GPU accelerators. The second part of the presentation will focus on two common optimization strategies, namely the use of shared memory and data streams. Some experience with writing serial code in C or FORTRAN will be helpful to follow the examples.

Introduction to Scientific Visualization

This session will discuss the visualization systems available to XSEDE users. This will include an overview of the hardware on-site, the services available for remote-access visualization, and the software available for scientific visualization, including Visit and ParaView. A hands-on training session for ParaView will be offered. Discussion of information visualization will also be included as interest indicates.

Bioinformatics

We will cover bioinformatics topics such as databases, structural bioinformatics, molecular dynamics and De Novo RNA-seq assembly using Trinity. During the hands-on exercises, we will help you run your own molecular dynamics simulation and Trinity job on an XSEDE supercomputer. Participants will benefit from attending the Linux/Unix Basics tutorial beforehand, but if unable to do so, we will help you get started during this session.

XSEDE Campus Bridging

This session for system administrators will provide an overview on the concept of campus bridging as well as how the XSEDE Campus Bridging team works toward its specific goals. Once we have discussed what Campus Bridging is, demos of several key bridging technologies including the XSEDE-Compatible Basic Cluster software stack, Globus Online, and the Global Federated File System. Bring your laptop.

Consulting

Have questions, need help? Bring your project ideas, code problems, questions about tools and resources and XSEDE staff will be there to answer. We can help you get started or solve a problem with an existing project using High Performance Computing.

Speaker Bios

Scott Lathrop of the Shodor Education Foundation, Inc. is the XSEDE Director of Education and Outreach and the Blue Waters Technical Program Manager for Education. Lathrop has been involved in high performance computing and communications activities since 1986. Lathrop coordinates education and outreach activities among the XSEDE Service Providers involved in the NSF funded XSEDE project. He coordinates undergraduate and graduate education activities for the Blue Waters project. Lathrop has been involved in the SC Conferences since 1989, has served as a member of the SC Steering Committee for more than six years, and was the SC11 General Conference Chair.

Jay Alameda is the lead for Advanced Application Support at the National Center for Supercomputing Applications. In this role, he works with the Extreme Science and Engineering Discovery Environment (XSEDE) which is a collaboration of NSF-funded high performance computing (HPC) resource providers, working to provide a common set of services, including the provisioning of advanced user support, to the science and engineering community. In particular, Jay leads the Extended Support for Training, Education, and Outreach Service of XSEDE, which provides the technical expertise to support Training, Education, and Outreach activities organized by XSEDE. Jay also works with the NSF-funded Track 1 project, Blue Waters, and in this role, has worked with advanced development tools (such as the Eclipse Parallel Tools Platform) to support development and optimization of HPC applications on the Blue Waters resource. He is also leading the NSF funded SI2 project, "A Productive and Accessible Development Workbench for HPC Applications Using the Eclipse Parallel Tools Platform", which is working on a user- and application-centric plan to improve Eclipse PTP as a platform for development of HPC applications, with a particular focus on broadening support of a diverse range of HPC resources (especially across XSEDE).

Dr. Vincent Charles Betro received his Ph.D. in Computational Engineering from the University of Tennessee SimCenter at Chattanooga in 2010, where he became research faculty and the STEM Outreach coordinator. Since 2012, he is a Computational Scientist at the University of Tennessee National Institute for Computational Sciences at Oak Ridge National Laboratory, where he focuses his research on porting and optimizing applications for several accelerator architectures and developing Computational Fluid Dynamics codes for the Application Acceleration Center of Excellence. Additionally, due to his background as a middle and high school mathematics teacher and college mathematics and engineering instructor, Vince enjoys working with students in his community to broaden their understanding of and interest in STEM careers and is actively involved in leadership in the XSEDE User Engagement and Campus Champions programs.

Barbara Hallock serves as Senior NET+ / Campus Bridging Systems Analyst at Indiana University, where her duties focus mainly on XSEDE Campus Bridging projects. She completed her Bachelor of Arts in Digital Media summa cum laude in 2008 from Saint Mary-of-the-Woods College in Indiana and then went on to receive her Master of Information Science from Indiana University in 2012. Her responsibilities include providing support by phone and e-mail for XSEDE Campus Bridging products and participating in various outreach activities. Prior to beginning her work on the XSEDE project, Hallock worked in technical support and then web development.

Daniel Lucio is a HPC consultant at the National Institute for Computational Sciences at Oak Ridge National Laboratory. He has more than fourteen years of experience working with Supercomputers and HPC resources. He has spent the last seven years working at NICS and being involved at several TeraGRID/XSEDE events. He loves training and helping other people to succeed. He is constantly looking for new ways to bring knowledge for the masses and get minorities to get engaged in science and engineering.

Marcela Madrid is a Senior Computational Scientist at the Pittsburgh Supercomputing Center. She specializes in computational biology and structural bioinformatics. She collaborates with other scientists with the aim of applying high performance computing to the advancement of scientific problems and teaching. Using quantum mechanics and molecular dynamics simulations, she has studied hemoglobin, HIV-1 Reverse Transcriptase, Peptide Nucleic Acids and heat transfer in Silicon thin films. She develops and delivers training and outreach activities on XSEDE leading-edge high performance computers. Dr. Madrid holds a Ph.D. in Physics from Instituto Balseiro, Argentina.