

Reference Materials for Customers of ASC's Unity Cluster

John W. Heimaster – April 25, 2021

Arts and Sciences' Unity system is a small supercomputer. Unity is a Linux cluster providing batch computing service to the general College community, plus additional services to groups that have purchased private computers to augment Unity. This document will help you learn to use Unity effectively.:

General background. Slides 20-72 of [What the Heck is Supercomputing?](#) from the *Supercomputing in Plain English* project at the University of Oklahoma explain the basic concepts of batch supercomputing. Some of the examples are dated but the diagrams are clear.

Basic online courses.

- Collections of training materials: [XSEDE catalog](#) [PSU catalog](#)
- Linux basics [Part 1](#) [Part 2](#)
- The [Unix shell](#)
- [Regular Expressions](#)
- Introduction to C [Part 1](#) [Part 2](#)
- [C++ programming](#) and [C++ for C Programmers](#)
- Introduction to Fortran [Part 1](#) [Part 2](#)
- Introduction to [Python](#) and [Python for high performance](#) and [Distributed Parallel Computing in Python](#)
- [R Tutorial](#) and [Parallel R](#) and [Hadley Wickham's Advanced R book](#)
- Introduction to Matlab [Basic](#) [Detailed](#) [Video](#)
- Introduction to Mathematica [Basic](#) [Video](#)
- [Use of make](#)
- [Git Tutorial](#) (version control software) - - see also OSU's git repository [code.osu.edu](#)
- [GPU programming](#)
- [Parallel Programming Concepts](#)
- [MPI - Parallel programming across multiple nodes](#)
- [OpenMP - Parallel programming on multiple cores of a single node](#)
- [Introduction to Data Management](#)
- [Parallel Input and Output](#)
- [Introduction to Deep Learning](#)
- [The Art of File Transfers](#) and [Globus data transfer](#)
- Introduction to Slurm [Simple examples](#) [Vendor videos](#)

The XSEDE catalog contains discipline-specific introductions for customers from chemistry, physics, engineering, and the life sciences. Both catalogs contain advanced material on parallel computing, debugging, performance tuning, visualization, GPUs, and input/output techniques.

Ohio Supercomputer Center. Unity's computing environment closely resembles the larger facilities of the Ohio Supercomputer Center, facilitating interchange of programs and data. Much of OSC's [introductory documentation](#) is applicable to Unity, especially the [HOWTOs](#) .

Free HPC Textbooks. The Texas Advanced Computing Center publishes two free HPC textbooks at <http://pages.tacc.utexas.edu/~eijkhout/istc/istc.html> . The first book covers a range of scientific computing topics from numerical mathematics to computational biology. The second book covers parallel computing, with emphasis on MPI and OpenMP. A third book will cover C++ and Fortran 2008.

Best HPC Textbook. *High Performance Computing – Modern Systems and Practices* by Thomas Sterling et al., Morgan Kaufmann 2018.

Free Data Science Textbook. Berkeley's [Fundamentals of Data Science](#) course provides an almost-painless introduction to data handling and simple statistics, with an excellent online textbook.

Small Introductory Books. The [MIT Press Essential Knowledge Series](#) has several dozen well-written introductory books on such topics as Deep Learning and Data Science. The books are small and inexpensive, and the collection is rapidly-expanding.

Intense Online Courses. The Argonne Training Program on Extreme Scale Computing produces several dozen excellent 30- to 90-minute training videos each year. Expert presenters provide in-depth coverage of HPC topics across hardware, applications, visualization, and performance. The [2019 slides and videos](#) are now available. These presentations are deep and fast-paced, but worth the effort.

A sobering example. Good supercomputer performance requires well-written code. Slides 18-69 of [this lecture](#) from MIT OpenCourseWare course 6.172 applies successive optimizations to a matrix multiplication routine originally written in Python. The resulting well-tuned C program is more than 50,000 times faster.

For additional information. Consult [ASCTech's Research Consulting site](#) to learn about Arts and Science's support for your research computing needs, or send a message to asctech@osu.edu to ask a question or meet with a consultant. Consult [OSU Research Computing Resources](#) to find other campus services to support your work. Consult ASCTech's extensive [HPC bibliography](#) to learn about additional materials, including web sites and periodical literature devoted to research computing topics.