

High-Performance Computing at NC State



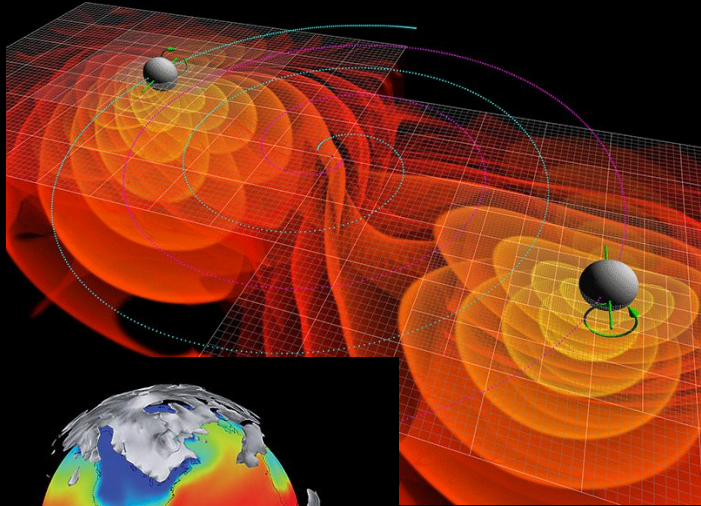
**LUG meeting
September 22, 2020
Lisa L. Lowe, PhD**

High-Performance Computing at NC State


- What is HPC and what is it used for?
- Introducing Henry2
- Our Users
- Challenges
- Outreach projects for your consideration

What is HPC and what is it used for?

High Performance Computing...

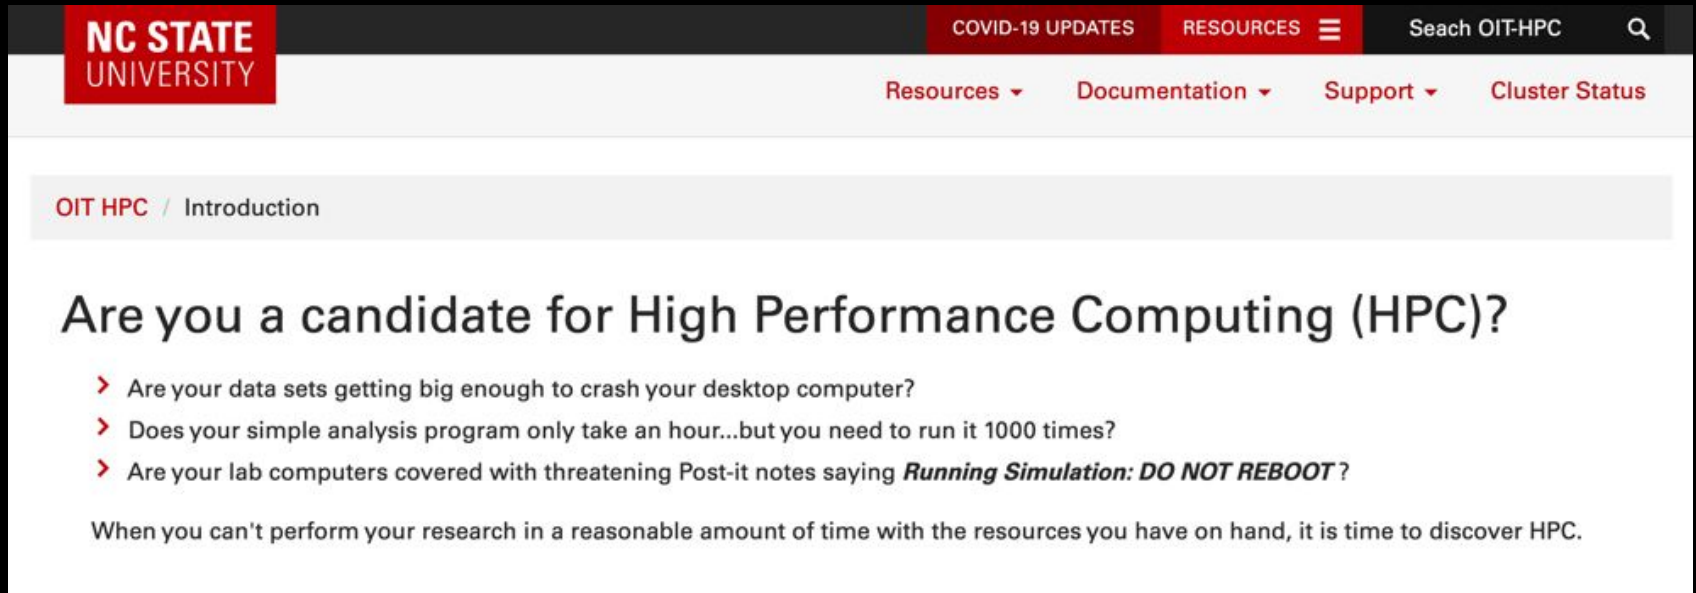


Before, it was only used for this stuff!!



Today, our HPC is actually used for:

Anything that can't be run on local resources due to time constraints or compute power.



The screenshot shows the top navigation bar of the NC State University OIT HPC website. The navigation bar includes the NC State University logo, a search bar labeled "Search OIT-HPC", and a menu with "RESOURCES" and "COVID-19 UPDATES". Below the navigation bar, there are links for "Resources", "Documentation", "Support", and "Cluster Status". The main content area features a breadcrumb trail "OIT HPC / Introduction" and a large heading "Are you a candidate for High Performance Computing (HPC)?". Below the heading is a list of three bullet points with red arrowheads, and a concluding paragraph.

NC STATE UNIVERSITY

COVID-19 UPDATES RESOURCES

Search OIT-HPC

Resources Documentation Support Cluster Status

OIT HPC / Introduction

Are you a candidate for High Performance Computing (HPC)?

- > Are your data sets getting big enough to crash your desktop computer?
- > Does your simple analysis program only take an hour...but you need to run it 1000 times?
- > Are your lab computers covered with threatening Post-it notes saying *Running Simulation: DO NOT REBOOT?*

When you can't perform your research in a reasonable amount of time with the resources you have on hand, it is time to discover HPC.

High-performance

Perform calculations in the fastest way possible using additional compute power

Additional compute power includes

- Multiple cores per computer
- Multiple computers networked together
- Fast networks
- GPUs

From here, computer == node

What's in a node?

Node

- A standalone "computer in a box"
- Comprised of
 - Multiple CPUs/processors/cores
 - Memory
 - Disk
 - Network interfaces
 - GPUs, FPGAs, etc.


Cluster

- Nodes are networked together to comprise a cluster.



Henry2





Network Fabric



Processor Model



GPU



Henry2 is an Intel Xeon based Linux cluster, and compute nodes are heterogeneous in processor model, number of cores, amount of memory, instruction set architecture (ISA), attached GPU, and network fabric.

Who uses HPC?

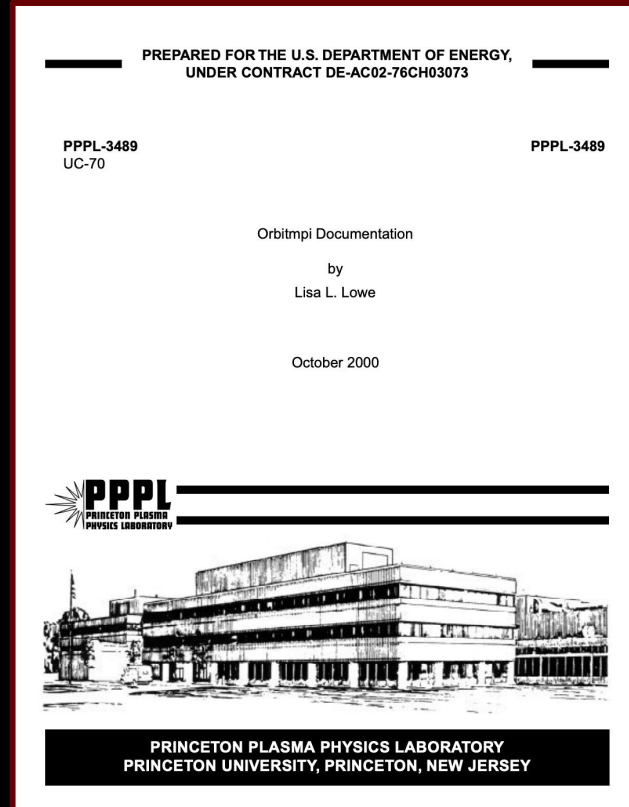
Before

- Physicists with training in numerical methods
- FORTRAN or C
- MPI (Message Passing Interface)

BTW...Today is my anniversary...20 years since I parallelized my first legacy FORTRAN code...

Orbitmpi Documentation

Lisa L. Lowe
Princeton Plasma Physics Laboratory
(September 22, 2000)



Who uses HPC now?

Everyone!

- Economists running **SAS**
- Sociologists running **R**
- Mechanical engineers running **MATLAB**
- Statisticians running **Python** (Tensorflow) on GPUs
- Veterinarians running **Perl, Java**
- Bioinformaticians running **all of these** (pipelines)
- Physicists running **FORTRAN** and **MPI**


Many of them

- Have always used a GUI interface (e.g. RStudio) on Windows or Mac
- Have only installed software by clicking an installer
- Have no Linux or command line experience
- Are **really annoyed** that they have to know what "node" means, let alone worry about *Instruction Set Architecture*

What are the challenges?

First hurdle:

- Learn Essential Linux
- Learn how to log in
- Learn how to move files
- Learn how to run basic jobs



YouTube channel page for OIT HPC (65 subscribers). The main video is "Henry2: Basic HPC Workshop: Parallel Jobs" (553 views, 5 months ago). Below are several other videos from the "Henry2" series, including "Getting HPC Access", "Introduction", "QIMEZ Moving Pictures Tutorial on the HPC", "Essential Linux Commands - Fluency Drill", and another "Parallel Jobs" video.

We have created lots of support to overcome the first hurdle.

What are the challenges?

Next hurdle:

INSTALLING SOFTWARE

Why is it hard?

- Can't click on an installer
- No, you can't move binaries from your Mac and run them on CentOS
- No elevated privileges, i.e. no `sudo yum install`
- Need to install from scratch, i.e. `compile`
- Compilers, flags, environment variables, oh-my!

And even if it compiles...

- One program needs CUDA X on GPU Y, crashes on some GPUs
- Another needs AVX2 instruction set architecture, crashes on some nodes
- Which compiler flags to optimize for performance? (It *is* HPC after all...)

Outreach Projects

- 1) **Linux Workshops: Beyond "Essential Linux"**
 - LUG could volunteer to lead some help sessions or workshops for new Linux users.
 - I have [basic material](#) for such a workshop.
 - Workshops on Git or scripting would also be in high demand.

- 2) **Installfests...but not like you're used to**
 - LUG could volunteer to help users install software on HPC or to build containers for HPC.

Good publicity for LUG...and good for the resume!

Basic Linux for using HPC

Need this to squeak by - Essential Linux

Topics:

- Basic Linux commands
- Key commands and shortcuts
- Parsing output
- An introduction to environment variables
- The text editor Vim

Need this to be efficient and maybe even *enjoy* the command line

Essential Linux

`pwd` - print working directory

`ls` - show contents of a directory

`cd` - change directory

`mkdir` - make a directory

`cp` - copy

`mv` - move or rename

`rm` - remove

`cat` - display the contents of a file

`less` - examine the contents of a larger file, use [space-bar] to page through, type [q] to quit.

`nano` - a text editor

This is all the Linux a user needs to run basic jobs on the HPC!

Installing Software

HPC provides some [basic guidance for installing software](#). We ask users to install their own software, and we help when they get stuck. If it turns out we need to do the entire install for them, it can take a long time...that leads to unhappy customers!

Basic methods of installation

- Download precompiled binaries
- Use Conda (usually precompiled binaries, packed with libraries and other dependencies)
- Compiling (usually configure/make/cmake, but could include modifying makefiles)

Another method of installation - Containers!

- HPC uses Singularity, but a Docker image can be converted to a Singularity image
- I have very little experience with this! But LUG members install software on their Linux machines all the time...that process can be saved in a container

Good way to learn about research groups...help them install their software!

Links

[Slides for a workshop](#) on learning Linux, including vi

HPC documentation - guidance for [installing software](#)

Job Description - [HPC Software Installation Intern](#) (*This [OIT PMLite](#) position was cancelled...maybe next year?*)

Instructions for [requesting HPC access](#), including an application form for users without a faculty advisor (undergraduates)

HPC's [YouTube Channel](#)

Questions, Comments, Concerns?

oit_hpc@help.ncsu.edu