

Introduction to High Performance Computing

Feng Chen

IT Analyst, User Services

HPC@LSU

Part of slide taken from Dr. Alexander Pacheco and Dr. Le Yan

Outline

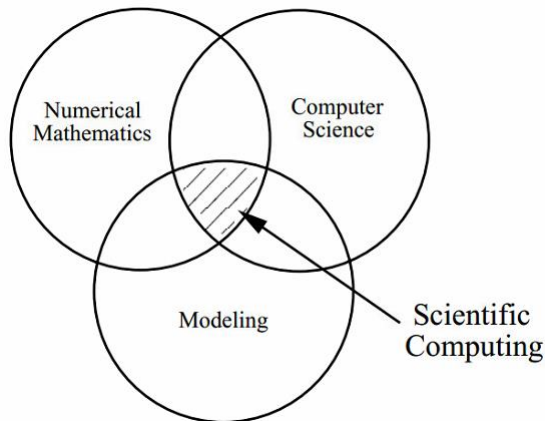
- What is HPC – background and definitions
- Science & Engineering Applications
- Available HPC Resources
- Where to get started - HPC @ LSU and LONI
- Wrap-up

Outline

- What is HPC – background and definitions
- Engineering & Science Applications
- Available HPC Resources
- Where to get started - HPC @ LSU and LONI
- Wrap-up

Computational Science

- Apply computational methods to gain insight not available through theoretical analysis and experiment.
- Began in earnest in the late 1950's for defense related work. Came into it's own in the 1970's.
- Typically, these models require massive amounts of calculations (usually floating-point) and are often executed on supercomputers or distributed computing platforms.



G. E Karniadakis & R. M. Kirby II, (2003),
 Parallel Scientific Computing in C++ and
 MPI: A Seamless Approach to Parallel
 Algorithms and Their Implementation

What is HPC

- High Performance Computing (HPC) is computation at the cutting edge of modern technology, often done on a supercomputer
- A supercomputer is in the class of machines that rank among the fastest in the world
 - Rule of thumb: a supercomputer could be defined to be at least 100 times as powerful as a PC
- How do you evaluate the performance of HPC?

Measure HPC performance-FLOPS

- Performance is measured in **F**loating point **O**perations **P**er **S**econd (FLOPS or flop/s)
- $$FLOPS = cores \times clock \times \frac{FLOPs}{cycle}$$
 - Most microprocessors today can do 4 FLOPs per clock cycle. Therefore a 2.5-GHz processor has a theoretical performance of 10 billion FLOPs = 10 GFLOPs
 - Dual Core? Quad Core?

Computer performance

Name	FLOPS
yottaFLOPS	10^{24}
zettaFLOPS	10^{21}
exaFLOPS	10^{18}
petaFLOPS	10^{15}
teraFLOPS	10^{12}
gigaFLOPS	10^9
megaFLOPS	10^6
kiloFLOPS	10^3

Supercomputing on a cell phone?


- Quad-core processors are coming to your phone
 - Nvidia, TI, Qualcomm...
 - Processing power in the neighborhood of 10 GigaFLOPS
 - Would make the top 500 list 15 years ago
 - What is your phone's FLOPS?
 - iPhone 5 1.3 GHz dual-core Cyclone
 - Compare to ENIAC (500 FLOPS)
 - Compare to top 500 in 1993 #1 (59.7 GFLOPS), #500 (0.42 GFLOPS)



What is driving the Change?

- What is driving this continuous change?
 - Moore's Law
 - What was a supercomputer 15 years ago now sits on your desk, or even in your hand.
 - Who measures/ranks the supercomputing performance?

Microprocessor Transistor Counts 1971-2011 & Moore's Law



Source: http://en.wikipedia.org/wiki/Moore's_law



The Top 500 List


- The TOP500 project provides a list of 500 fastest super computers in the world ranked by their LINPACK performance.
- Semi-annually published (in the public domain)
- As of Nov 2013, China's Tianhe-2 (Milky Way-2) supercomputer is the fastest in the world.
 - Nodes: 16,000
 - Cores: 3,120,000
 - Peak Performance: 33.86 PFlop/s





The Top 500 List

Performance Development




Source: <http://www.top500.org/statistics/perfdevel/>



The Top 500 List

Projected Performance Development

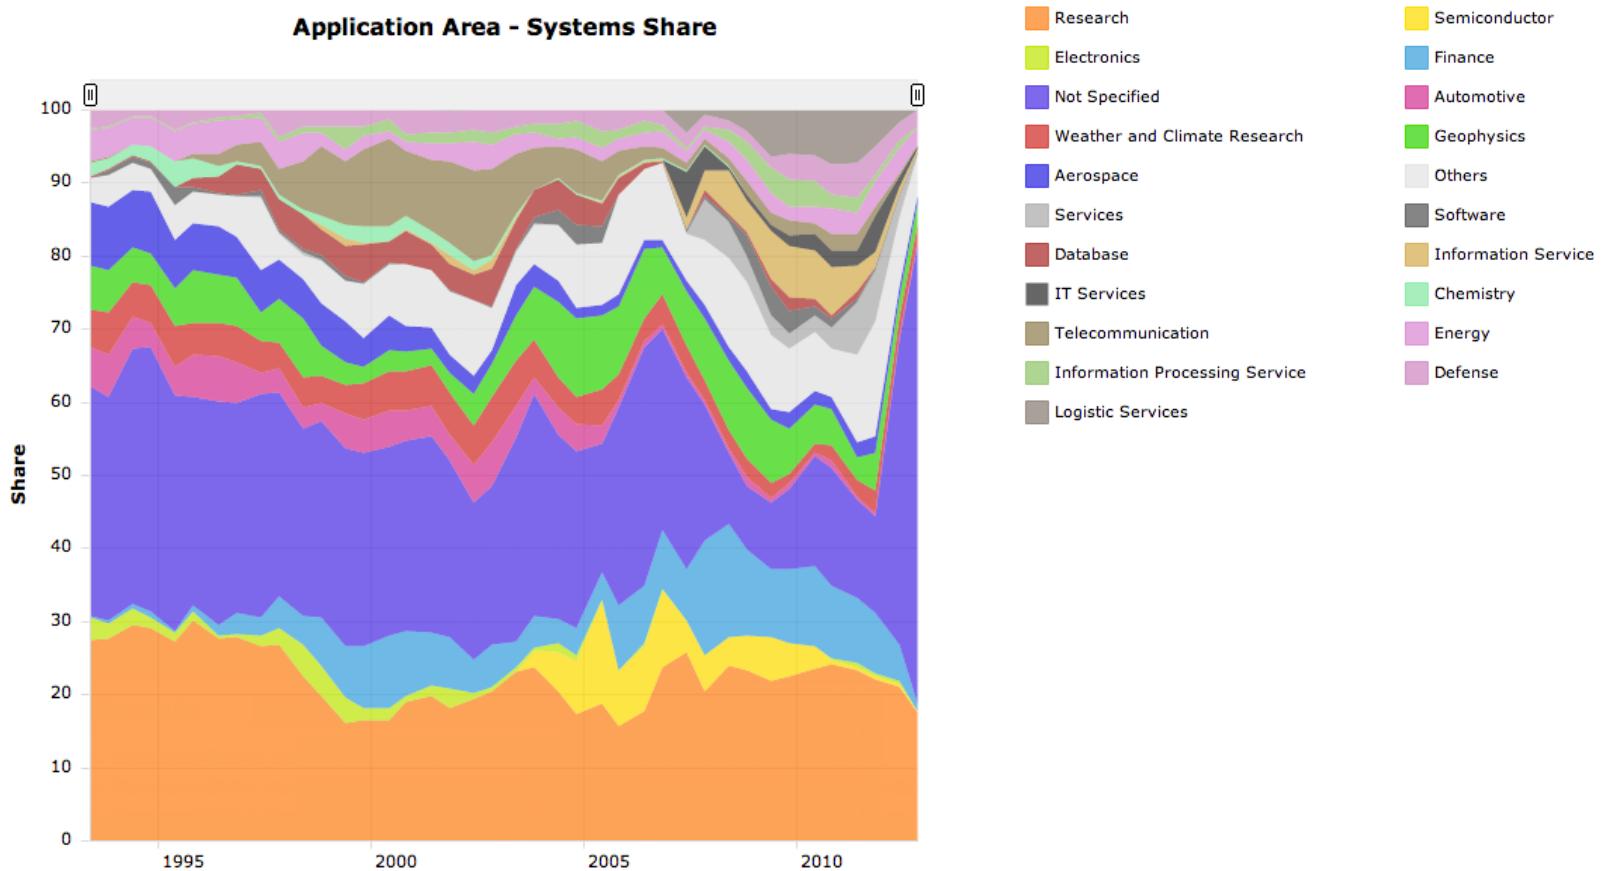


Source: <http://www.top500.org/statistics/perfdevel/>



The Top 500 List

Application Area - Systems Share




Source: <http://www.top500.org/statistics/overtime/>



The Top 500 List

Countries - Systems Share




Source:

<http://www.top500.org/statistics/overtime/>



The Top 500 List

Operating System - Systems Share




Source:

<http://www.top500.org/statistics/overtime/>



The Top 500 List

Operating system Family - Systems Share



Source:

<http://www.top500.org/statistics/overtime/>

What are HPC Strength Problems?

Size



Speed






HPC Problem Discriminators?

- Size
 - Many HPC problems **can't fit on a PC** - usually because they need more than a few GB of RAM, or more than a few TB of disk
- Speed
 - Many HPC problems **would take a very long time to run on a PC**: months or even years, but a problem that would take a month on a PC might only take a few hours on a supercomputer



What Does HPC do? (1)

- Simulation of physical phenomena
 - Storm surge prediction
 - Black holes colliding
 - Molecular dynamics




What Does HPC do? (2)

- Visualization
 - Turning (huge amount of) data into pictures and/or animations that scientists can better understand



Using HPC for Design

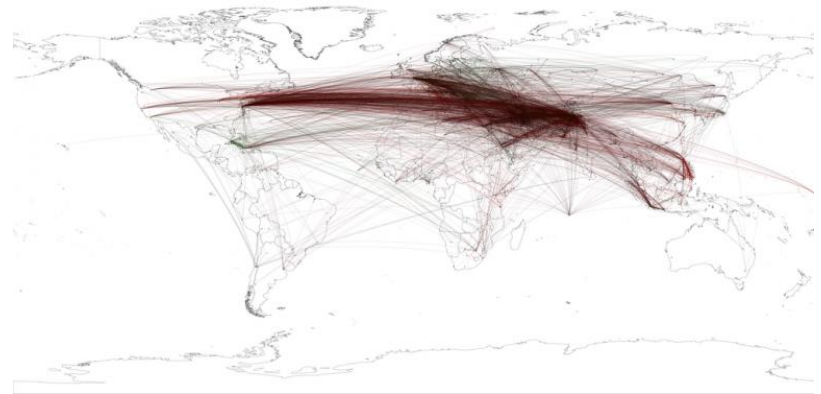
- Boeing 787 design
- Personalized Drugs
- Oil Exploration and Production
- Automotive design
- Art and Entertainment



Gilmanov, LSU, 2009

HPC By Disciplines

- Traditional disciplines
 - Science, such as chemistry, physics and biology
 - Engineering
- Not-so-traditional disciplines
 - Finance
 - Predictive analytics
 - Trading
 - Humanities
 - Culturomics or cultural analytics: study human behavior and cultural trends through quantitative analysis of digitized texts, images and videos



Outline

- What is HPC – background and definitions
- **Engineering & Science Applications**
- Available HPC Resources
- Where to get started - HPC @ LSU and LONI
- Wrap-up


Who Needs a Supercomputer?

- What type of user needs a supercomputer? Recall the types:
 - One with a large database and many calculations over each element of data.
 - Example: Bioinformatics
 - Complex problem using methods that must be iterated over many, many times.
 - Example: Time marching problems-Computational Fluid Dynamics/Molecular dynamics/Discrete Particle Physics

Fluid Dynamics Example


- Simulation of interaction of a large wave with an oil rig using Smoothed Particle Hydrodynamics (DualSphysics, 2012)
- [Ref: http://vimeo.com/53932398](http://vimeo.com/53932398)

- 12 seconds of physical time
- 1,015 million particles
- 64 GPUs Tesla 2090
- 237,342 steps running 91.9 hours.



GeoScience Example

- Collapse of a soil cube (ESyS-Particle develop team) 40 sec debris flow
[https://twiki.esscc.uq.edu.au/index.php/Medium Block, Low Friction](https://twiki.esscc.uq.edu.au/index.php/Medium_Block,_Low_Friction)



Molecular Dynamics Example

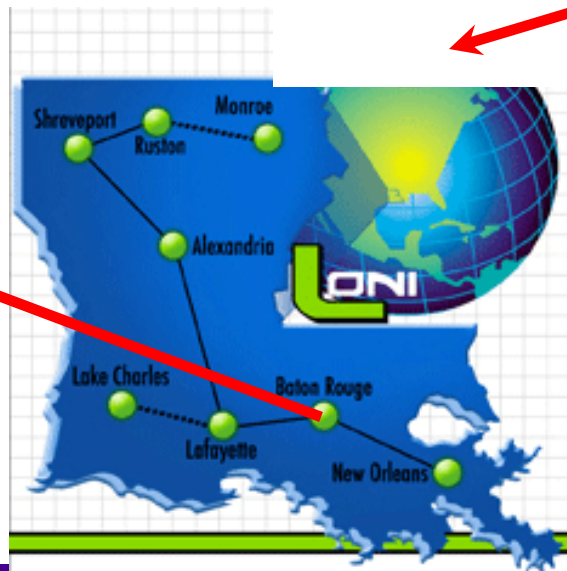

- NAMD is a commonly used MD program.
- MD works by repeated repetition of basic steps:
 - compute the forces acting on each atom.
 - simulate their movement with Newton's Laws for a *short* time, 1 femto-second (1×10^{-15} sec) or so
 - rinse and repeat.
- 1000 steps to reach a pico-second, 1M steps to reach a nano-second, 1B steps to reach a micro-second

Outline

- What is HPC – background and definitions
- Examples from bioinformatics and molecular dynamics
- **Where to get started - HPC resources**
- Wrap-up

Available Computing Resources

- Nation wide-XSEDE
- State wide-LONI
- University wide-HPC@LSU



National Level

- National Level: Extreme Science and Engineering Discovery Environment (xSEDE)
 - 5 year, \$121M project supported by NSF
 - Supports 16 supercomputers and high-end visualization and data analysis resources across the country.

XSEDE

Extreme Science and Engineering
Discovery Environment

State Level-Louisiana Optical Network Initiative (LONI)

- A state-of-the-art fiber optic network that runs throughout Louisiana and connects Louisiana and Mississippi research universities.
- \$40M Optical Network, 10Gb Ethernet over fiber optics.
- \$10M Supercomputers installed at 6 sites.



LONI-Louisiana Optical Network Initiative

- LONI connects supercomputers at Louisiana's universities:
 - Louisiana State University
 - Louisiana Tech University
 - LSU Health Sciences Center in New Orleans
 - LSU Health Sciences Center in Shreveport
 - Southern University
 - Tulane University
 - University of Louisiana at Lafayette
 - University of New Orleans

University Level: HPC@LSU

- University Level: LSU HPC resources available to LSU Faculty and their affiliates.
- LONI and LSU HPC administered and supported by HPC@LSU



Clusters for LSU and LONI Users

	Name	Performance (TFLOPS)	Location	Vendor	Architecture
LONI	Queen Bee	50.7	ISB	Dell	Linux x86_64
	Eric	4.8	LSU	Dell	Linux x86_64
	Oliver	4.8	ULL	Dell	Linux x86_64
	Louie	4.8	Tulane	Dell	Linux x86_64
	Poseidon	4.8	UNO	Dell	Linux x86_64
	Painter	4.8	LaTech	Dell	Linux x86_64
LSU	Philip	3.5	LSU	Dell	Linux x86_64
	SuperMIC	1000	LSU	Dell	Linux x86_64
	SuperMike	212 (CPU+GPU)	LSU	Dell	Linux x86_64
	Pandora	6.8	LSU	IBM	Power7


SuperMike II



Queen Bee



Painter Wiring



When a 2TB Disk Just Won't Do

- This is part of a 240TB storage system for Queen Bee.
- It uses the Lustre Parallel file system to improve I/O throughput.



LONI Cluster Hardware

- Queen Bee
 - 668 nodes: 8 Intel Xeon cores @ 2.33 GHz
 - 8 GB RAM per node
 - 192 TB storage
- Other LONI clusters
 - 128 nodes: 4 Intel Xeons cores @ 2.33 GHz
 - 4 GB RAM per node
 - 9 TB storage

HPC@LSU - SuperMike II

- Ranked 250 in Nov 2012 Top 500 List (393th by June 2013).
- 146 CPU TFlops and 66 double-precision GPU TFlops,
- 440 nodes, dual 8-core Intel Sandy bridge Xeon cores @2.6 GHz
- 382 standard nodes with 32GB RAM (16 cores per node),
- 50 GPU nodes with 64GB RAM and dual NVIDIA Tesla M2090 6GB GPUs,
- 8 big memory nodes with 256GB RAM, capable of aggregation into a single virtual symmetric processing
- (vSMP) node using ScaleMP
- Mellanox Infiniband QDR network of 2:1 over-subscription.

HPC@LSU – Philip & Pandora

- Pandora
 - 8 Power7 nodes, 8 IBM Power7 processors @ 3.33 GHz
 - 128 GB Ram
 - 19 TB storage
- Philip
 - 37 nodes, 8 Intel Xeon cores @ 2.93 GHz
 - 3 nodes w/96 GB Ram; 2 nodes w/48GB Ram; 32 Nodes w/24GB Ram
 - 2 GPU nodes, 12 Intel Xeon core @ 2.66GHz with hyperthreading with 3 Tesla 2070 GPU's each
 - Tesla M2070: 448 CUDA cores @ 1.15GHz and 5.25GB Total Memory

HPC@LSU - SuperMIC

- Expected to be in production by April 1, 2014
- 40 percent resources are reserved for XSEDE
- 360 Compute Nodes
 - Two 2.8GHz 10-Core Ivy Bridge-EP E5-2680 Xeon 64-bit Processors
 - Two Intel Xeon Phi 7120P Coprocessors
 - 64GB DDR3 1866MHz Ram
- 20 Hybrid Compute Nodes
 - Two 2.8GHz 10-Core Ivy Bridge-EP E5-2680 Xeon 64-bit Processors
 - One Intel Xeon Phi 7120P Coprocessors
 - One NVIDIA Tesla K20X 6GB GPU with GPU Direct Support
 - 64GB DDR3 1866MHz Ram

Account Eligibility-*LONI*

- All faculty and research staff at a LONI Member Institution, as well as students pursuing sponsored research activities at these facilities, are eligible for a LONI account.
- Requests for accounts by research associates not affiliated with a LONI Member Institution will be handled on a case by case basis.
- For prospective LONI Users from a non-LONI Member Institution, you are required to have a faculty or research staff in one of LONI Member Institutions as your Collaborator to sponsor you a LONI account.

Account Eligibility-*LSU HPC*

- All faculty and research staff at Louisiana State University, as well as students pursuing sponsored research activities at LSU, are eligible for a LSU HPC account.
- For prospective LSU HPC Users from outside LSU, you are required to have a faculty or research staff at LSU as your Collaborator to sponsor you a LSU HPC account.

LONI & LSU HPC Accounts

- LSU HPC and LONI systems are two distinct computational resources administered by HPC@LSU.
- Having an account on one does not grant the user access to the other.

How do I get a *LONI* Account?

- Visit https://allocations.loni.org/login_request.php
- Enter your INSTITUTIONAL Email Address and captcha code.
- Check your email and click on the link provided (link is active for 24hrs only)
- Fill the form provided
- For LONI Contact/Collaborator field enter the name of your research advisor/supervisor who must be a Full Time Faculty member at a LONI member institution.
- Click Submit button
- Your account will be activated once we have verified your credentials.

How do I get a *LSU HPC* Account?

- Visit https://accounts.hpc.lsu.edu/login_request.php
- Enter your INSTITUTIONAL Email Address and captcha code.
- Check your email and click on the link provided (link is active for 24hrs only)
- Fill the form provided
- For HPC Contact/Collaborator field enter the name of your research advisor/supervisor who must be a Full Time Faculty member at LSU
- Click Submit button
- Your account will be activated once we have verified your credentials

Allocation

- An allocation is a block of service unit (SUs) that allow a user to run jobs on a cluster
 - One SU is one cpu-hour
 - Example
 - 40 SUs will be charged for a job that runs 10 hours on 4 cores
- LONI & HPC users: All jobs need to be charged to valid allocation.

Allocation Types

- **Startup: Allocations upto 50K SUs**
 - Can be requested at any time during the year.
 - Reviewed and Approved by the LONI Resource Allocation Committee.
 - Only two active allocations per PI at any time.
 - Expired Allocations are considered active if the end date is in the future.
- **Large: Allocations between 50K - 4M SUs.**
 - Reviewed and Approved by the LONI Resource Allocation Committee every Quarter.
 - Decision will be made on January 1, April 1, July 1 and October 1 of each year
 - Users can have multiple Large Allocations.
 - **LSU HPC:** Each request is limited to 3 million SUs, and a PI may have a total of 5 million SUs active at any given time.
 - **LONI:** Each requests is limited to 4 million SU, and a PI may have a total of 6M SU active at any given time

Who can request an Allocation?

- Only Full Time Faculty member at LONI member institutions can act as Principle Investigators (PI) and request Allocations.
- Rule of Thumb: If you can sponsor user accounts, you can request allocations.
- Everyone else will need to join an existing allocation of a PI, usually your advisor/supervision or course instructor (if your course requires a LONI account).

How to request/join an Allocation

- Login to your LONI Profile at <https://allocations.loni.org>
- Click on "Request Allocation" in the right sidebar.
- Click "New Allocation" to request a New Allocation.
 - Fill out the form provided.
 - All requests require submission of a proposal justifying the use of the resources.
 - Click "Submit Request" button.
- Click "Join Allocation" to join an existing Allocation.
 - Search for PI using his/her email address, full name or LONI username
 - Click "Join Projects" button associated with the PI's information.
 - You will be presented with a list of allocations associated with the PI. Click "Join" for the allocation you wish to join.
 - Your PI will receive an email requesting him to confirm adding you to the allocation.
 - Please do not contact the helpdesk to do this.

Account Management

- LONI account
 - <https://allocations.loni.org>
- LSU HPC account
 - <https://accounts.hpc.lsu.edu>
- The default Login shell is bash
 - Supported Shells: bash, tcsh, ksh, csh & sh
 - Change Login Shell at the profile page

How do I reset my password?

- LONI: Visit https://allocations.loni.org/user_reset.php
- LSU HPC: Visit https://accounts.hpc.lsu.edu/user_reset.php
- Enter the email address attached to your account and captcha code
- You will receive an email with link to reset your password, link must be used within 24 hours.
- Once you have entered your password, one of the HPC Admins need to approve the password reset.
- The Password approval can take anything from 10 mins to a few hours depending on the schedule of the Admins and also time of day
- You will receive a confirmation email stating that your password reset has been approved.

Password Security

- Passwords should be changed as soon as your account is activated for added security.
- Password must be at least 12 and at most 32 characters long, must contain three of the four classes of characters:
 - lowercase letters,
 - uppercase letters,
 - digits, and
 - other special characters (punctuation, spaces, etc.).
- Do not use a word or phrase from a dictionary,
- Do not use a word that can be obviously tied to the user which are less likely to be compromised.
- Changing the password on a regular basis also helps to maintain security.
 - <http://www.thegeekstuff.com/2008/06/the-ultimate-guide-for-creating-strong-passwords/>
 - http://en.wikipedia.org/wiki/Password_policy

HPC@LSU User Services

- Hardware resources
 - Currently manages 10 clusters
- Software stack
 - Communication software
 - Programming support: compilers and libraries
 - Application software
- User services

Programming Support

- Compilers
 - Intel Fortran and C/C++
 - GNU compiler suite
 - Portland group Fortran and C/C++ (PGI)
 - CUDA
- Scripting languages:
 - Perl, Python, Bash, etc.
- Scientific and utility libraries
 - FFTW, HDF5, NetCDF, PetSc, Intel MKL
- Debugging and profiling tools:
 - Totalview, DDT, TAU

Communication Software

- Shared Memory Programming – OpenMP
 - Good for programs that exhibit *data parallelism*.
 - Manage by the compiler via special programming statements.
- Message Passing Programming – MPI
 - Good for programs that exhibit *task parallelism*.
 - Managed by the programmer with library function calls.
- Hybrid Programming – OpenMP + MPI
- GPU-CUDA

Scientific Libraries

- Some things are common in scientific codes
 - Let the experts develop best-of-class common methods for things like
 - Matrix operation
 - Fast Fourier Transform
- Many scientific libraries are available
 - Linear algebra: BLAS, ATLAS
 - Linear solvers: Scalapack, SuperLU, HYPRE, Intel MKL
 - Fast Fourier transform: FFTW, Intel MKL
 - Boost
- Do not reinvent the wheel!

I/O

- Read and writing data is another common problem
 - ASCII portable, but very slow
 - Binary is fast, but not portable across machine architectures (byte-order problems, floating point format)
 - And what about metadata
- Common HPC I/O libraries
 - HDF - Hierarchical Data Format
 - NetCDF - Network Common Data Format
 - Manage format conversions between machines, can be annotated with metadata
 - Used by many applications, such as third-party visualization applications

Why Worry About Programming

- Getting your hands dirty:
 - Roll your own code.
 - Install a source code application release.
 - Modify an existing code.
 - Understand what the code does *for* and *to* you.
- ELSE, use an existing or installed packages and hope it satisfies all your research needs.
- Try to “***Stand on the shoulders of giants***”

Application software

- Quantum chemistry
 - Gaussian, GAMESS, CPMD, NWChem
- Molecular dynamics
 - NAMD, Amber, LAMMPS, Gromacs
- Engineering
 - Fluent (LSU only), OpenFOAM
- Mathematics and statistics
 - Matlab (LSU only), Mathematica (LSU only), Octave, R, PETSC
- Visualization
 - GaussView, VisIt, VMD, GNUPLOT

User Services

- Consulting services
 - Types of help: Usage Problems, Program Optimization, Installation of Custom Software, Program development advice
- User Guides:
 - HPC: <http://www.hpc.lsu.edu/docs/guides.php#hpc>
 - LONI: <http://www.hpc.lsu.edu/docs/guides.php#loni>
- Documentation: <http://www.hpc.lsu.edu/docs/index.php>
- On-line courses: <http://docs.loni.org/moodle>
- Contact us:
 - Email ticket system: sys-help@loni.org
 - Help Desk: 225-578-0900
 - Instant messaging: lsuhpchelp (AIM, GoogleTalk, Yahoo Messenger)

Educational Activities

- **Weekly Trainings**
 - Introductory: User Environment, Linux
 - Programming: C/C++, Fortran, Shell Scripting, Perl, Python, MPI, OpenMP
 - Software Development: Debugging, Profiling, Make, Subversion
 - Software Applications: Molecular Dynamics, Computational
 - Chemistry & Biology, Octave, MatLab
- **Workshops**
 - Programming: Fortran, C/C++
 - Parallel Programming: MPI, OpenMP, GPU
 - Support Workshops organized through other Departments
- **User Symposium: Researchers from various LONI institutions presented their research via invited talks and poster sessions**
 - <http://www.hpc.lsu.edu/users/symposium/index.php>

Outline

- What is HPC – background and definitions
- Examples from bioinformatics and molecular dynamics
- Where to get started - HPC @ LSU and LONI
- **Wrap-up**

US Council On Competitiveness

- While it is easy to view only the research side of computational science, the field of HPC has much broader modern day impacts.
- The US Council on Competitiveness has a long history of extolling the benefits of HPC use in industry for the competitive edge it provides.
- HPC is becoming an industrial tool, not just a science/research tool

Why HPC is Worth the Effort

- What HPC gives you that you won't get elsewhere is the ability to do larger, better and more exciting science
- Only hope of staying ahead of the data deluge problem we face with modern instrumentation
- Develop techniques for faster, cheaper design
- Prepares you for the day when 1 PetaFLOPS is available on your desktop

The Future is Now

Remember:

Whatever happens in supercomputing today will be on your desktop in 10 – 15 years.

So, if you have experience with HPC, you'll be ahead of the curve when things get to the desktop.

The End

- Thank you for your attention!
 - Any Questions?
- **Next Week (Important!):**
Introduction to Linux