High-Performance Numerical Computing Summer Workshop

Shahriar Afkhami
Dept. Mathematical Sciences

Goals

- Exposure to tools and techniques
- Improve productivity
- Develop new ideas

We have a big/diverse group of people, from undergrad to postdoc, and from engineering to math and physics

- I cannot cover everything deeply: there will be a wide range of topics.
- There might be some basic and preliminary topics; be patient!
- There will be some involved and new topics, I promise!
- There will be hands-on sessions, mainly on Thursdays - ask questions as much as you can. Meet your neighbors, they may have some great ideas!

Getting started with HPC

- Intro to NJIT HPC resources:
  `https://wiki.hpc.arcs.njit.edu/index.php/HPC_and_BD`
  `https://ist.njit.edu/high-performance-computing-machine-specifications/`
- Intro to parallel computing
- Logging in and using the computer clusters

Requirements

- Hands-on exercises will require computers; participants are expected to bring a laptop computer and power supply.
- You will need an account on stheno and/or kong clusters.
  - for accessing kong, send email to arcs@njit.edu
  - for accessing stheno (only open to affiliates of Dept. Mathematical Sciences), faculty can directly email arcs@njit.edu; students, postdocs, and other researchers must have their faculty advisor requesting access on their behalf.
- You will need an ssh client with XWindows.
  - for Mac users, see: `https://ist.njit.edu/using-SSH-from-Mac-OS-X/`
  - for Windows: use MobaXterm `http://ist.njit.edu/software/display.php?id=632`
  - ssh from outside the NJIT network will require a VPN client; see: `https://ist.njit.edu/vpn/`

AFS

Andrew File System (AFS) is the primary academic computing environment at NJIT. It is a distributed file system comprised of multiple file and database. Avery wide spectrum of both open source and commercial applications, compilers, libraries,and utilities is available in AFS:
`https://ist.njit.edu/afs/`

Reset AFS password:
`https://ist.njit.edu/password-resets/`

What is on AFS:

- Storage Space: Each student receives storage space on AFS, but only 1 GB!!!
- Web Pages: Personal web pages are stored on AFS as `http://web.njit.edu/~yourUCID` public_html directory in your home directory.
- Software: A wide spectrum of open source and commercial applications, compilers, libraries, and utilities is available in AFS.

40 Years of Microprocessor Trend Data

40 Years of Microprocessor Trend Data

One of the most popular plots when it comes to technologic advancements in microprocessors in general and Moore's Law in particular is a plot entitled *35 Years of Microprocessor Trend Data* based on data by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten. Later, trend lines with some (speculative) extrapolation were added by C. Moore. One can find the plot with and without trend lines at various locations in the web (and further down). However, the plot suffers from one the sands of time: Data is only plotted up until the year 2010, missing out the last five years.

Update Feb. 15, 2018: Check out the newer 42 Years of Microprocessor Trend Data.

In order to add more recent data and in order to judge whether the extrapolations from 2010 were appropriate, I extracted the original data using g3data and added data from AMD Opteron processors, Intel Xeon processors, Power7+ and Power8 processors as well as Xeon Phi to the plot. Details on the new data can be found in a zip archive with all raw data. The resulting plot is as follows:

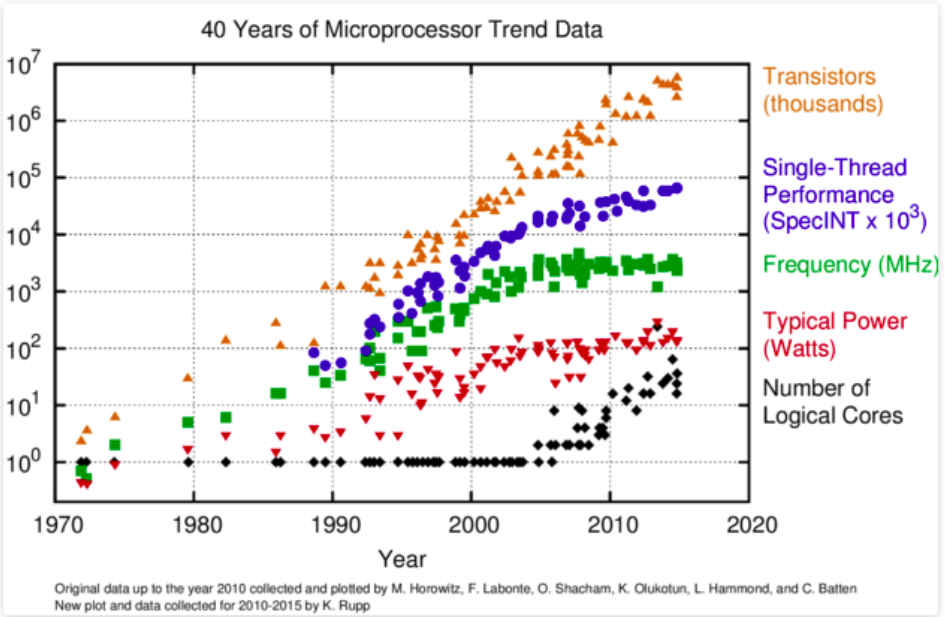

Some terminologies

- Cluster: connected collection of computers known as computing nodes
- Server: single-unconnected node
- Machine: entire clusters and stand alone servers

Some terminologies: Clusters

- Servers are stand alone computers: Each have their own OS, CPU, RAM, HD, ...
- Stored in racks in GITC
- Connected together via interconnects
  kong via ethernet $<$ stheno via infiniband

Some terminologies: SMP

- SMP: Symmetric Multi-Processing
- Two or more identical processors are connected to a single shared memory
- Have full access to all input and output devices, and are controlled by a single operating system instance that treats all processors equally

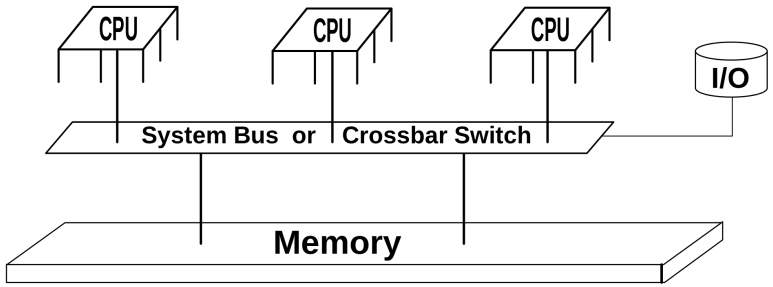

Basic Cluster Layout

- Login node(s): edit, debug, compile, and interact with scheduler
- Scheduler: a mechanism for distributing jobs across the nodes
  - You tell the scheduler the number of cpus, ram, number of hours, ...
  - Scheduler then assigns hardware
  - SGE and SLURM are current schedulers
    `https://wiki.hpc.arcs.njit.edu/index.php/SchedulerIntro`
  - !!!Login nodes ARE ABSOLUTELY NOT for long running jobs!!!

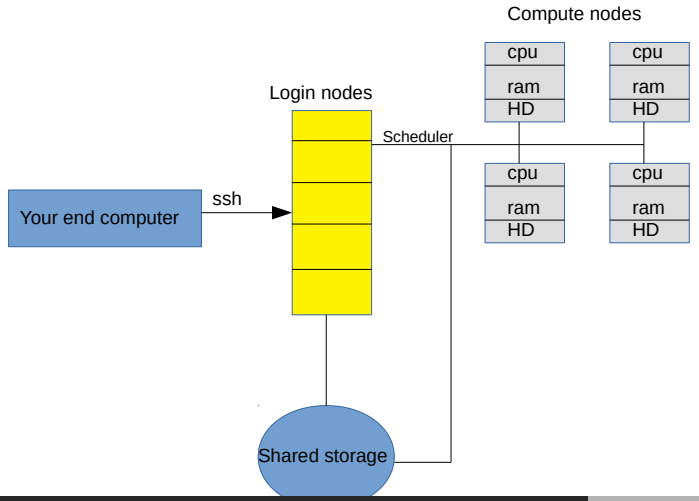

Modern Clusters

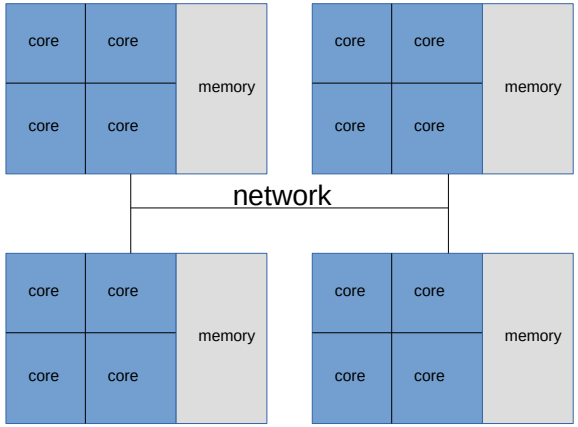

- Serial codes: only use one core
- Multithreading: only uses one node - OpenMP
- Message Passing Interface (MPI): can use all the discrete nodes

Emerging Architectures: Clusters with accelerators

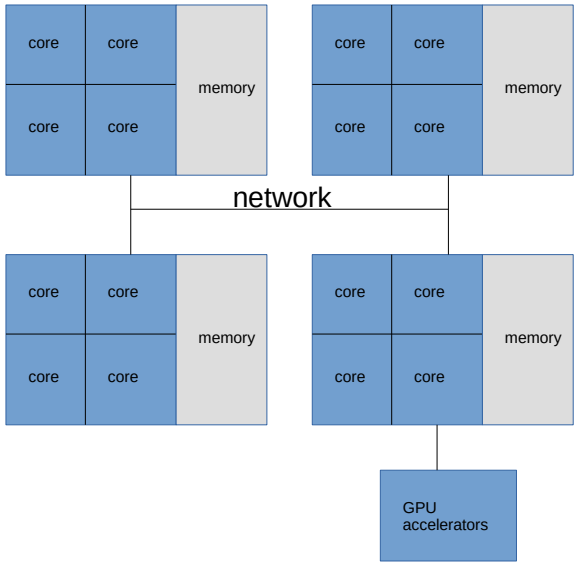

- Graphic Processing Unit (GPU) programmable using Cuda, OpenCL

Emerging Architectures: Cloud Computing

- Regular servers sitting somewhere
- Users can log in remotely to virtual machines of containers
- Usage: almost everything, from hosting websites, to machine learning and HPC
- Major providers are AWS, Google, MS, HP, IBM, ...
  - Save on buying machines and pay as you go
  - Provides access to a wide range of hardware and software