Phys 7411
Computational Physics:
Computing for Petascale Systems

Karen Tomko - ktomko@osc.edu
and
Juana Moreno - moreno@phys.lsu.edu

Spring 2009

Today's topics

- Goals of this course
- Syllabus / grading
- Online Resources
- OSC Glenn Opteron cluster
- Introduction/Motivation for HPC

Goals

- Primary:
  - Provide an introduction to the computing systems, programming approaches and some common numerical methods used for high performance computing
- Secondary:
  - Provide hands-on parallel programming experience

Course Basics

- Course website
  - http://www.osc.edu/~ktomko/HPC_course/CoursePage.html
- Google Group
  - Email: phys7411@googlegroups.com
  - Website: http://groups.google.com/group/phys7411

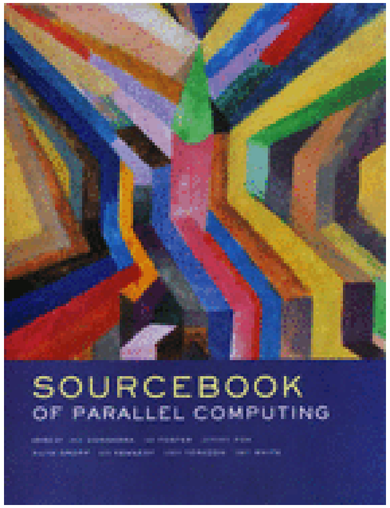

- Text book
  - Jack Dongarra, Ian Foster, Geoffrey Fox, William Gropp, Ken Kennedy, Linda Torczon, Andy White, Soucebook of Parallel Computing, Morgan Kaufmann/Elsevier, 2003.

Workload/Grading

- Programming Assignments (3)
  - MPI programming, OpenMP, performance measurement
  - 20% of final grade each
- Final project
  - Small research project in groups of 2
  - 40% of final grade

OSC Glenn Opteron Cluster

- Glenn links
  - Online guide to using Glenn: http://www.osc.edu/supercomputing/computing/opt/index.shtml
  - Description of the machine http://www.osc.edu/supercomputing/hardware/#opt
  - A set of slides from a workshop on using Glenn: http://www.osc.edu/supercomputing/training/opt/
- Accounts
  - Everyone should have an account
  - Send me an email if you do not (ktomko@osc.edu)
- FAQ for OSC supercomputer users
  - http://www.osc.edu/supercomputing/faq.shtml

Top 500 List

Rank	Site	Computer
1	DOE/NNSA/LANL United States	Roadrunner - BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz , Voltaire Infiniband IBM
2	Oak Ridge National Laboratory United States	Jaguar - Cray XT5 QC 2.3 GHz Cray Inc.
3	NASA/Ames Research Center/NAS United States	Pleiades - SGI Altix ICE 8200EX, Xeon QC 3.0/2.66 GHz SGI
4	DOE/NNSA/LLNL United States	BlueGene/L - eServer Blue Gene Solution IBM
5	Argonne National Laboratory United States	Blue Gene/P Solution IBM

Top 500 List (continued)

Rank	Site	Computer
6	Texas Advanced Computing Center/Univ. of Texas United States	Ranger - SunBlade x6420, Opteron QC 2.3 Ghz, Infiniband Sun Microsystems
7	NERSC/LBNL United States	Franklin - Cray XT4 QuadCore 2.3 GHz Cray Inc.
8	Oak Ridge National Laboratory United States	Jaguar - Cray XT4 QuadCore 2.1 GHz Cray Inc.
9	NNSA/Sandia National Laboratories United States	Red Storm - Sandia/ Cray Red Storm, XT3/4, 2.4/2.2 GHz dual/quad core Cray Inc.
10	Shanghai Supercomputer Center China	Dawning 5000A - Dawning 5000A, QC Opteron 1.9 Ghz, Infiniband, Windows HPC 2008 Dawning

Top 500

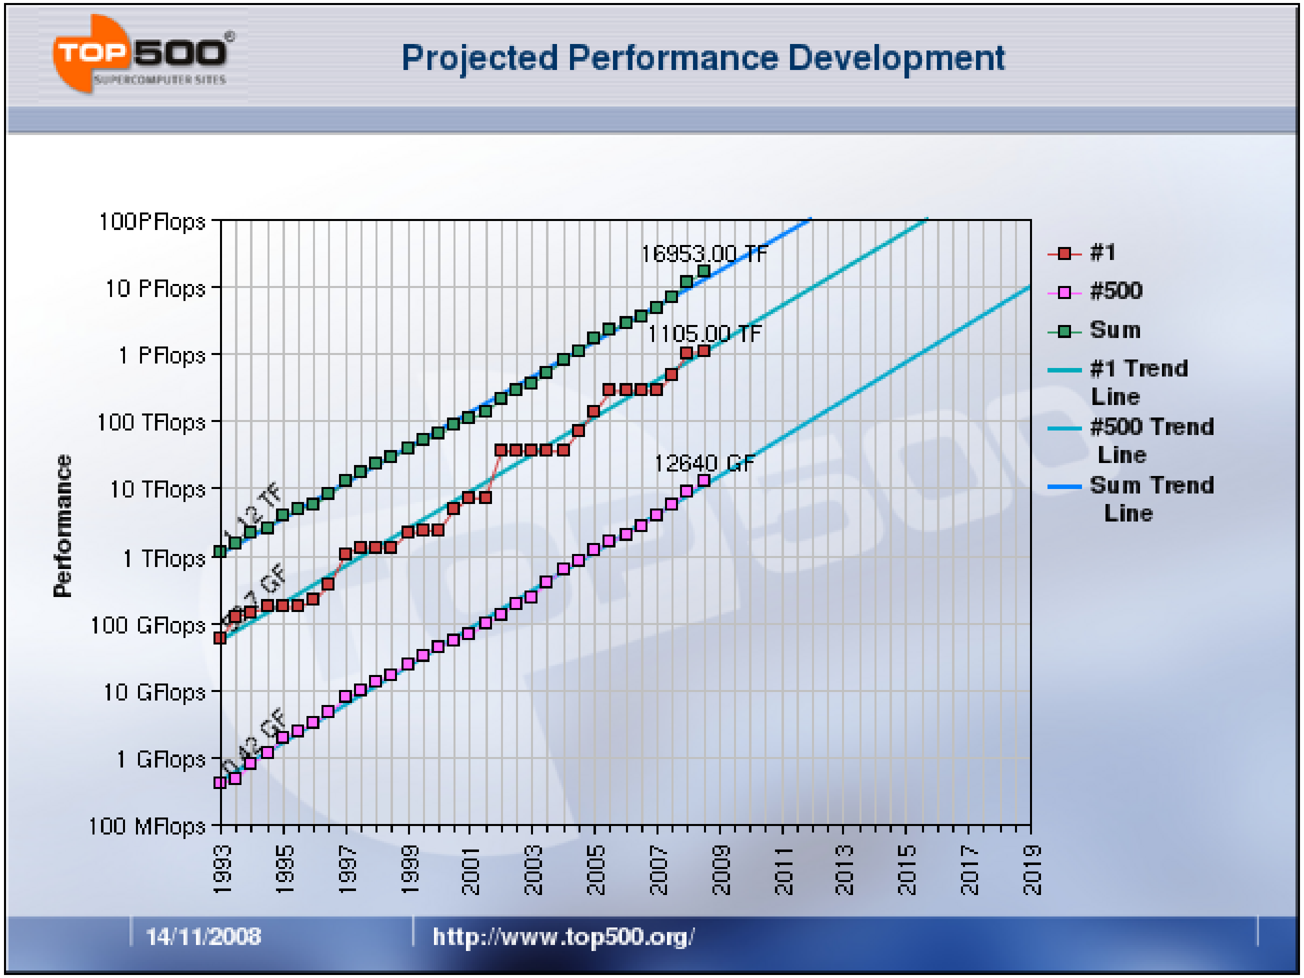

More nice charts -
http://www.top500.org/static/lists/2008/11/TOP500_200811_BOF-2.pdf

Why HPC?

- Who needs a roomful of computers when you can have more than a GFLOPS capacity in your hands?

OSC IBM Opteron Cluster, Glenn
http://www.osc.edu/press/logos/index.shtml#glenn

Who are the users?

- People who study:
  - Materials / Superconductivity
  - Fluid Flow
  - Weather/Climate
  - Structural Deformation
  - Genetics / Protein interactions
  - Seismic Modeling
  - and others…

Why are the problems so large?

- Many Length Scales (both time and space)
  - If you want to observe the interactions between very small local phenomenon and larger more global phenomenon
- Multi-Physics Simulation
  - Many interacting phenomenon must be modeled for a realistic treatment of a system
- Physical problems (3-D in space, plus time)
  - If you increase the level of spatial resolution by factor of 10, problem geometry increases by 10^3
  - An increase in spatial resolution usually must have a corresponding increase in temporal resolution

How can you solve these problems?

- Take advantage of parallelism
  - Large problems generally have many operations which can be performed concurrently
- Parallelism can be exploited at many levels by the Computer hardware
  - Within the CPU core, multiple functional units, pipelining
  - Within the Chip, many cores
  - On a node, multiple chips
  - In a system, many nodes

So why don’t I get 22TFLOPS when I run on Glenn

- Parallelism has overheads
  - At the core and chip level the cost is complexity
  - Keeping pipelines full of data is difficult
  - Most applications get only a fraction of peak performance (10%-20%)
  - Between nodes, the communication infrastructure is typically much slower than the CPU

Other constraining issues

- Heat dissipation and power
  - See http://www.green500.org/index.php
- Reliability
- I/O

Other options besides Linux clusters

- FPGA/Reconfigurable computing accelerators
- GPUs (now 64-bit)
- Cell or Array processors (many simpler cores)