

Advanced Topics in Numerical Analysis: High Performance Computing

MATH-GA 2012.001 & CSCI-GA 2945.001

Georg (Stadler)
Courant Institute, NYU
stadler@cims.nyu.edu

Spring 2017, Thursday, 5:10–7:00PM, WWH #512

Feb. 2, 2017

Outline

Organization issues

A tour through HPC

Demo time

Tool/command of the week

Organization issues

- ▶ **Time and location:** Thursday 5:10–7:00PM, WWH 512
- ▶ Public course **webpage:**
<http://cims.nyu.edu/~stadler/hpc17/>

Organization issues

- ▶ **Time and location:** Thursday 5:10–7:00PM, WWH 512
- ▶ Public course **webpage:**
<http://cims.nyu.edu/~stadler/hpc17/>
- ▶ We will mainly use Piazza for organization/communication in this class. Let me know if you want to be added.
- ▶ TA/grader: Bill Bao (yxb201@nyu.edu)
- ▶ **Email:** If you email me (stadler@cims.nyu.edu) about anything related to this course, please put [hpc17] in the email's subject line.

Organization issues

- ▶ **Time and location:** Thursday 5:10–7:00PM, WWH 512
- ▶ Public course **webpage:**
<http://cims.nyu.edu/~stadler/hpc17/>
- ▶ We will mainly use Piazza for organization/communication in this class. Let me know if you want to be added.
- ▶ TA/grader: Bill Bao (yxb201@nyu.edu)
- ▶ **Email:** If you email me (stadler@cims.nyu.edu) about anything related to this course, please put [hpc17] in the email's subject line.
- ▶ **Office hours:** Wednesday, 10AM-noon, stop by or make an appointment (please email). My office number is #1111.
- ▶ We have different backgrounds (Maths/CS/Economy/(?)). Please ask during class! Feedback and comments of any kind are welcome.

Organization issues

Prerequisites:

- ▶ Some basic experience in serial programming in either C or FORTRAN (this is not primarily a programming course!)
- ▶ Basic command line experience (shell, ssh, . . .)
- ▶ Basic knowledge/interest in numerical algorithms.

Organization issues

Prerequisites:

- ▶ Some basic experience in serial programming in either C or FORTRAN (this is not primarily a programming course!)
- ▶ Basic command line experience (shell, ssh, ...)
- ▶ Basic knowledge/interest in numerical algorithms.



Class topics (intended):

- ▶ **Architecture and algorithms:** hardware basics, memory hierarchy, programming models, networks
- ▶ **Programming:** MPI, OpenMP, OpenCL, performance optimization
- ▶ **Tools:** Make, debugging, valgrind, git, paraview, job schedulers, ...

Organization issues

Recommended textbooks/literature:

- ▶ T. Rauber and G. Runger: *Parallel Programming for Multicore and Cluster Systems*, Springer, 2nd edition 2013. Available online on NYU Campus.
- ▶ A. Grama, A. Gupta, G. Karypis and V. Kumar: *Introduction to Parallel Computing*, Pearson, 2003.
- ▶ To refresh C programming language: Z. Shaw: *Learn C the hard way*, 2016.
- ▶ Command line crash course: Youtube, *From Bash to Z Shell* (see links on Piazza).
- ▶ (+ some online references for later in the course)



Organization issues

Required work:

- ▶ Practical homework assignments every 2-3 weeks (50% of grade)—first short assignment posted, due next week.
- ▶ Final project (50% of grade):
 - ▶ by yourself or in teams of 2 or 3
 - ▶ 10min project presentations at the end of the class, and handing in a short report
 - ▶ project can be related to your research project, but it needs to bring in some of the topics of this course!
 - ▶ please start thinking about your projects (would like to finalize project topics by mid March)
 - ▶ pitch your project to me during the next weeks

Even if you are only auditing this class, I encourage you to “get your hands dirty” (work on homework assignments, try out code, work on a project)!

Organization issues

Get the most out of this class!

- ▶ I'll “force” you to do some work (homework, final project)
- ▶ Will additionally try to open doors that allow you to explore and learn more things—it's up to you to get the most out of the class

Organization issues

Get the most out of this class!

- ▶ I'll “force” you to do some work (homework, final project)
- ▶ Will additionally try to open doors that allow you to explore and learn more things—it's up to you to get the most out of the class
- ▶ Ideally, your final project helps for your research, industry or academic applications
- ▶ Interactive hands-on class. Please contribute (in class, on Piazza, etc)
- ▶ **Stop me and ask** if I do something that does not make sense to you!

Outline

Organization issues

A tour through HPC

Demo time

Tool/command of the week

Units and performance measures

- ▶ **flop**: floating point operation, usually double precision (add or multiply; fused multiply-add)
- ▶ **flop/s**: floating point operation per second
- ▶ **bytes**: size of data (double floating point is 8 bytes), each byte has 8 bits


Things are getting large...

Mega:	Mflop/s= 10^6 flop/sec	Mbyte= $2^{20} \approx 10^6$ bytes
Giga:	Gflop/s= 10^9 flop/sec	Gbyte= $2^{30} \approx 10^9$ bytes
Tera:	Tflop/s= 10^{12} flop/sec	Tbyte= $2^{40} \approx 10^{12}$ bytes
Peta:	Pflop/s= 10^{15} flop/sec	Pbyte= $2^{50} \approx 10^{15}$ bytes
Exa:	Eflop/s= 10^{18} flop/sec	Ebyte= $2^{60} \approx 10^{18}$ bytes
Zetta:	...	
Yotta:	...	

Currently fastest machine: Tianhe-2: ~ 125 Pflop/s, 10.6 million processor cores, 1.3Pbyte of memory (we'll talk about the fastest machines later)


Limits of transistor scaling

- ▶ exponential growth of cost for tools for chip manufacturing (Moore's 2nd law)




Limits of transistor scaling

- ▶ exponential growth of cost for tools for chip manufacturing (Moore's 2nd law)
- ▶ reliability of production (high yield; number of Intel's Xeon Phi, "60+" cores)




Limits of transistor scaling

- ▶ exponential growth of cost for tools for chip manufacturing (Moore's 2nd law)
- ▶ reliability of production (high yield; number of Intel's Xeon Phi, "60+" cores)
- ▶ power density is cubic in frequency (leakage due to quantum mechanical effects becomes a problem)



Limits of transistor scaling


- ▶ exponential growth of cost for tools for chip manufacturing (Moore's 2nd law)
- ▶ reliability of production (high yield; number of Intel's Xeon Phi, "60+" cores)
- ▶ power density is cubic in frequency (leakage due to quantum mechanical effects becomes a problem)



Thus: more transistors, but not faster serial processors

Moore's law today

- ▶ Frequency/clock speed stopped growing in ~ 2004
- ▶ Number of cores per CPU
- ▶ Moore's law still holds
- ▶ Energy use \sim bounded



Source: CS Department, UC Berkeley.

Parallel computing \subset high-performance computing

performance = single core & multicore performance

- ▶ All major vendors produce **multicore** chips.
- ▶ How well can applications and algorithms **exploit parallelism**?
- ▶ Do we need to **think differently** about algorithms?
- ▶ How do we divide problems into smaller chunks?

Parallel computing \subset high-performance computing

performance = single core & multicore performance

- ▶ All major vendors produce **multicore** chips.
- ▶ How well can applications and algorithms **exploit parallelism**?
- ▶ Do we need to **think differently** about algorithms?
- ▶ How do we divide problems into smaller chunks?

How about automatic parallelization?

- ▶ compilers help with other levels of parallelism (we will talk about those)
- ▶ automatic parallelization across different CPU cores has mostly been a failure (OK, some tools/libraries are in development)
- ▶ need to **think parallel**; parallelizing a sequential code can be difficult or impossible

Parallel computing \subset high-performance computing

Beyond floating point operations per second

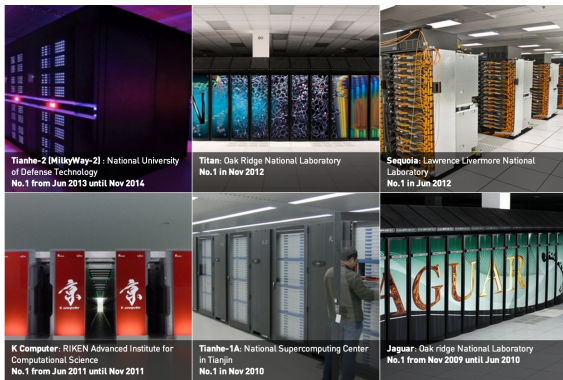
Memory density and time to access memory:

- ▶ The **memory density** only doubles every ~ 3 years (Moore's law at a slower rate)
- ▶ **Memory-bound algorithms**: For many mathematical algorithms, loading/writing data from/to the memory is the bottleneck
- ▶ **Memory hierarchies** have been established (L1/L2/L3-cache), to allow faster read and write of memory
- ▶ Will be the topic of one of the next classes

The Top 500 List: <http://www.top500.org/>

Listing of most powerful supercomputers in the world

- ▶ updated twice a year at the European and US Supercomputing Conferences
- ▶ Reports theoretical flop/s (*RPEAK*), and “sustained” flop/s (*RMAX*) on a benchmark problem (dense matrix-vector multiplication)

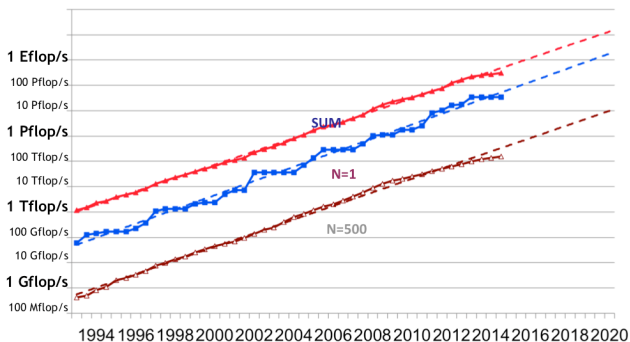


<https://youtu.be/HBpkSIJ9KwU?t=5m18s>

The Top 500 List: <http://www.top500.org/>

Development over time (observe Moore's law)

- ▶ #1 machine on Top500
- ▶ #500 machine on Top 500
- ▶ Sum over all machines on Top 500 list



The US Department of Energy (DOE) plans a 1 Eflop/s machine (with ≤ 20 MW power consumption) by 2020 (?). . .

The Top Green 500 List: <http://www.green500.org/>

Listing of most energy-efficient computers in the world

- ▶ Reports floating point operations per Watt of energy used
- ▶ Energy usage is a serious challenge in processor design (also for portable devices)

The Top Green 500 List: <http://www.green500.org/>

Listing of most energy-efficient computers in the world

- ▶ Reports floating point operations per Watt of energy used
- ▶ Energy usage is a serious challenge in processor design (also for portable devices)
- ▶ rough estimate of per person electricity use in the US: 1kW
 - ▶ costs ~ 10 cents per hour (~ 20 in Europe)
- ▶ Thus, running a 20MW machine costs about: $20 \times 1000 \times 10$ cents = 2000\$ per hour $\sim 2M$ \$ per year.
- ▶ Roughly same amount of energy (and cost) is necessary for cooling (often, water cooling)


Do we really need larger and faster?

- ▶ Increased computational power allows simulations **beyond what experiments or theory** (paper & pencil) can do.
- ▶ Allows to go beyond “simulations”, for instance by solving
 - ▶ **inverse problems**: infer unknown phenomena that cannot be observed/measured from observations (e.g., earth structure; weather prediction)
 - ▶ **control/design/decision problems**: influence/control systems, simulate different scenarios
- ▶ Exponential growth in computing power brings along enhances in sensors, i.e., we have much more data that needs to be analyzed to make it useful (“finding the signal in the noise”) ⇒ **big data** requires big (and smart) computing

Do we really need larger and faster?

<http://blog.tamest.org/computational-science-the-third-pillar-of-science/>


Traditional scientific method:



Do we really need larger and faster?

<http://blog.tamest.org/computational-science-the-third-pillar-of-science/>

Traditional scientific method:




Limitations:

- ▶ difficult (wind tunnels)
- ▶ expensive (build experimental passenger jets; crash tests)
- ▶ slow (climate change)
- ▶ dangerous (weapons; drug design)

Do we really need larger and faster?

<http://blog.tamest.org/computational-science-the-third-pillar-of-science/>


Traditional scientific method:



Limitations:

- ▶ difficult (wind tunnels)
- ▶ expensive (build experimental passenger jets; crash tests)
- ▶ slow (climate change)
- ▶ dangerous (weapons; drug design)

Simulation has become the **third pillar of Science**:



Particularly challenging computations

- ▶ Science

- ▶ weather prediction, climate modeling
- ▶ genomics, drug design
- ▶ astrophysical simulation (supernova/black holes, . . .)
- ▶ human brain modeling
<https://www.humanbrainproject.eu/>
- ▶ earthquake modeling and global imaging
- ▶ global mantle convection modeling

Particularly challenging computations

▶ Science

- ▶ weather prediction, climate modeling
- ▶ genomics, drug design
- ▶ astrophysical simulation (supernova/black holes, . . .)
- ▶ human brain modeling
<https://www.humanbrainproject.eu/>
- ▶ earthquake modeling and global imaging
- ▶ global mantle convection modeling

▶ Engineering

- ▶ earthquake engineering
- ▶ semiconductor design
- ▶ computational fluid dynamics for airplane design
- ▶ crash test simulations
- ▶ weapons testing (US nuclear weapon tests stopped in 1992)

Particularly challenging computations

▶ Science

- ▶ weather prediction, climate modeling
- ▶ genomics, drug design
- ▶ astrophysical simulation (supernova/black holes, . . .)
- ▶ human brain modeling
<https://www.humanbrainproject.eu/>
- ▶ earthquake modeling and global imaging
- ▶ global mantle convection modeling

▶ Engineering


- ▶ earthquake engineering
- ▶ semiconductor design
- ▶ computational fluid dynamics for airplane design
- ▶ crash test simulations
- ▶ weapons testing (US nuclear weapon tests stopped in 1992)

▶ Business

- ▶ financial modeling; speed trading
- ▶ transaction processing

Example 1: Seismic wave propagation

<http://vimeo.com/16807465>



Simulation of waves through the earth at a frequency of f Hertz

$$\sim 1600^3 p^3 f^3 \text{ gridpoints,}$$

where p is the number of points per wavelength (e.g., 6), and 1600 is computed from the earth radius (about 6700km) and the average wave speed 5-10km/sec).


f	#gridpints	per timestep
1/8Hz	1.7×10^9	$\sim 400\text{Gflop}$
1/4Hz	1.4×10^{10}	$\sim 3\text{Tflop}$
1/2Hz	1.1×10^{11}	$\sim 24\text{Tflop}$
1Hz	8.7×10^{12}	$\sim 280\text{Tflop}$

Example 2: Weather/Climate modeling

Target: Compute/predict weather or future climate.

Approach: Discretize governing PDEs (atmosphere, ocean, land ice, land models. . .)

Uses: Predict major events (Sandy, Katrina) and evaluate global warming scenarios




Simulated radar reflectivity (Source: NOAA).

Example 2: Weather/Climate modeling

Target: Compute/predict weather or future climate.

Approach: Discretize governing PDEs (atmosphere, ocean, land ice, land models. . .)

Uses: Predict major events (Sandy, Katrina) and evaluate global warming scenarios




Simulated radar reflectivity (Source: NOAA).

Example 2: Weather/Climate modeling

Target: Compute/predict weather or future climate.

Approach: Discretize governing PDEs (atmosphere, ocean, land ice, land models. . .)

Uses: Predict major events (Sandy, Katrina) and evaluate global warming scenarios




Simulated radar reflectivity (Source: NOAA).

Example 2: Weather/Climate modeling

Target: Compute/predict weather or future climate.

Approach: Discretize governing PDEs (atmosphere, ocean, land ice, land models. . .)

Uses: Predict major events (Sandy, Katrina) and evaluate global warming scenarios




Simulated radar reflectivity (Source: NOAA).

Example 2: Weather/Climate modeling

Target: Compute/predict weather or future climate.

Approach: Discretize governing PDEs (atmosphere, ocean, land ice, land models. . .)

Uses: Predict major events (Sandy, Katrina) and evaluate global warming scenarios



Simulated radar reflectivity (Source: NOAA).

Weather prediction modeling (atmosphere only) requirements:

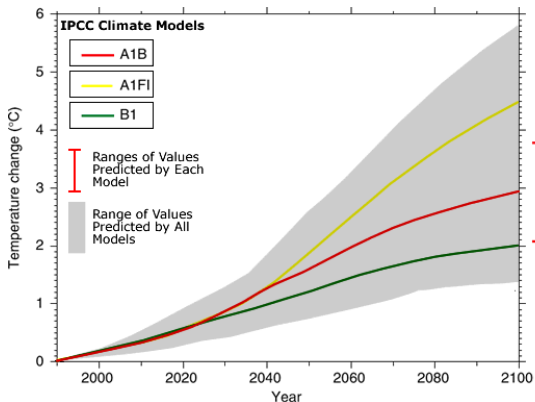
- ▶ For “real time” simulation: 8Gflop/s
- ▶ for weather prediction (7 days in 24h simulation): 56Gflop/s
- ▶ for climate prediction (50 years in 12 days) 288Tflop/s

Current weather models are often inaccurate due to coarse resolution (e.g., clouds not resolved). About last week’s “blizzard”:

[Read Techonomy article.](#)


Example 2: Weather/Climate modeling

Climate predictions based on different CO₂ scenaria (adapted from IPCC, 2001)




Climate models require coupled simulations of climate system (atmosphere+chemistry, land ice, oceans, sea ice) over long time scales (100, 200 years). **Very challenging computations.**

Example 3: Understanding tectonic plate motion




- ▶ Driver for plate tectonics
- ▶ Stress release between plates causes earthquakes and tsunamis
- ▶ Subduction (left image) causes volcanic activity (chemical processes lead to rock melting)

- ▶ Mantle convection is the thermal convection in the Earth's upper ~ 3000 km
- ▶ It controls the thermal and geological evolution of the Earth
- ▶ Solid rock in the mantle moves like viscous incompressible fluid on time scales of millions of years



Example 3: Understanding tectonic plate motion



- ▶ requires very fine mesh to solve nonlinear partial differential equations
- ▶ $> 10^8$ unknowns, strong nonlinearity, iterative solvers

Outline

Organization issues

A tour through HPC

Demo time

Tool/command of the week

Demos, and preview

Can be downloaded from:

<https://github.com/NYU-HPC17/lecture1>

Outline

Organization issues

A tour through HPC

Demo time

Tool/command of the week

The module command

Allows to switch the user environment (programs, compilers etc).
Available on all UNIX-based systems, i.e., on CIMS computers,
compute servers etc.

```
module list
```

```
module avail ...
```

```
module load python/3.4
```

```
module unload texlive-2016
```

```
module whatis gcc-6.1.0
```