Parallel Programming & Cluster Computing

Overview: What the Heck is Supercomputing?

Joshua Alexander, U Oklahoma
Ivan Babic, Earlham College
Michial Green, Contra Costa College
Mobeen Ludin, Earlham College
Tom Murphy, Contra Costa College
Kristin Muterspaw, Earlham College
Henry Neeman, U Oklahoma
Charlie Peck, Earlham College

People


EXIT
Fall Semester 2003
Integration of HPC in Nanotechnology
National Science Foundation

Things

Thanks for your attention!

Questions?

www.oscer.ou.edu

What is Supercomputing?

Supercomputing is the **biggest, fastest computing right this minute**.

Likewise, a ***supercomputer*** is one of the biggest, fastest computers right this minute.

So, the definition of supercomputing is **constantly changing**.

Rule of Thumb: A supercomputer is typically at least 100 times as powerful as a PC.

Jargon: Supercomputing is also known as ***High Performance Computing*** **(HPC)** or ***High End Computing*** **(HEC)** or ***Cyberinfrastructure*** **(CI).**

Fastest Supercomputer vs. Moore

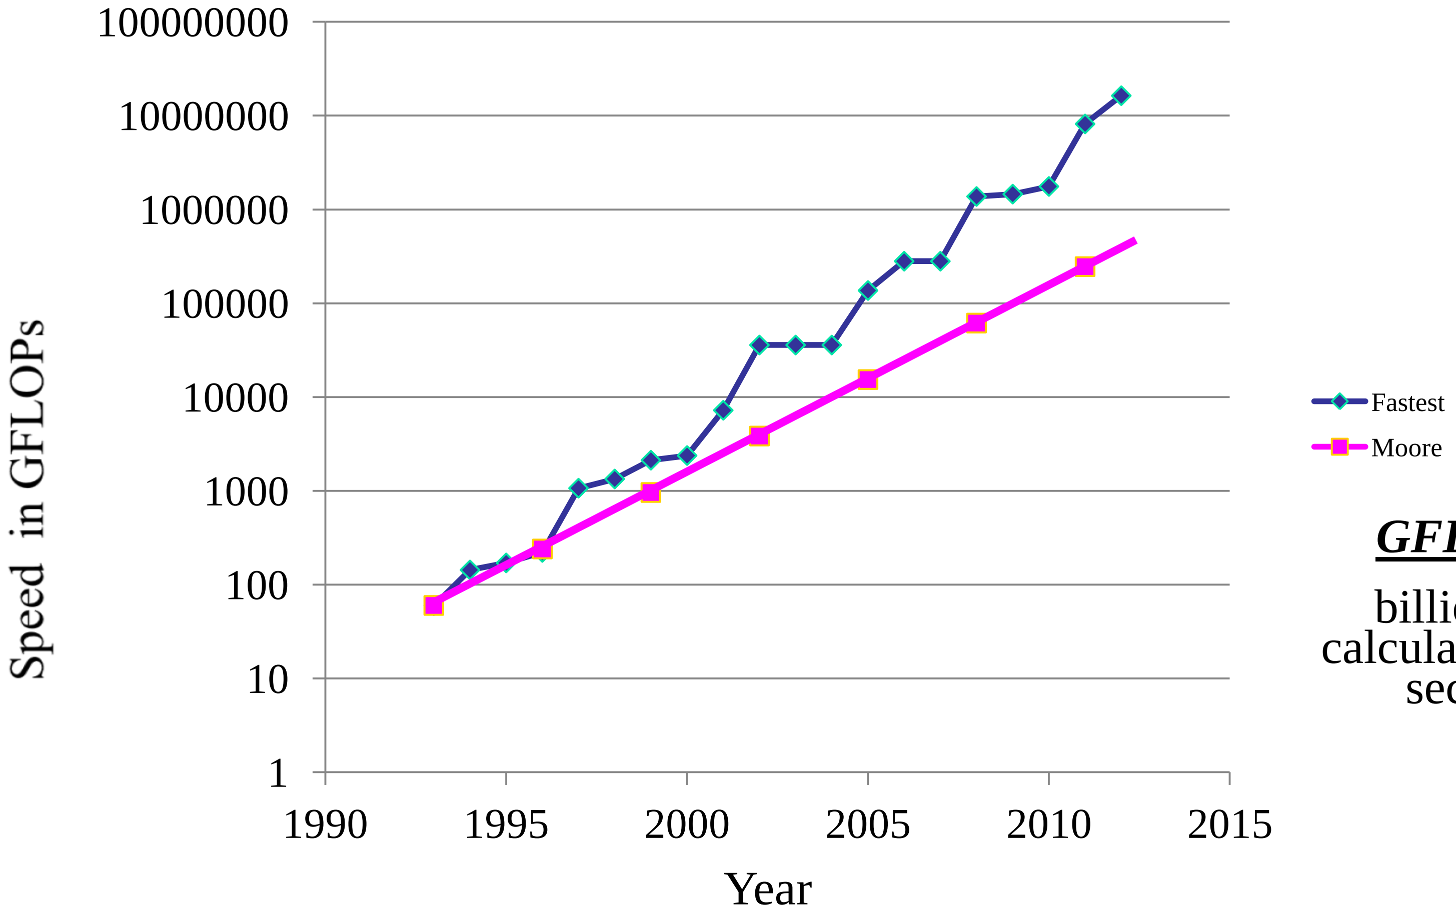

GFLOPs:

billions of calculations per second

What is Supercomputing About?

Size

Speed

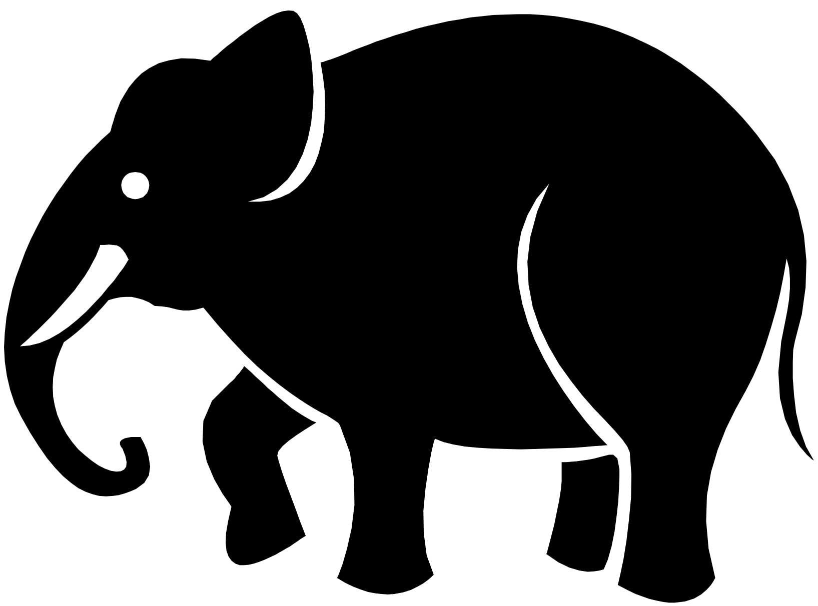

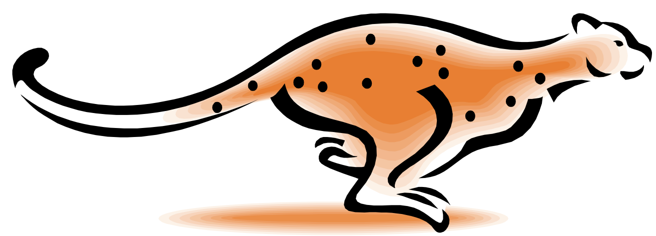

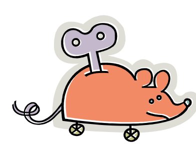

Laptop

What is Supercomputing About?

- **Size**: Many problems that are interesting to scientists and engineers **can't fit on a PC** – usually because they need more than a few GB of RAM, or more than a few 100 GB of disk.

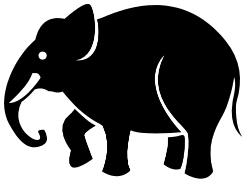

- **Speed**: Many problems that are interesting to scientists and engineers would take a very very long time to run on a PC: months or even years. But a problem that would take **a month on a PC** might take only **an hour on a supercomputer**.

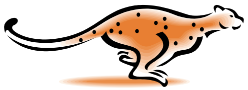

What Is HPC Used For?

- ***Simulation*** of physical phenomena, such as
  - Weather forecasting
  - Galaxy formation
  - Oil reservoir management

[1]

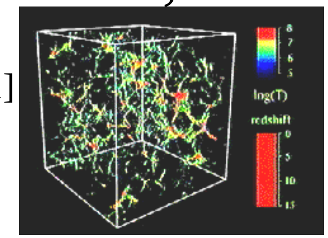

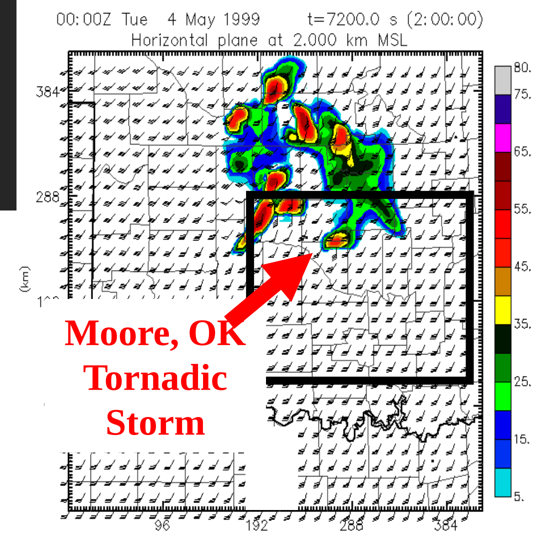

May 3 1999[2]

- ***Data mining***: finding **needles** of information in a **haystack** of data, such as
  - Gene sequencing
  - Signal processing
  - Detecting storms that might produce tornados
- ***Visualization***: turning a vast sea of **data** into **pictures** that a scientist can understand

[3]

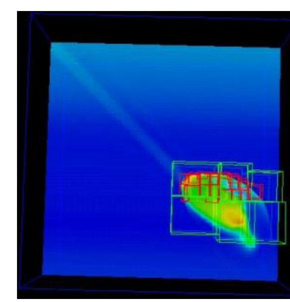

Supercomputing Issues

- The tyranny of the ***storage hierarchy***
- ***Parallelism***: doing multiple things at the same time

OSCER

What is OSCER?

- Multidisciplinary center
- Division of OU Information Technology
- Provides:
  - Supercomputing **education**
  - Supercomputing **expertise**
  - Supercomputing **resources**: hardware, storage, software
- For:
  - Undergrad students
  - Grad students
  - Staff
  - Faculty
  - Their collaborators (including **off campus**)

Who is OSCER? Academic Depts

- Aerospace & Mechanical Engr
- Anthropology
- Biochemistry & Molecular Biology
- Biological Survey
- Botany & Microbiology
- Chemical, Biological & Materials Engr
- Chemistry & Biochemistry
- Civil Engr & Environmental Science
- Computer Science
- Economics
- Electrical & Computer Engr
- Finance
- Health & Sport Sciences
- History of Science
- Industrial Engr
- Geography
- Geology & Geophysics
- Library & Information Studies
- Mathematics
- Meteorology
- Petroleum & Geological Engr
- Physics & Astronomy
- Psychology
- Radiological Sciences
- Surgery
- Zoology

E m ∃ w

More than 150 faculty & staff in **26 depts** in Colleges of Arts & Sciences, Atmospheric & Geographic Sciences, Business, Earth & Energy, Engineering, and Medicine – with **more to come**!

Who is OSCER? Groups

- Advanced Center for Genome Technology
- Center for Analysis & Prediction of Storms
- Center for Aircraft & Systems/Support Infrastructure
- Cooperative Institute for Mesoscale Meteorological Studies
- Center for Engineering Optimization
- Fears Structural Engineering Laboratory
- Human Technology Interaction Center
- Institute of Exploration & Development Geosciences
- Instructional Development Program
- Interaction, Discovery, Exploration, Adaptation Laboratory
- Microarray Core Facility
- OU Information Technology
- OU Office of the VP for Research
- Oklahoma Center for High Energy Physics
- Robotics, Evolution, Adaptation, and Learning Laboratory
- Sasaki Applied Meteorology Research Institute
- Symbiotic Computing Laboratory

E m ɜ w

Who? Oklahoma Collaborators

1. Cameron University
2. East Central University
3. Langston University
4. Northeastern State University
5. Northwestern Oklahoma State University
6. Oklahoma Baptist University
7. Oklahoma City University
20. Oklahoma Panhandle State University
21. Oklahoma School of Science & Mathematics
22. Oklahoma State University
23. Rogers State University
24. St. Gregory's University
25. Southeastern Oklahoma State University
26. Southwestern Oklahoma State University
27. University of Central Oklahoma
28. University of Oklahoma (Norman, HSC, Tulsa)
29. University of Science & Arts of Oklahoma
30. University of Tulsa

1. NOAA National Severe Storms Laboratory
2. NOAA Storm Prediction Center
3. Oklahoma Climatological Survey
4. Oklahoma Medical Research Foundation
5. OneNet
6. Samuel Roberts Noble Foundation
7. Tinker Air Force Base

OSCER has supercomputer users at <u>every public university</u> in Oklahoma, plus at many private universities and one high school.

Who Are the Users?

Over 750 users so far, including:

- Roughly equal split between students vs faculty/staff (students are the bulk of the active users);
- many off campus users (roughly 20%);
- … more being added **every month**.

Comparison: XSEDE, consisting of 7 resource provide sites across the US, has ~7500 unique users.

Biggest Consumers

- **Center for Analysis & Prediction of Storms**: daily real time weather forecasting
- **Oklahoma Center for High Energy Physics**: simulation and data analysis of banging tiny particles together at unbelievably high speeds
- **Chemical Engineering**: lots and lots of molecular dynamics

Why OSCER?

- Computational Science & Engineering has become **sophisticated enough** to take its place alongside experimentation and theory.
- **Most students** – and most faculty and staff – **don't learn much CSE**, because CSE is seen as needing too much computing background, and as needing HPC, which is seen as very hard to learn.
- **HPC can be hard to learn**: few materials for novices; most documents written for experts as reference guides.
- **We need a new approach**: HPC and CSE for computing novices – **OSCER's mandate!**

Why Bother Teaching Novices?

- Application scientists & engineers typically know their applications very well, much better than a collaborating computer scientist ever would.
- Commercial software lags far behind the research community.
- Many potential CSE users don't need full time CSE and HPC staff, just some help.
- One HPC expert can help dozens of research groups.
- Today's novices are tomorrow's top researchers, especially because today's top researchers will eventually retire.

What Does OSCER Do? Teaching

Science and engineering faculty from all over America learn supercomputing at OU by playing with a jigsaw puzzle (NCSI @ OU 2004).

What Does OSCER Do? Rounds

OU undergrads, grad students, staff and faculty learn how to use supercomputing in their specific research.

OSCER Resources

OK Cyberinfrastructure Initiative

- All academic institutions in Oklahoma are eligible to sign up for free use of OU's and OSU's centrally-owned CI resources.
- Other kinds of institutions (government, NGO, commercial) are eligible to use, though not necessarily for free.
- Everyone can participate in our CI education initiative.
- The Oklahoma Supercomputing Symposium, our annual conference, continues to be offered to all.

OCII Goals

- **Reach** institutions outside the mainstream of advanced computing needs.
- **Serve** every higher education institution in Oklahoma that has relevant curricula.
- **Educate** Oklahomans about advanced computing.
- **Attract** underrepresented populations and institution types into advanced computing.

OCII Service Methodologies Part 1

- **<u>Access (A)</u>**: to supercomputers and related technologies (20 academic institutions to date).
- **<u>Dissemination (D)</u>**: Oklahoma Supercomputing Symposium – annual advanced computing conference at OU (25).
- **<u>Education (E)</u>**: "Supercomputing in Plain English" (SiPE) workshop series: 11 talks about advanced computing, taught with stories, analogies and play rather than deep technical jargon. Have reached 166 institutions (academic, government, industry, nonprofit) in 42 US states and territories and 5 other countries (14 academic institutions in OK to date).

OCII Service Methodologies Part 2

- **Faculty Development (F)**: Workshops held at OU and OSU on advanced computing and computational science topics, sponsored by the National Computational Science Institute, the SC supercomputing conference series and the Linux Clusters Institute. Oklahoma is the only state to have hosted and co-taught multiple events sponsored by each of these (18).
- **Outreach (O)**: "Supercomputing in Plain English" (SiPE) overview talk (24).
- **Proposal Support (P)**: Letters of commitment for access to OCII resources; collaborations with OCII lead institutions (4).

OCII Service Methodologies Part 3

- **<u>Technology (T)</u>**: Got or helped get technology (e.g., network upgrade, mini-supercomputer, hi def video camera for telepresence) for that institution (14).
- **<u>Workforce Development (W)</u>** – (26)
  - Oklahoma Information Technology Mentorship Program (OITMP)
  - "A Day in the Life of an IT Professional" presentations to courses across the full spectrum of higher education.
  - Job shadowing opportunities and direct mentoring of individual students.
  - Institution Types: career techs, community colleges, regional universities, PhD-granting universities.
- Special effort to reach underrepresented populations: underrepresented minorities, non-PhD-granting, rural

OCII Institutions

1. **Bacone College (<u>MSI</u>, 30.9% AI, 24.0% AA): T**
2. **Cameron U (8.1% AI, 15.4% AA): A, D, E, F, O, T, W**
   Teaching <u>advanced computing course</u> using OSCER's supercomputer.
3. Canadian Valley Technology Center: W
4. **College of the Muscogee Nation (<u>Tribal</u>): O, T**
5. **Comanche Nation College (<u>Tribal</u>): D, O, T**
6. DeVry U Oklahoma City: D, F, O
7. **East Central U (<u>NASNI</u>, 20.4% AI): A, D, E, F, O, P, T, W**
   Taught <u>advanced computing course</u> using OSCER's supercomputer.
8. **Eastern Oklahoma State College (24.5% AI): W**

Average: ~3 (mean 3.4, median 3, mode 1)

9. **Eastern Oklahoma County Tech Center (10.4% AI): W**
10. Francis Tuttle Technology Center: D
11. **Great Plains Tech Center (11.7% AI): T, W**
12. **Gordon Cooper Technology Center (18.5% AI): D, O, W**
13. **Langston U (<u>HBCU</u>, 82.8% AA): A, D, E, F, O, P, T, W**
    <u>NSF Major Research Instrumentation</u> proposal for supercomputer submitted in 2012.

<u>Note</u>: Langston U (HBCU) and East Central U (NASNI) are the only two non-PhD-granting institutions to have benefited from every category of service that OCII provides.

AA = African American (7.4% OK population, 12.6% US population)
AI = American Indian (8.6% OK, 0.9% US)
H = Hispanic (8.9% OK, 16.3% US)
ALL = 24.9% OK, 29.8% US

HBCU: Historically Black College or University
NASNI = Native American Serving Non-Tribal Institution
MSI = Minority Serving Institution

OCII Institutions (cont'd)

14. Lawton Christian School (high school): W
15. **Metro Technology Centers (30.6% AA): D**
16. **Mid-America Technology Center (23.5% AI): D, T, W**
17. Moore Norman Technology Center: D
18. **Northeastern State U (NASNI, 28.3% AI): A, D, E, F, O, W**

    Taught computational chemistry course using OSCER's supercomputer.
19. Northwestern Oklahoma State U: A, F
20. Oklahoma Baptist U: A, D, E, F, O
21. Oklahoma Christian U: W

Average: ~3 (mean 3.4, median 3, mode 1)

22. Oklahoma City U: A, D, E, F, O, T, W

    *Educational Alliance for a Parallel Future mini-supercomputer proposal **funded** in 2011.*

    Teaching advanced computing course using OSCER's supercomputer (several times).
23. Oklahoma City Community College: W
24. **Oklahoma Panhandle State U (15.4% H): A, D, O, W**
25. **Oklahoma School of Science & Mathematics (high school): A, D, E, O, W**
26. **Oklahoma State U (PhD, 8.3% AI): A, D, E, F, O, T, W**

    *NSF Major Research Instrumentation proposal for supercomputer **funded** in 2011.*
27. **Oklahoma State U Institute of Technology (Comm College, 24.2% AI): W**

AA = African American (7.4% OK population, 12.6% US population)
AI = American Indian (8.6% OK, 0.9% US)
H = Hispanic (8.9% OK, 16.3% US)
ALL = 24.9% OK, 29.8% US

HBCU: Historically Black College or University
NASNI = Native American Serving Non-Tribal Institution
MSI = Minority Serving Institution

OCII Institutions (cont'd)

28. Oklahoma State U Oklahoma City (Comm College): O, W
29. Oral Roberts U: A, F, O, W
30. **Pawnee Nation College (Tribal): T**
31. **Pontotoc Technology Center (30.4% AI): W**
32. **Rogers State U (13.9% AI): A, D, F, O**
33. **Rose State College (17.4% AA): W**
34. St. Gregory's U: A, D, E, F, O
35. **Southeastern Oklahoma State U (NASNI, 29.6% AI): A, D, E, F, O, T, W**

    Educational Alliance for a Parallel Future mini-supercomputer proposal ***funded*** *in 2011.*
36. Southern Nazarene U: A, D, F, O, P, T, W

    Teaching computational chemistry course using OSCER's supercomputer.
37. Southwestern Oklahoma State U: A, D, E, F, O
38. U Central Oklahoma: A, D, E, F, O, W

    NSF Major Research Instrumentation proposal for supercomputer submitted in 2011-12.
39. U Oklahoma (PhD): A, D, E, F, O, P, T, W

    NSF Major Research Instrumentation proposal for large scale storage ***funded*** *in 2010.*
40. U Phoenix: D
41. **U of Science & Arts of Oklahoma (14.1% AI): A, O**
42. U Tulsa (PhD): A, D, E, F, O

    Taught bioinformatics course using OSCER's supercomputer.

Average: ~3 (mean 3.4, median 3, mode 1)

AA = African American (7.4% OK population, 12.6% US population)
AI = American Indian (8.6% OK, 0.9% US)
H = Hispanic (8.9% OK, 16.3% US)
ALL = 24.9% OK, 29.8% US

HBCU: Historically Black College or University
NASNI = Native American Serving Non-Tribal Institution
MSI = Minority Serving Institution

NEW SUPERCOMPUTER!

874 Intel Xeon CPU chips/6992 cores

412 dual socket/oct core Sandy Bridge 2.0 GHz, 32 GB
23 dual socket/oct core Sandy Bridge 2.0 GHz, 64 GB
1 quad socket/oct core Westmere, 2.13 GHz, 1 TB

15,680 GB RAM

~360 TB global disk

QLogic Infiniband
(16.67 Gbps, ~1 microsec latency)

Dell Force10 Gigabit/10G Ethernet

Red Hat Enterprise Linux 6

Peak speed: 111.6 TFLOPs*

*TFLOPs: trillion calculations per second

Just over 3x (300%) as fast as our 2008-12 supercomputer.

Just over 100x (10,000%) as fast as our first cluster supercomputer in 2002.

boomer.oscer.ou.edu

What is a Cluster Supercomputer?

"… [W]hat a ship is … It's not just a keel and hull and a deck and sails. That's what a ship needs. But what a ship is ... is freedom."

– Captain Jack Sparrow

"Pirates of the Caribbean"

What a Cluster is ….

A cluster **needs** of a collection of small computers, called ***nodes***, hooked together by an ***interconnection network*** (or ***interconnect*** for short).

It also **needs** software that allows the nodes to communicate over the interconnect.

But what a cluster **is** … is all of these components working together as if they're one big computer ... a **super** computer.

An Actual Cluster

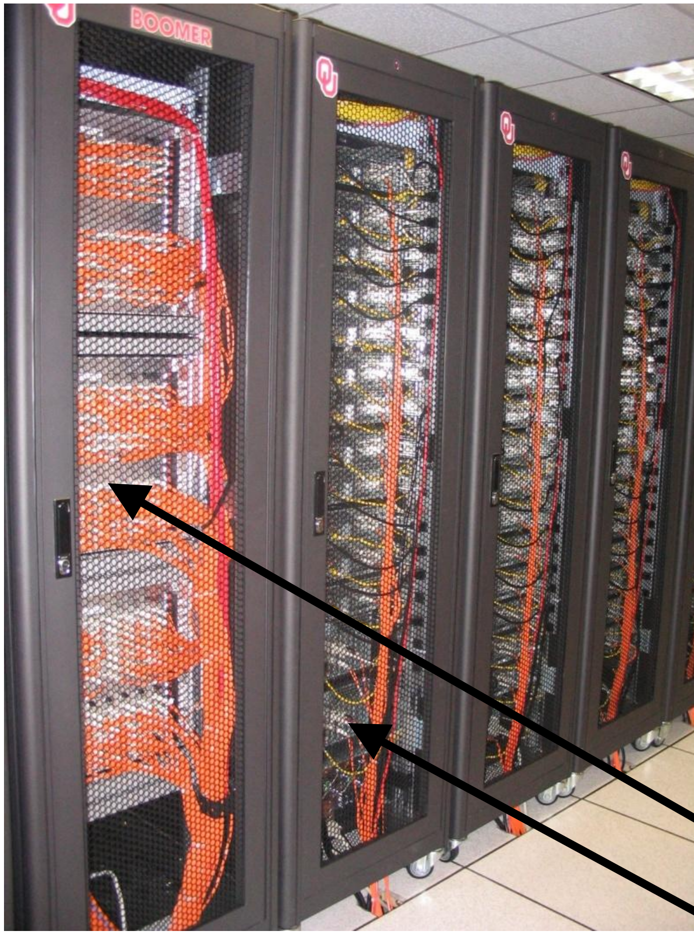

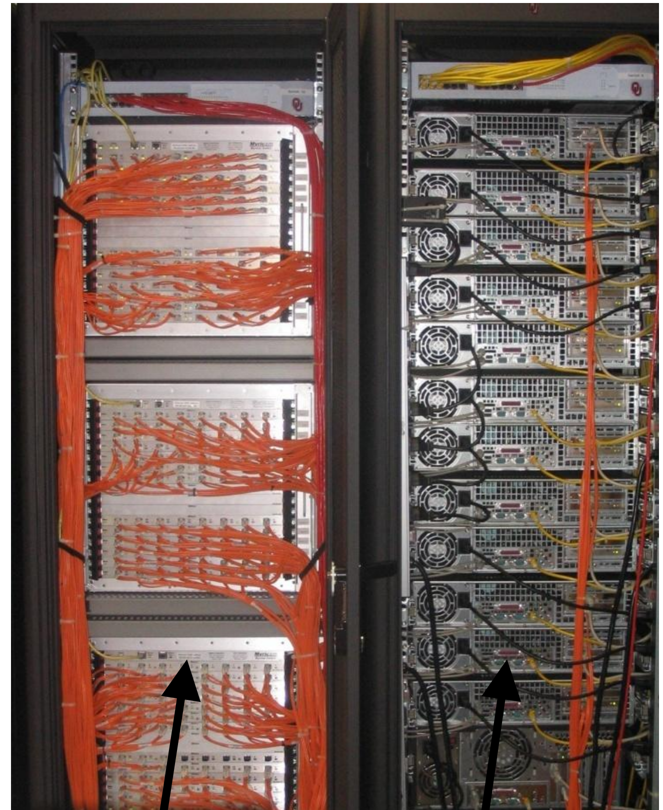

Interconnect

Nodes

Also named Boomer, in service 2002-5.

Condor Pool

- Condor is a software technology that allows idle desktop PCs to be used for number crunching.
- OU IT has deployed a large Condor pool (795 desktop PCs in IT student labs all over campus).
- It provides a huge amount of additional computing power – more than was available in all of OSCER in 2005.
- 20+ TFLOPs peak compute speed.
- And, the cost is very very low – almost literally free.
- Also, we’ve been seeing empirically that Condor gets about 80% of each PC’s time.

National Lambda Rail

Internet2

Internet2 Network

ciena | INDIANA UNIVERSITY | infinera | Juniper NETWORKS | Level(3)

CONNECTORS

- 3ROX
- CENIC
- CIC OmniPoP
- Drexel University
- GPN
- Indiana GigaPoP
- KyRON
- LEARN
- LONI
- MAGPI
- MAX
- MCNC
- Merit Network
- MREN
- NOX
- NYSERNet
- Oregon Gigapop
- Pacific Northwest GigaPoP
- SoX
- University of Memphis
- University of New Mexico
- University of South Florida
- University of Utah/UEN

www.internet2.edu

NSF EPSCoR C2 Grant

Oklahoma has been awarded a National Science Foundation EPSCoR RII Intra- campus and Inter-campus Cyber Connectivity (C2) grant (PI Neeman), a collaboration among OU, OneNet and several other academic and nonprofit institutions, which is:

- upgrading the statewide ring from routed components to optical components, making it straightforward and affordable to provision dedicated "lambda" circuits within the state;
- upgrading several institutions' connections;
- providing telepresence capability to institutions statewide;
- providing IT professionals to speak to IT and CS courses about what it's like to do IT for a living.

NSF MRI Grant: Petascale Storage

OU has been awarded an National Science Foundation Major Research Instrumentation (MRI) grant (PI Neeman).

We'll purchase and deploy a combined disk/tape bulk storage archive:

- the NSF budget pays for the hardware, software and warranties/maintenance for 3 years;
- OU cost share and institutional commitment pay for space, power, cooling and labor, as well as maintenance after the 3 year project period;
- individual users (e.g., faculty across Oklahoma) pay for the media (disk drives and tape cartridges).

A Quick Primer on Hardware

Henry's Laptop

Dell Latitude Z600[4]

- Intel Core2 Duo SU9600
  1.6 GHz w/3 MB L2 Cache
- 4 GB 1066 MHz DDR3 SDRAM
- 256 GB SSD Hard Drive
- DVD±RW/CD-RW Drive (8x)
- 1 Gbps Ethernet Adapter

Typical Computer Hardware

- Central Processing Unit
- Primary storage
- Secondary storage
- Input devices
- Output devices

Central Processing Unit

Also called ***CPU*** or ***processor***: the "brain"

Components

- ***Control Unit***: figures out what to do next – for example, whether to load data from memory, or to add two values together, or to store data into memory, or to decide which of two possible actions to perform (***branching***)
- ***Arithmetic/Logic Unit***: performs calculations – for example, adding, multiplying, checking whether two values are equal
- ***Registers***: where data reside that are **being used right now**

Primary Storage

- ***Main Memory***
  - Also called ***RAM*** (“Random Access Memory”)
  - Where data reside when they’re **being used by a program that’s currently running**
- ***Cache***
  - Small area of much faster memory
  - Where data reside when they’re **about to be used** and/or **have been used recently**
- Primary storage is ***volatile***: values in primary storage disappear when the power is turned off.

Secondary Storage

- Where data and programs reside that are going to be used **in the future**
- Secondary storage is ***non-volatile***: values **don't** disappear when power is turned off.
- Examples: hard disk, CD, DVD, Blu-ray, magnetic tape, floppy disk
- Many are ***portable***: can pop out the CD/DVD/tape/floppy and take it with you

Input/Output

- Input devices – for example, keyboard, mouse, touchpad, joystick, scanner
- Output devices – for example, monitor, printer, speakers

The Tyranny of the Storage Hierarchy

The Storage Hierarchy

Fast, expensive, few

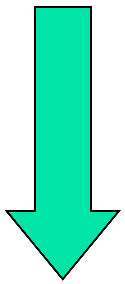

Slow, cheap, a lot

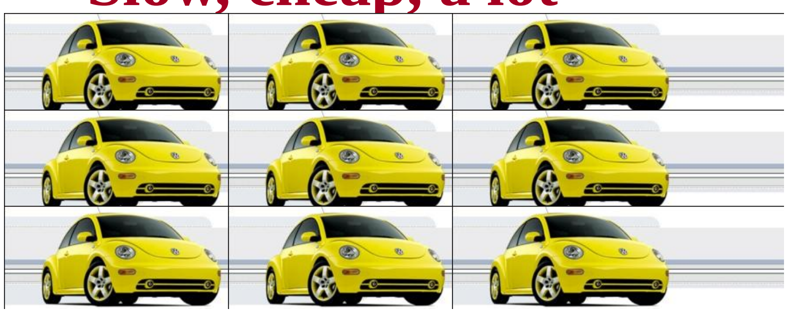

[5]

- Registers
- Cache memory
- Main memory (RAM)
- Hard disk
- Removable media (CD, DVD etc)
- Internet

RAM is Slow

The speed of data transfer between Main Memory and the CPU is much slower than the speed of calculating, so the CPU spends most of its time waiting for data to come in or go out.

CPU 307 GB/sec[6]

Bottleneck

4.4 GB/sec[7] (1.4%)

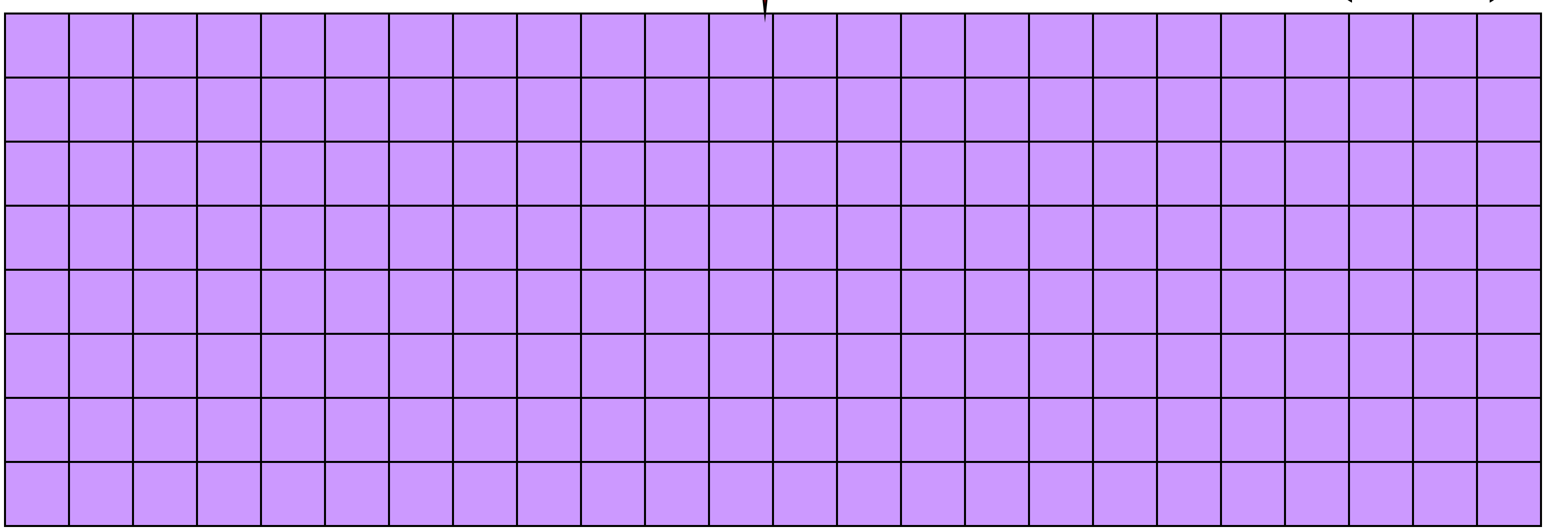

Why Have Cache?

Cache is much closer to the speed of the CPU, so the CPU doesn't have to wait nearly as long for stuff that's already in cache: it can do more operations per second!

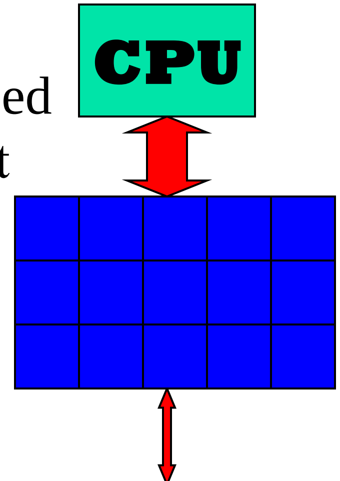

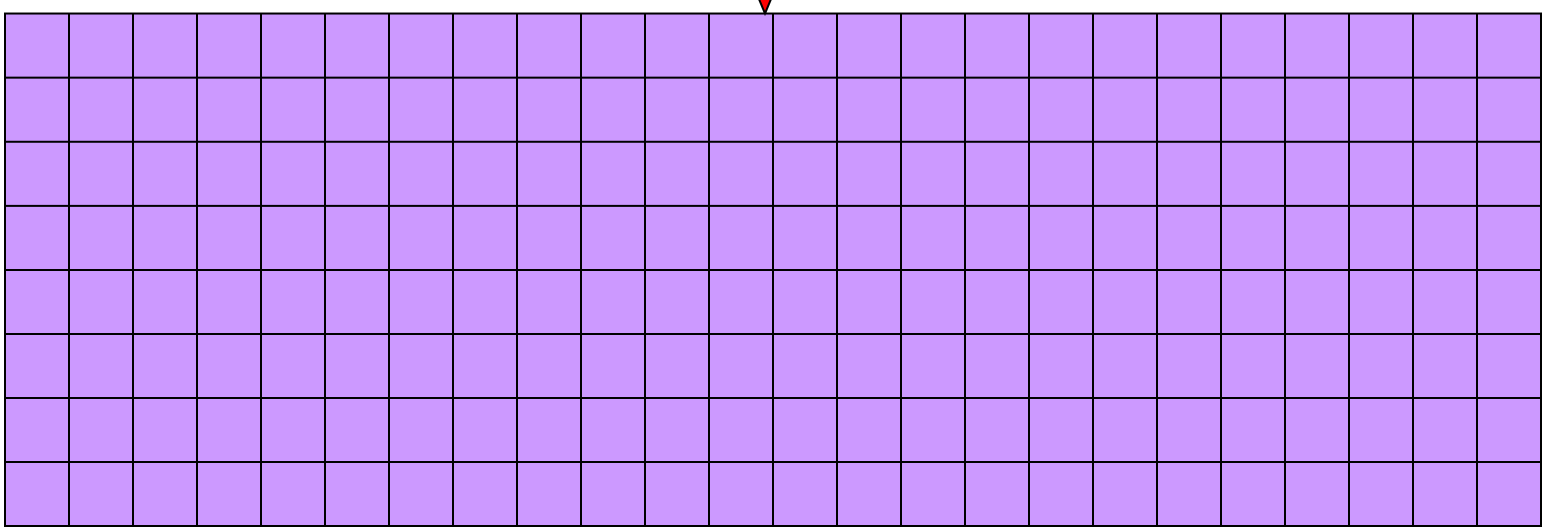

Henry's Laptop

Dell Latitude Z600[4]

- Intel Core2 Duo SU9600
  1.6 GHz w/3 MB L2 Cache
- 4 GB 1066 MHz DDR3 SDRAM
- 256 GB SSD Hard Drive
- DVD±RW/CD-RW Drive (8x)
- 1 Gbps Ethernet Adapter

Storage Speed, Size, Cost

Henry's Laptop	Registers (Intel Core2 Duo 1.6 GHz)	Cache Memory (L2)	Main Memory (1066MHz DDR3 SDRAM)	Hard Drive (SSD)	Ethernet (1000 Mbps)	DVD±R (16x)	Phone Modem (56 Kbps)
Speed (MB/sec) [peak]	314,573[6] (12,800 MFLOP/s*)	27,276 [7]	4500 [7]	250 [9]	125	22 [10]	0.007
Size (MB)	464 bytes** [11]	3	4096	256,000	unlimited	unlimited	unlimited
Cost ($/MB)	–	$285 [12]	$0.03 [12]	$0.002 [12]	charged per month (typically)	$0.00005 [12]	charged per month (typically)

* MFLOP/s: millions of floating point operations per second

** 16 64-bit general purpose registers, 8 80-bit floating point registers, 16 128-bit floating point vector registers

Why the Storage Hierarchy?

Why does the Storage Hierarchy always work? Why are faster forms of storage more expensive and slower forms cheaper?

Proof by contradiction:

Suppose there were a storage technology that was **slow** and **expensive**.

How much of it would you buy?

Comparison

- Zip: Cartridge $7.15 (2.9 cents per MB), speed 2.4 MB/sec
- Blu-Ray: Disk $5 ($0.0002 per MB), speed 19 MB/sec

Not surprisingly, no one buys Zip drives any more.

Parallelism

Parallelism

Parallelism means doing multiple things at the same time: you can get more work done in the same time.

Less fish …

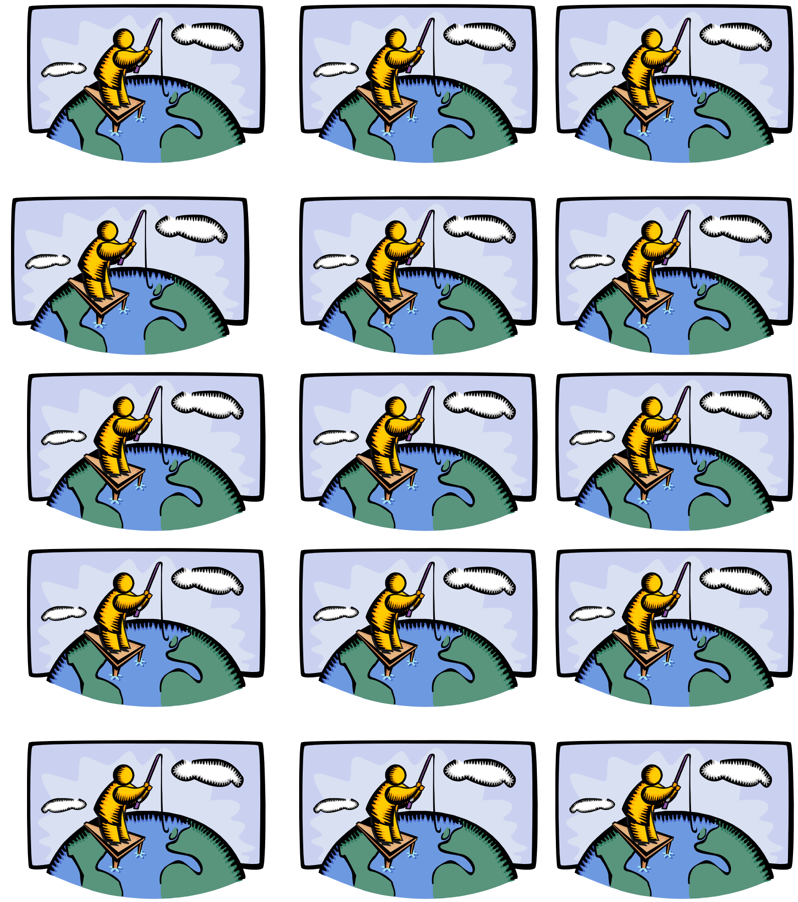

More fish!

The Jigsaw Puzzle Analogy

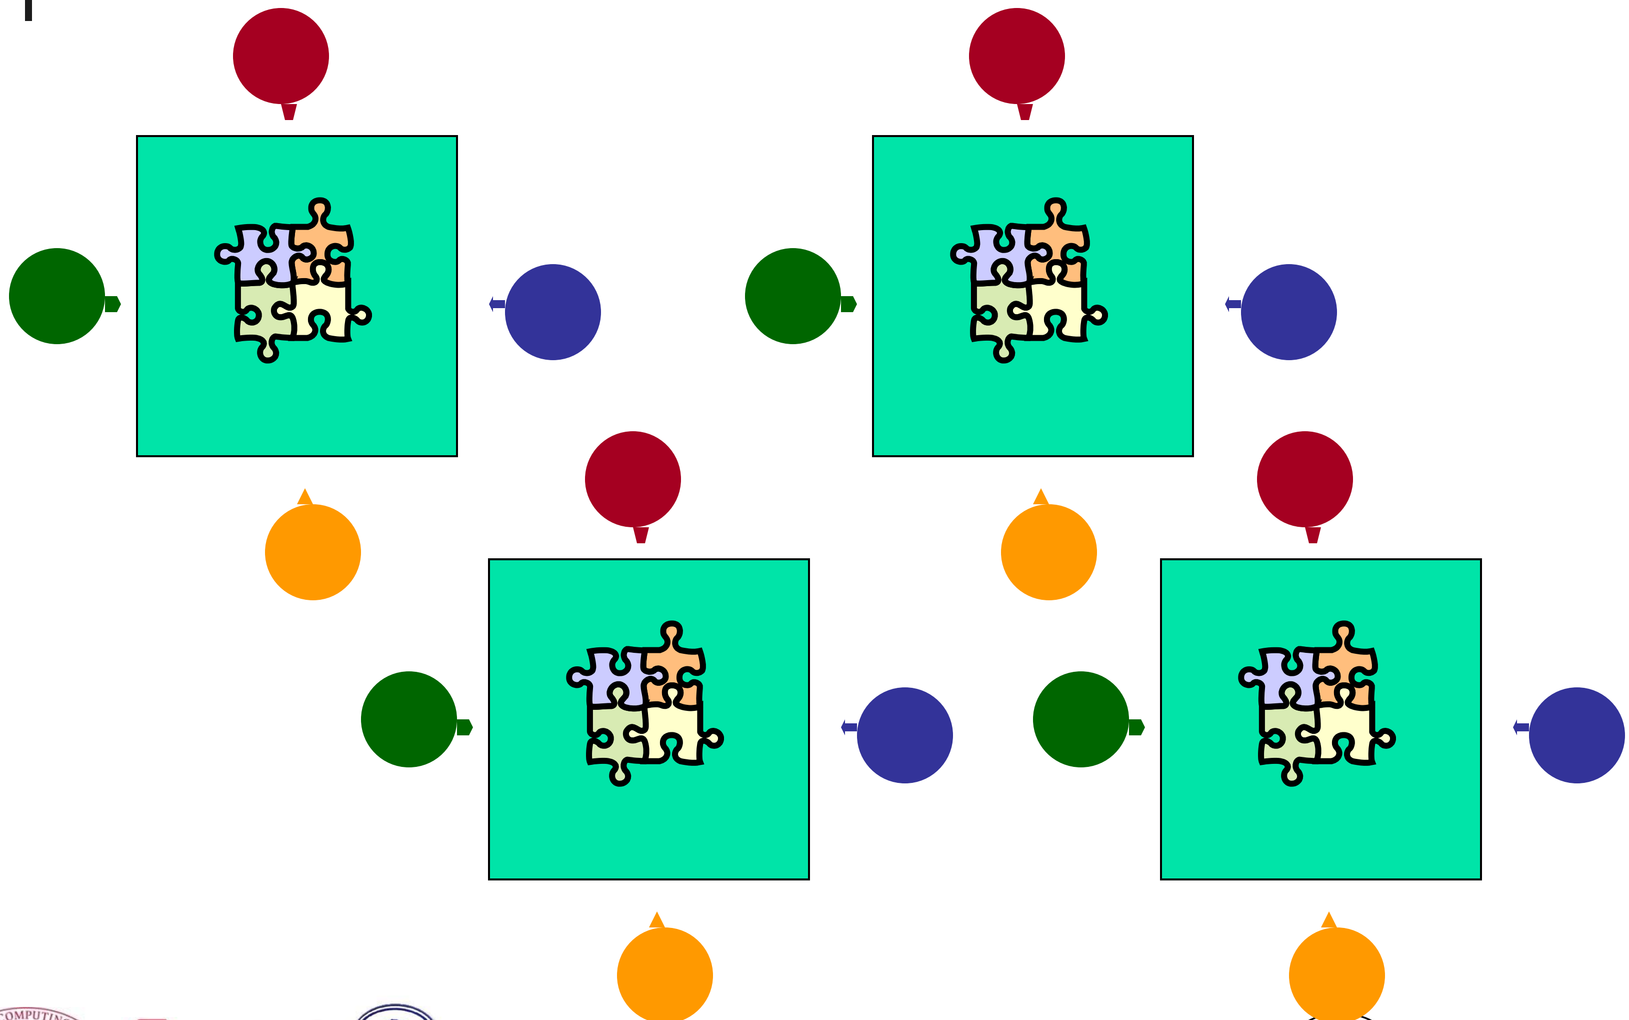

Serial Computing

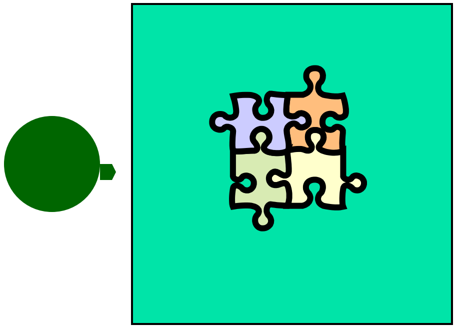

Suppose you want to do a jigsaw puzzle that has, say, a thousand pieces.

We can imagine that it'll take you a certain amount of time. Let's say that you can put the puzzle together in an hour.

Shared Memory Parallelism

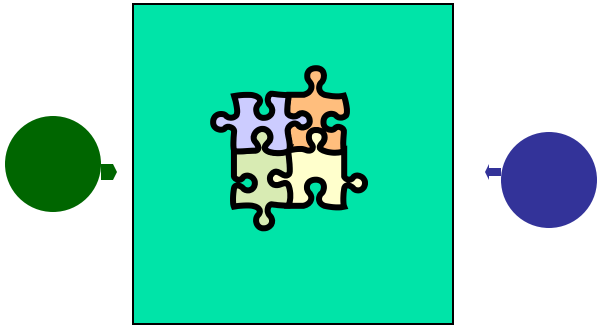

If Scott sits across the table from you, then he can work on his half of the puzzle and you can work on yours. Once in a while, you'll both reach into the pile of pieces at the same time (you'll ***contend*** for the same resource), which will cause a little bit of slowdown. And from time to time you'll have to work together (***communicate***) at the interface between his half and yours. The speedup will be nearly 2-to-1: y'all might take 35 minutes instead of 30.

The More the Merrier?

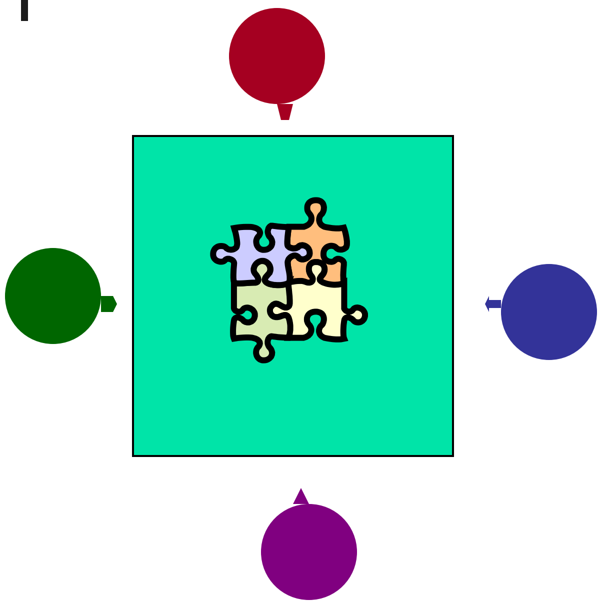

Now let's put Paul and Charlie on the other two sides of the table. Each of you can work on a part of the puzzle, but there'll be a lot more contention for the shared resource (the pile of puzzle pieces) and a lot more communication at the interfaces. So y'all will get noticeably less than a 4-to-1 speedup, but you'll still have an improvement, maybe something like 3-to-1: the four of you can get it done in 20 minutes instead of an hour.

Diminishing Returns

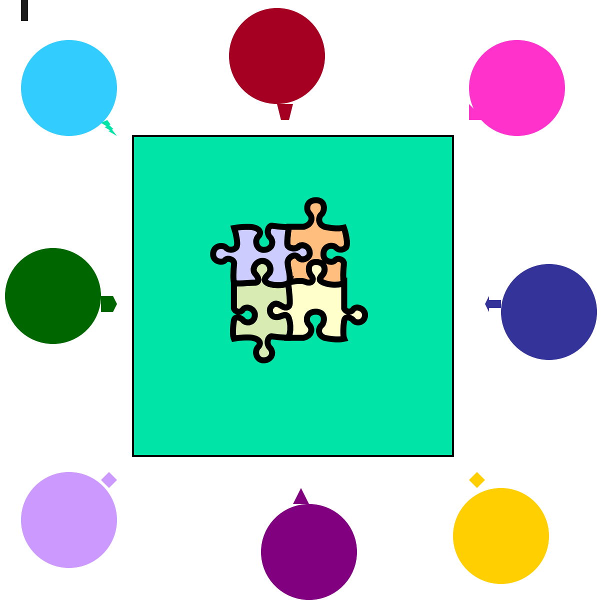

If we now put Dave and Tom and Horst and Brandon on the corners of the table, there's going to be a whole lot of contention for the shared resource, and a lot of communication at the many interfaces. So the speedup y'all get will be much less than we'd like; you'll be lucky to get 5-to-1.

So we can see that adding more and more workers onto a shared resource is eventually going to have a diminishing return.

Distributed Parallelism

Now let’s try something a little different. Let’s set up two tables, and let’s put you at one of them and Scott at the other. Let’s put half of the puzzle pieces on your table and the other half of the pieces on Scott’s. Now y’all can work completely independently, without any contention for a shared resource. **BUT**, the cost per communication is **MUCH** higher (you have to scootch your tables together), and you need the ability to split up (***decompose***) the puzzle pieces reasonably evenly, which may be tricky to do for some puzzles.

More Distributed Processors

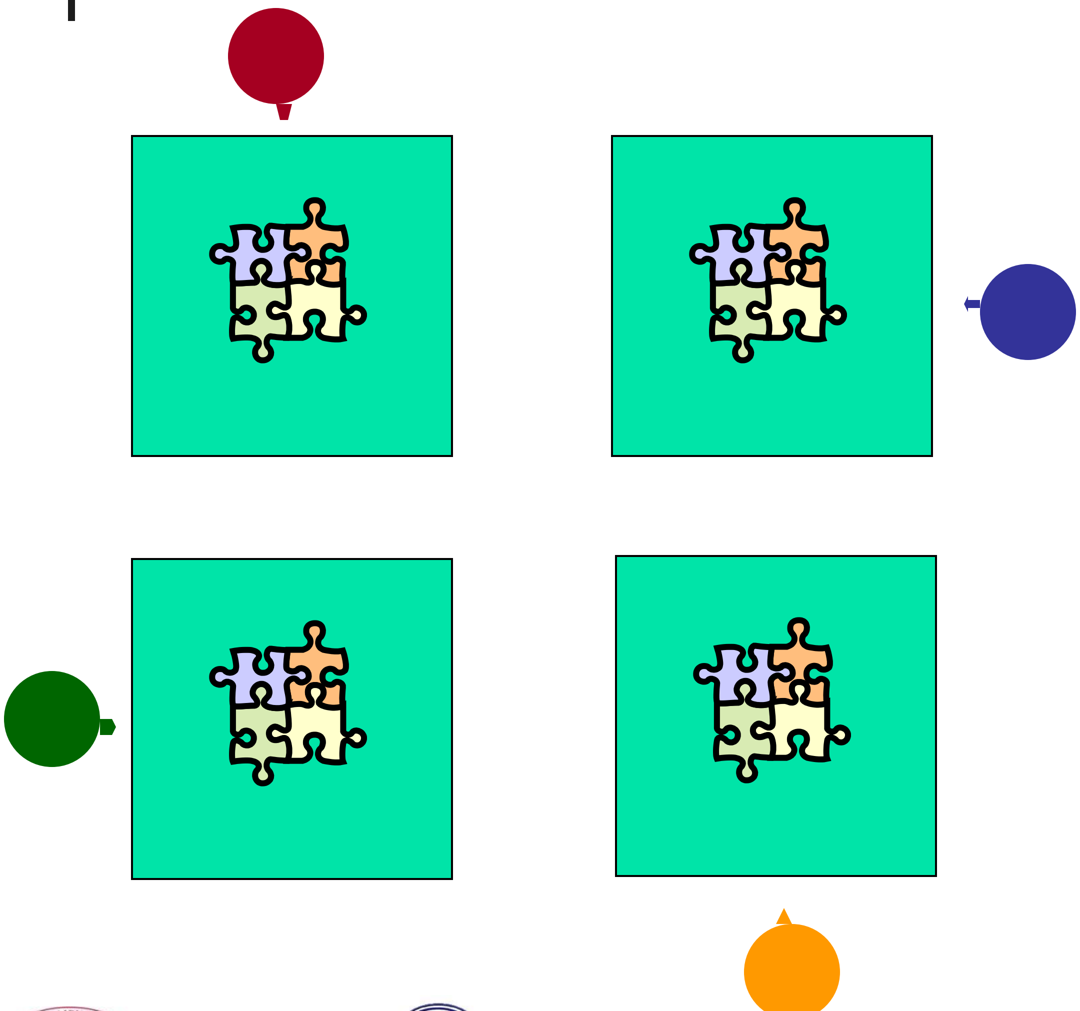

It's a lot easier to add more processors in distributed parallelism. But, you always have to be aware of the need to decompose the problem and to communicate among the processors. Also, as you add more processors, it may be harder to ***load balance*** the amount of work that each processor gets.

Load Balancing

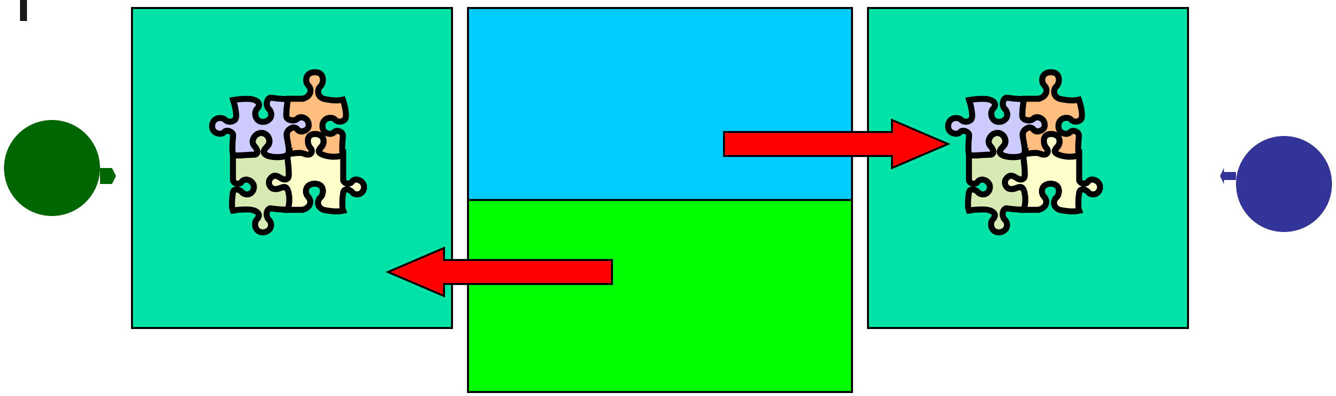

Load balancing means ensuring that everyone completes their workload at roughly the same time.

For example, if the jigsaw puzzle is half grass and half sky, then you can do the grass and Scott can do the sky, and then y'all only have to communicate at the horizon – and the amount of work that each of you does on your own is roughly equal. So you'll get pretty good speedup.

Load Balancing

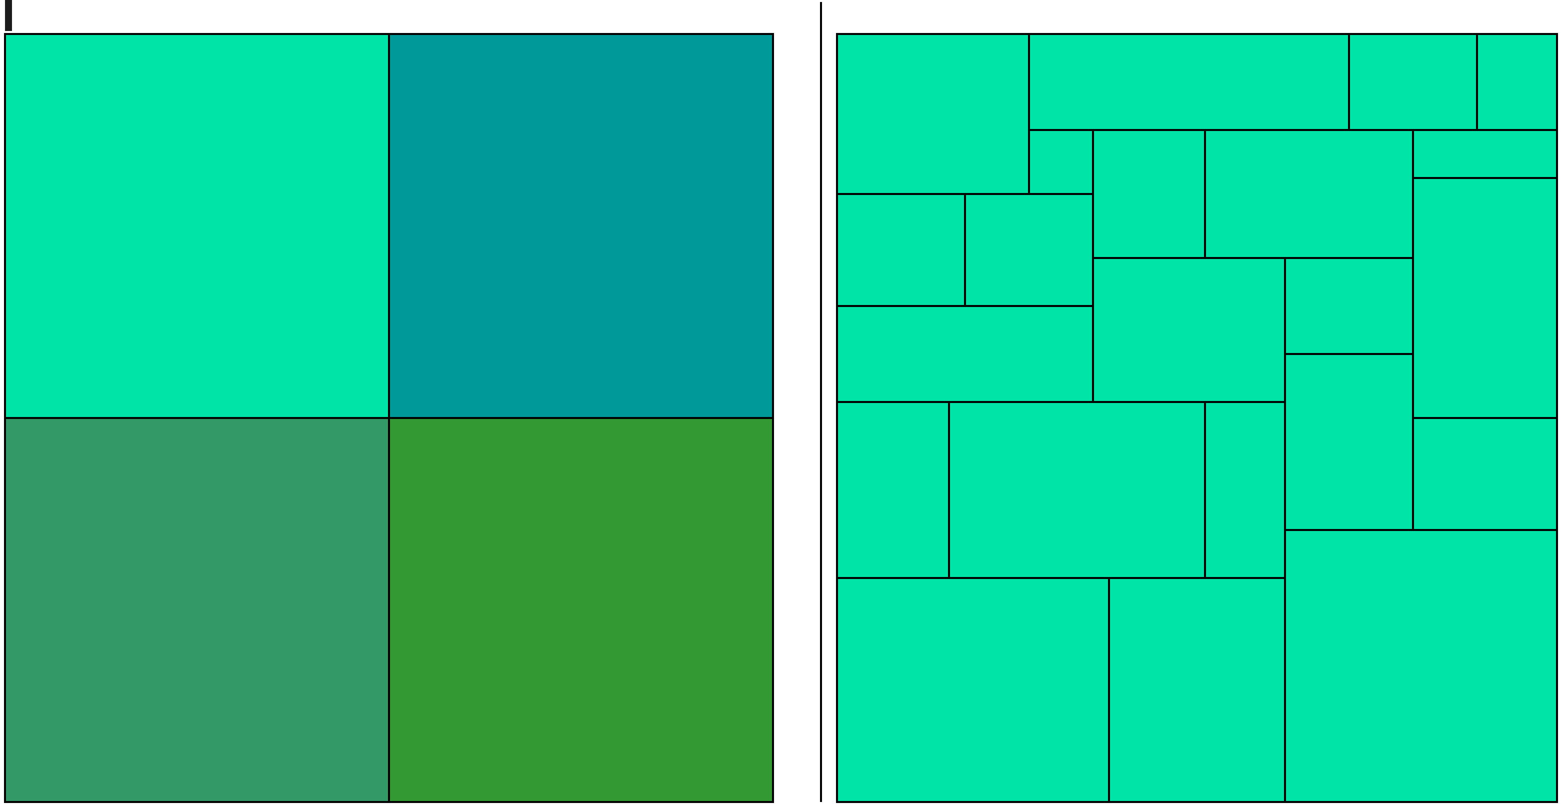

Load balancing can be easy, if the problem splits up into chunks of roughly equal size, with one chunk per processor. Or load balancing can be very hard.

Load Balancing

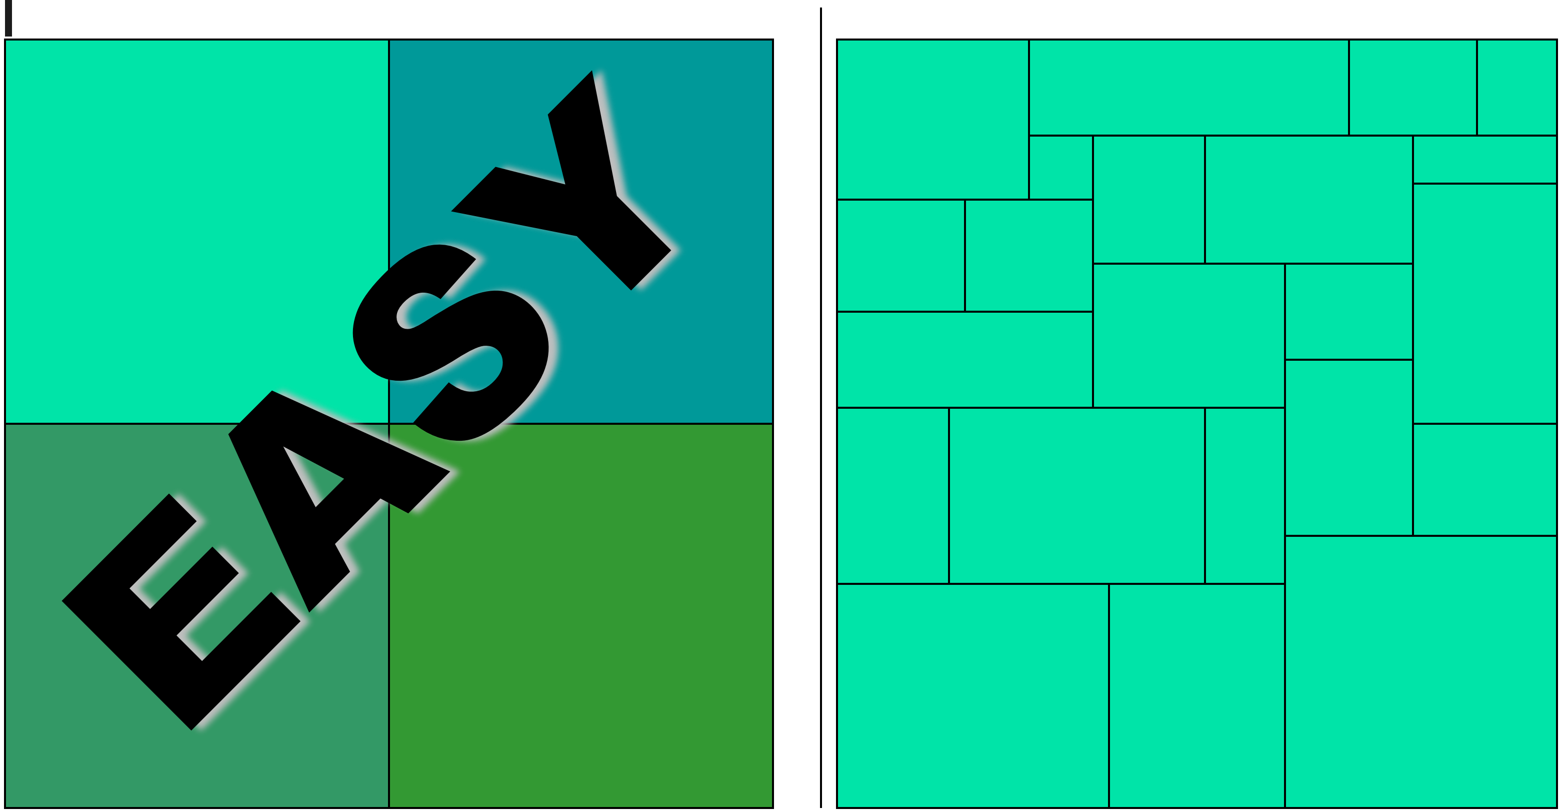

Load balancing can be easy, if the problem splits up into chunks of roughly equal size, with one chunk per processor. Or load balancing can be very hard.

Load Balancing

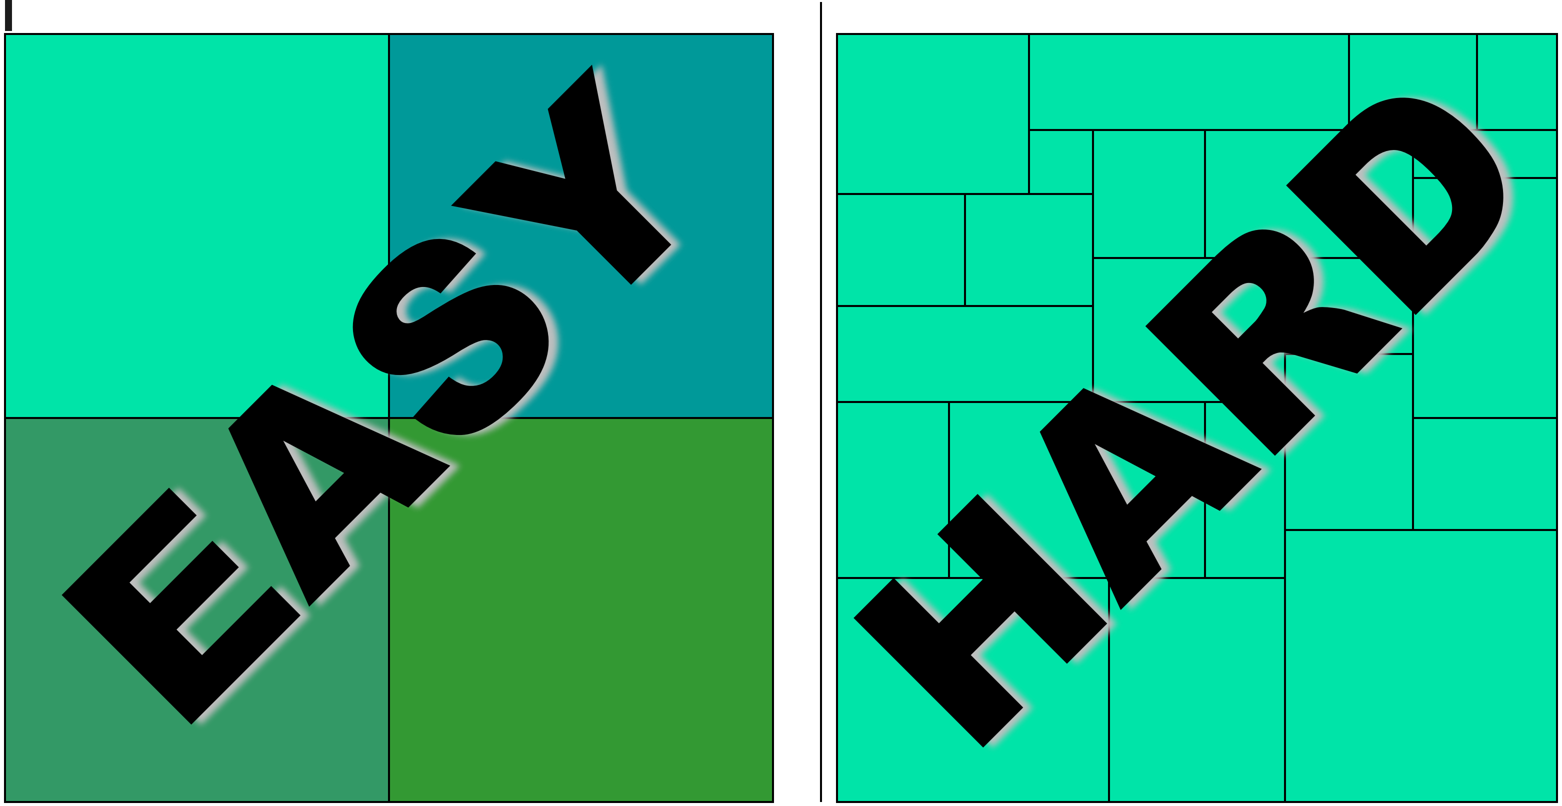

Load balancing can be easy, if the problem splits up into chunks of roughly equal size, with one chunk per processor. Or load balancing can be very hard.

Moore’s Law

Moore's Law

In 1965, Gordon Moore was an engineer at Fairchild Semiconductor.

He noticed that the number of transistors that could be squeezed onto a chip was doubling about every 2 years.

It turns out that computer speed is roughly proportional to the number of transistors per unit area.

Moore wrote a paper about this concept, which became known as ***"Moore's Law."***

Fastest Supercomputer vs. Moore

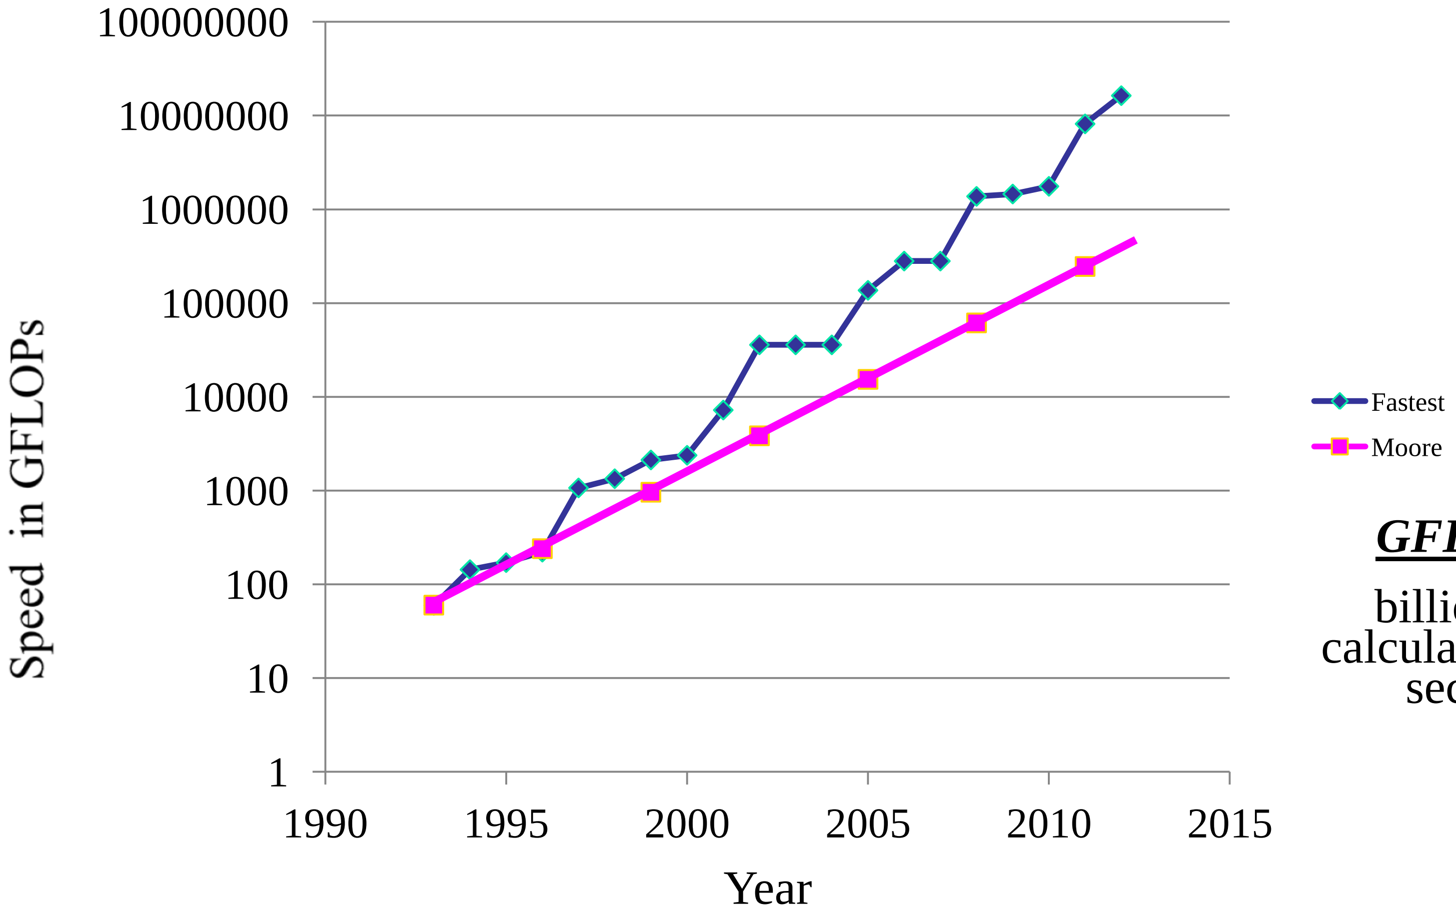

GFLOPs:

billions of calculations per second

Fastest Supercomputer vs. Moore

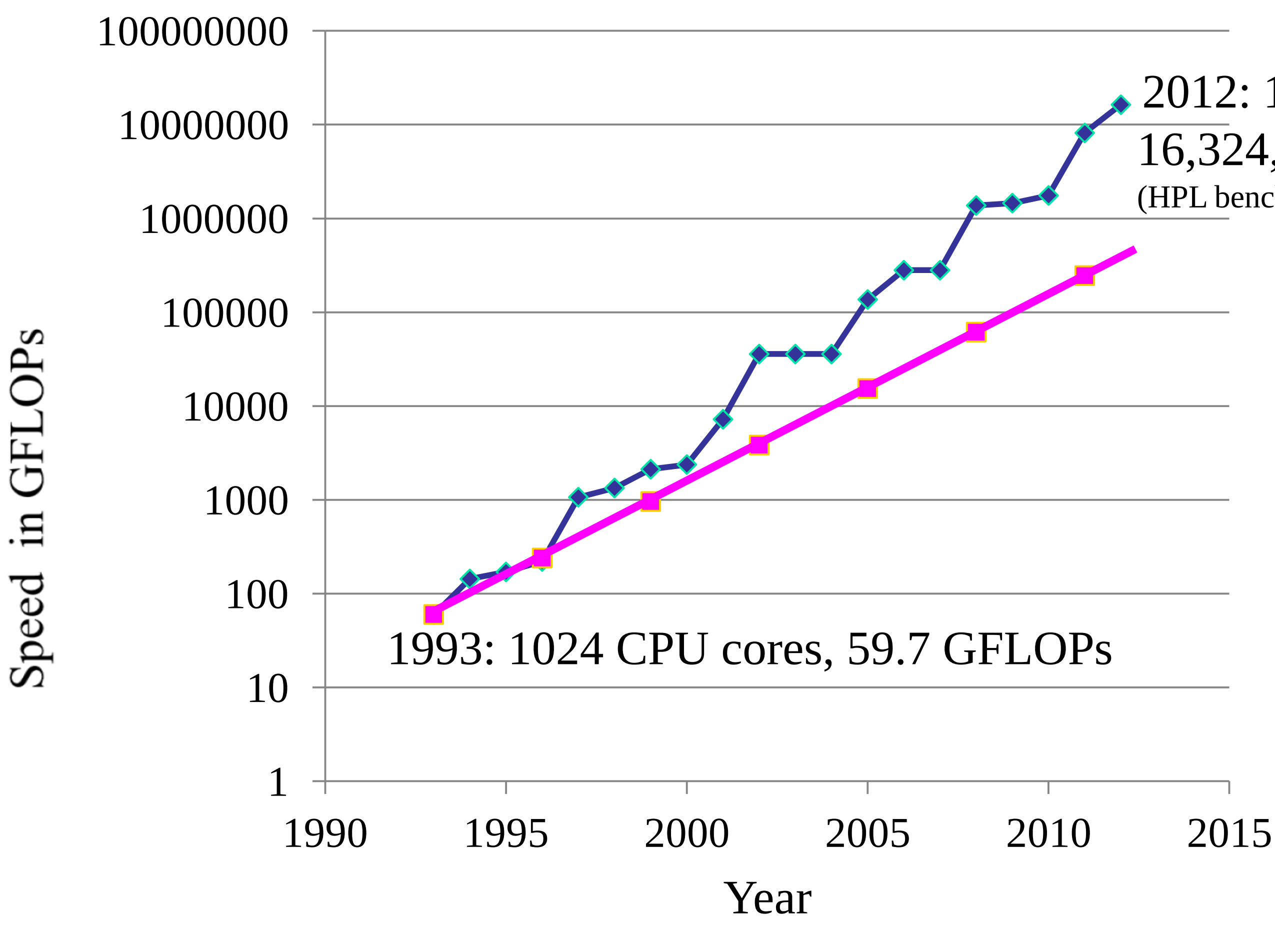

GFLOPs:

billions of calculations per second

Gap: Supercomputers were 35x higher than Moore in 2011.

Moore: Uncanny!

- Nov 1971: Intel 4004 – 2300 transistors
- March 2010: Intel Nehalem Beckton – 2.3 billion transistors
- Factor of 1M improvement in 38 1/3 years
- $2^{(38.33 \text{ years} / 1.9232455)} = 1{,}000{,}000$

So, transistor density has doubled every 23 months:

UNCANNILY ACCURATE PREDICTION!

Moore's Law in Practice

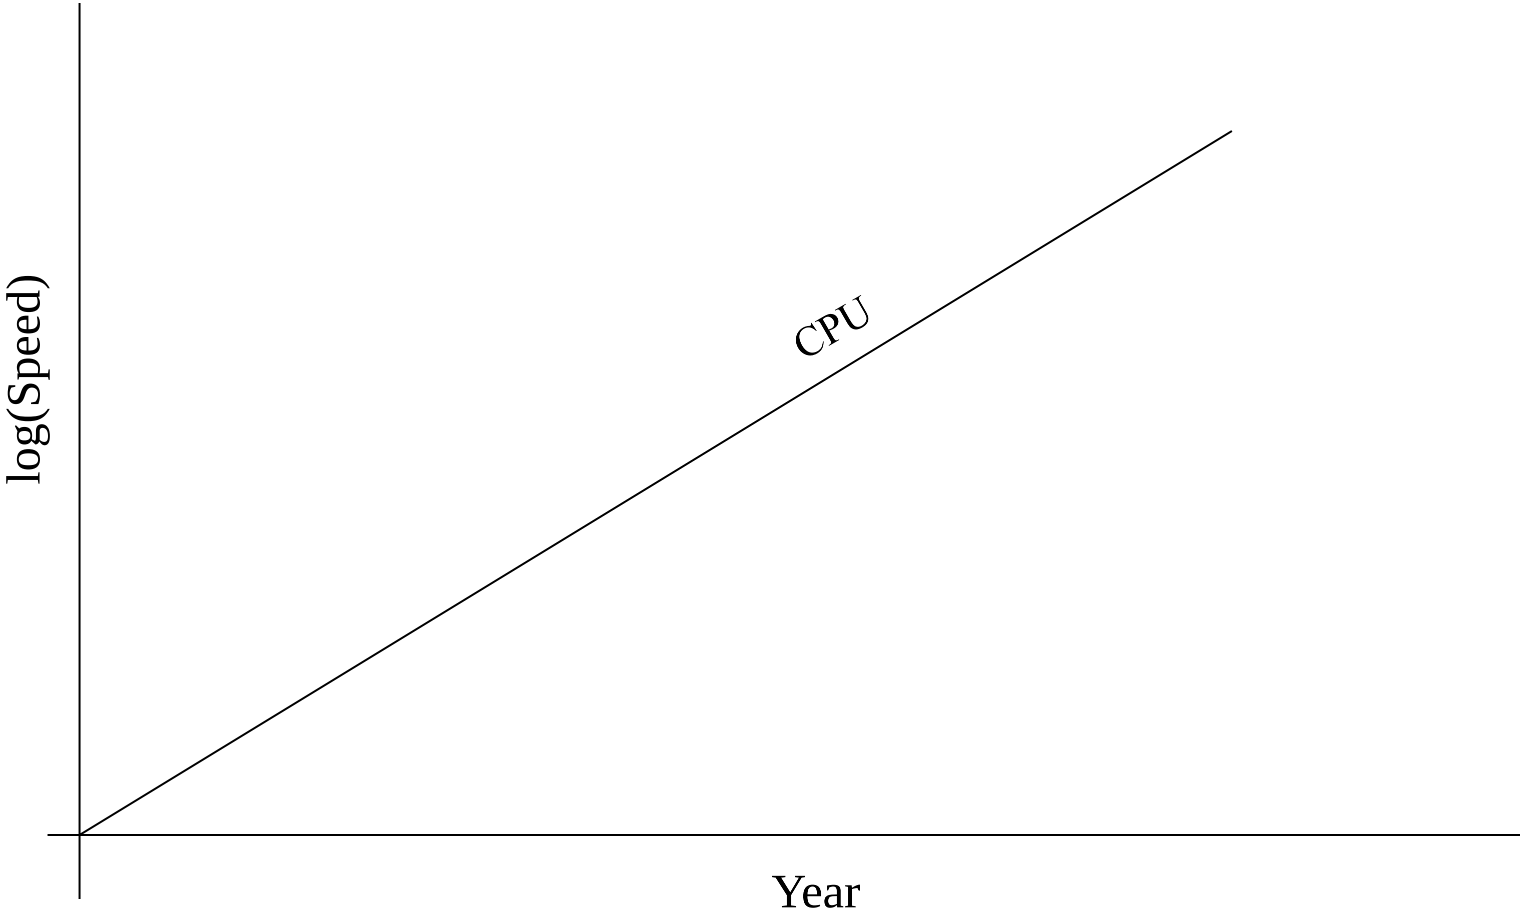

Moore's Law in Practice

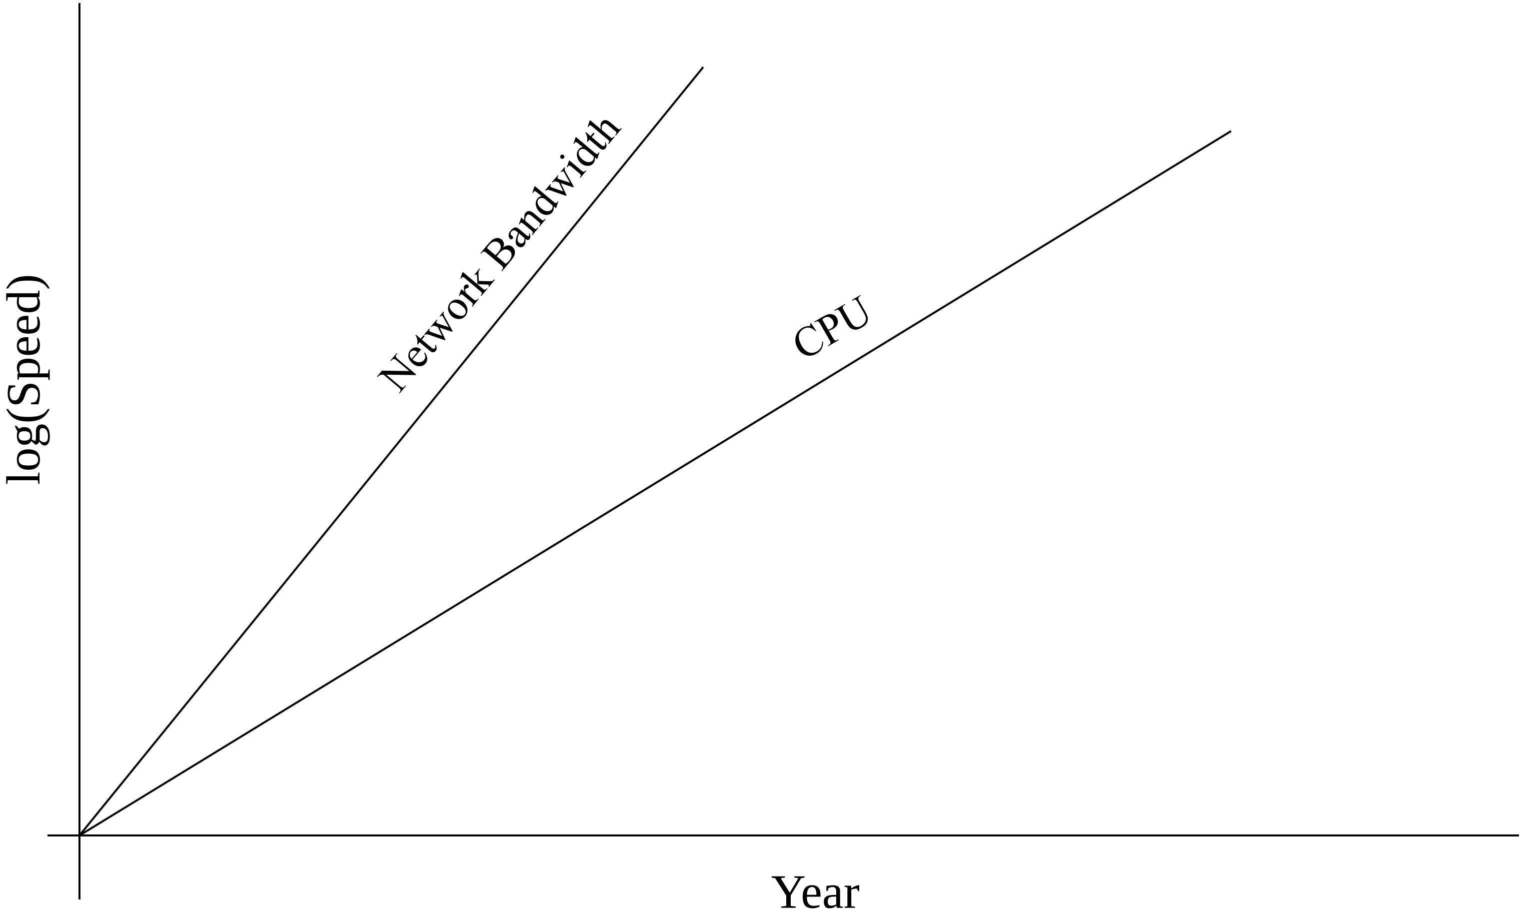

Moore's Law in Practice

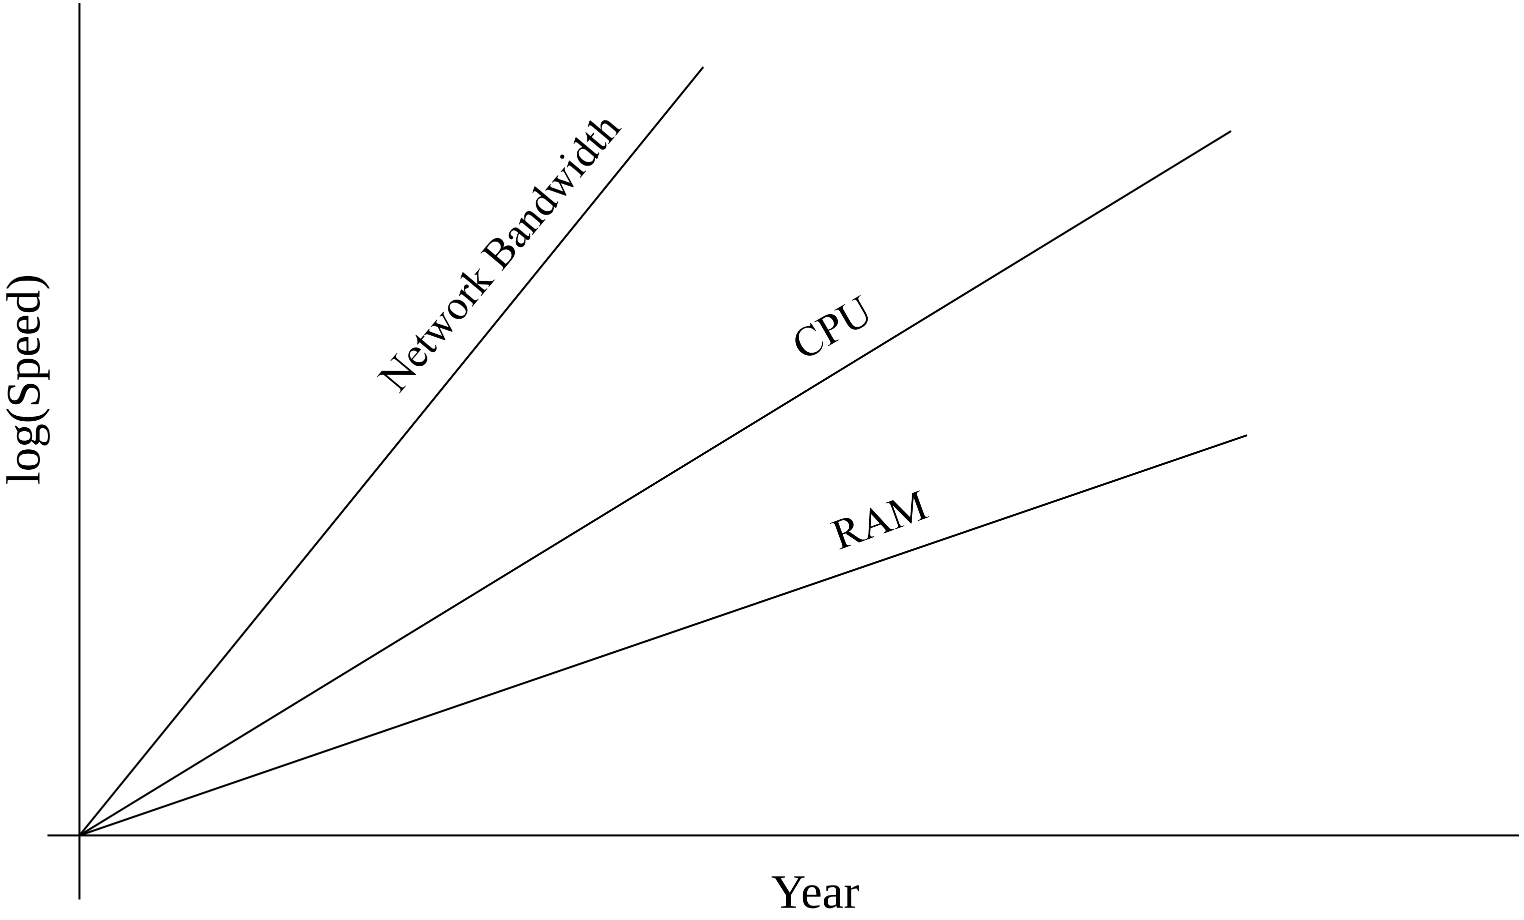

Moore’s Law in Practice

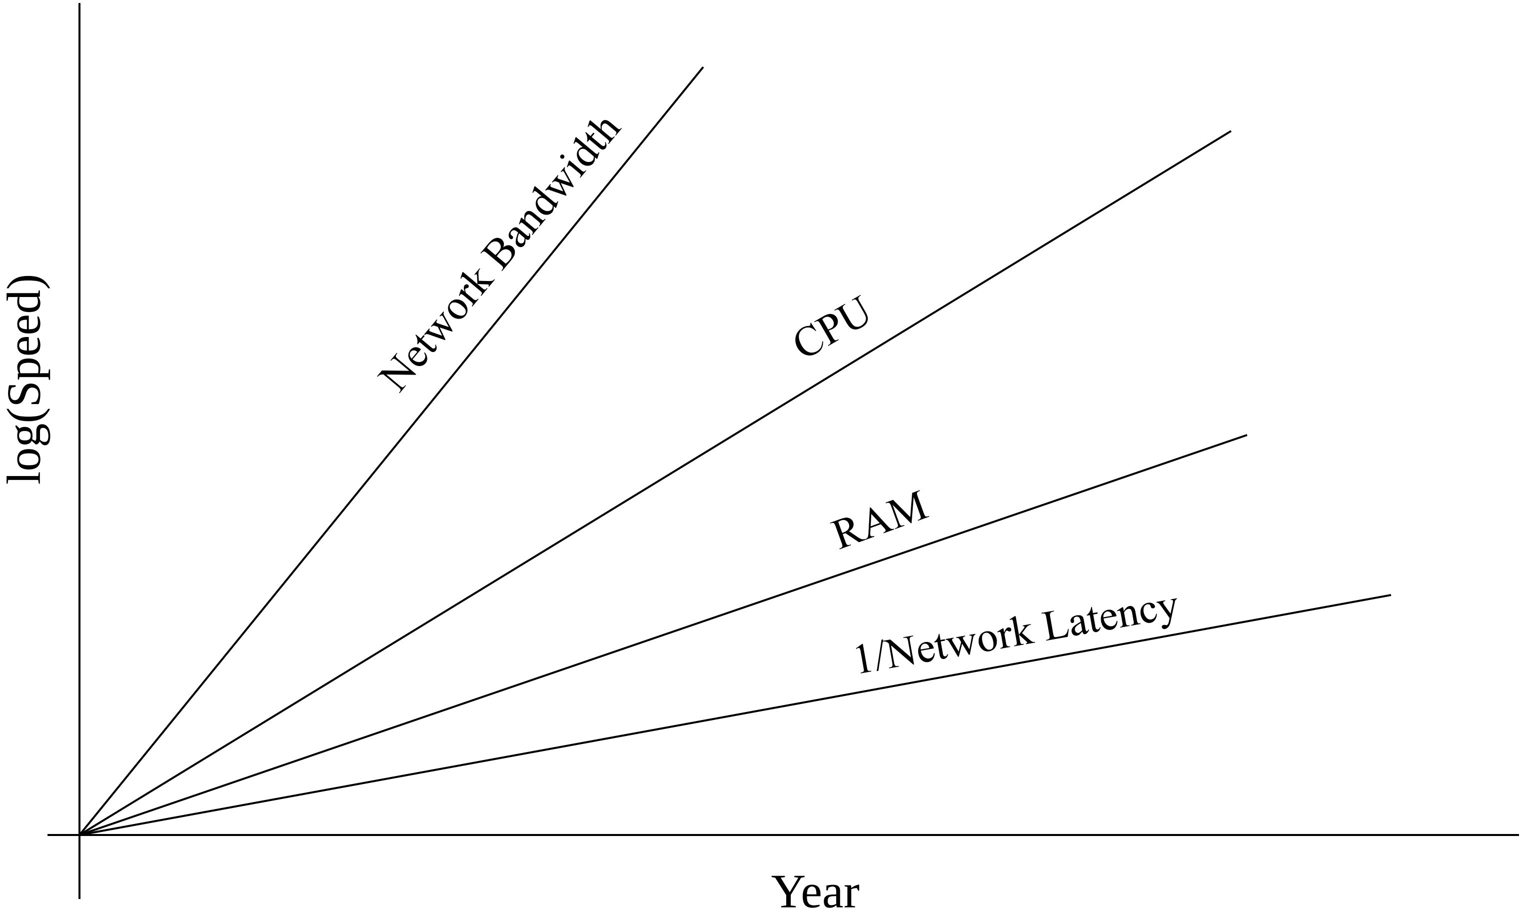

Moore's Law in Practice

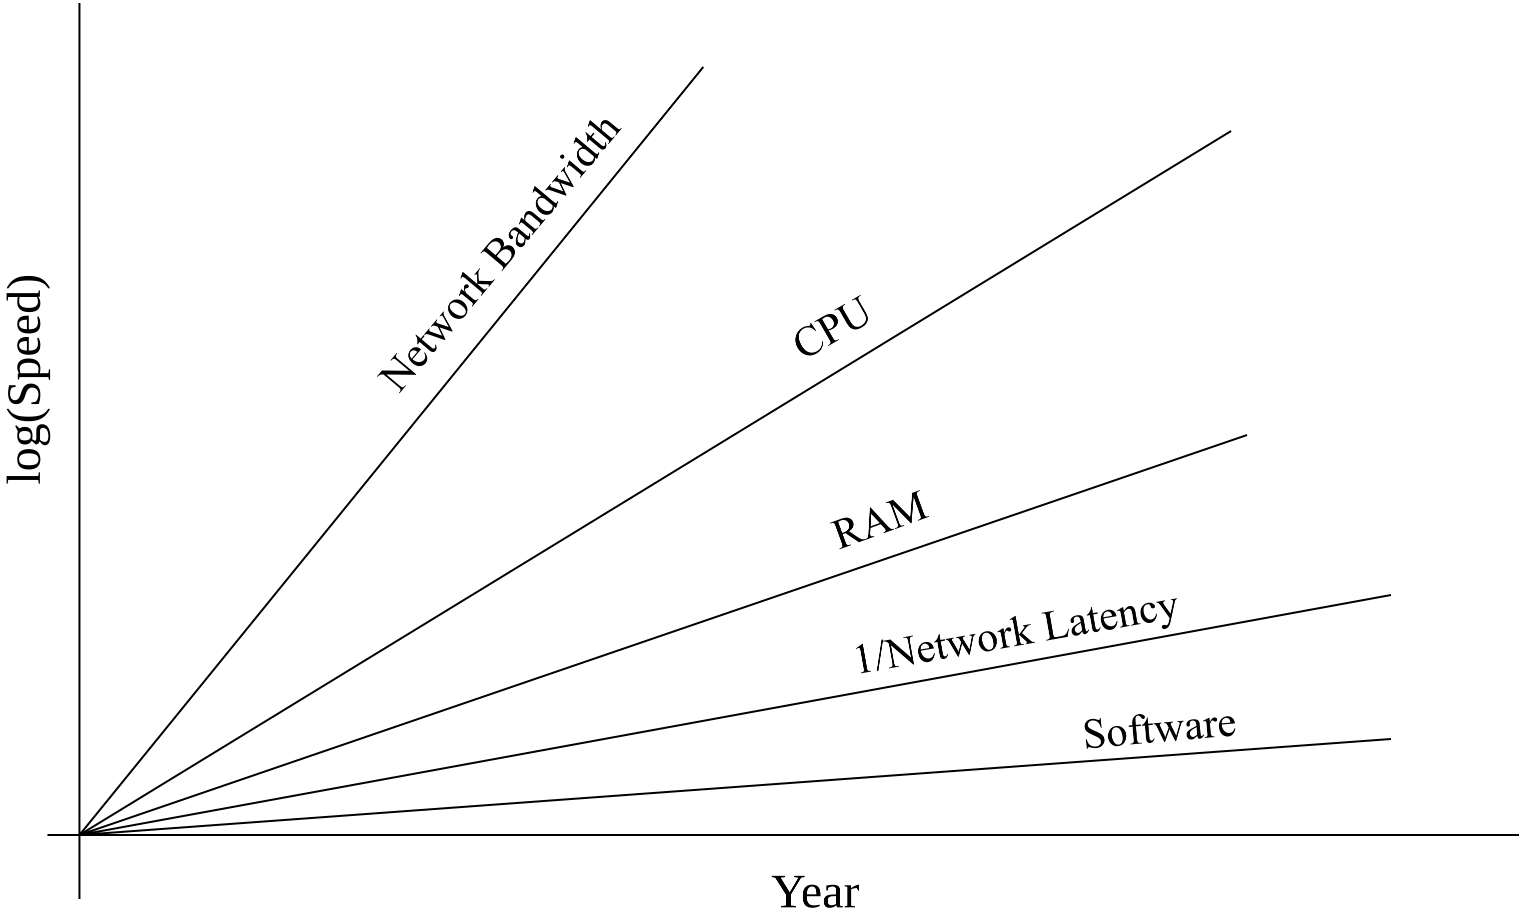

Moore's Law on Gene Sequencers

Increases 10x every 18 months, compared to 2x every 18 months for CPUs.

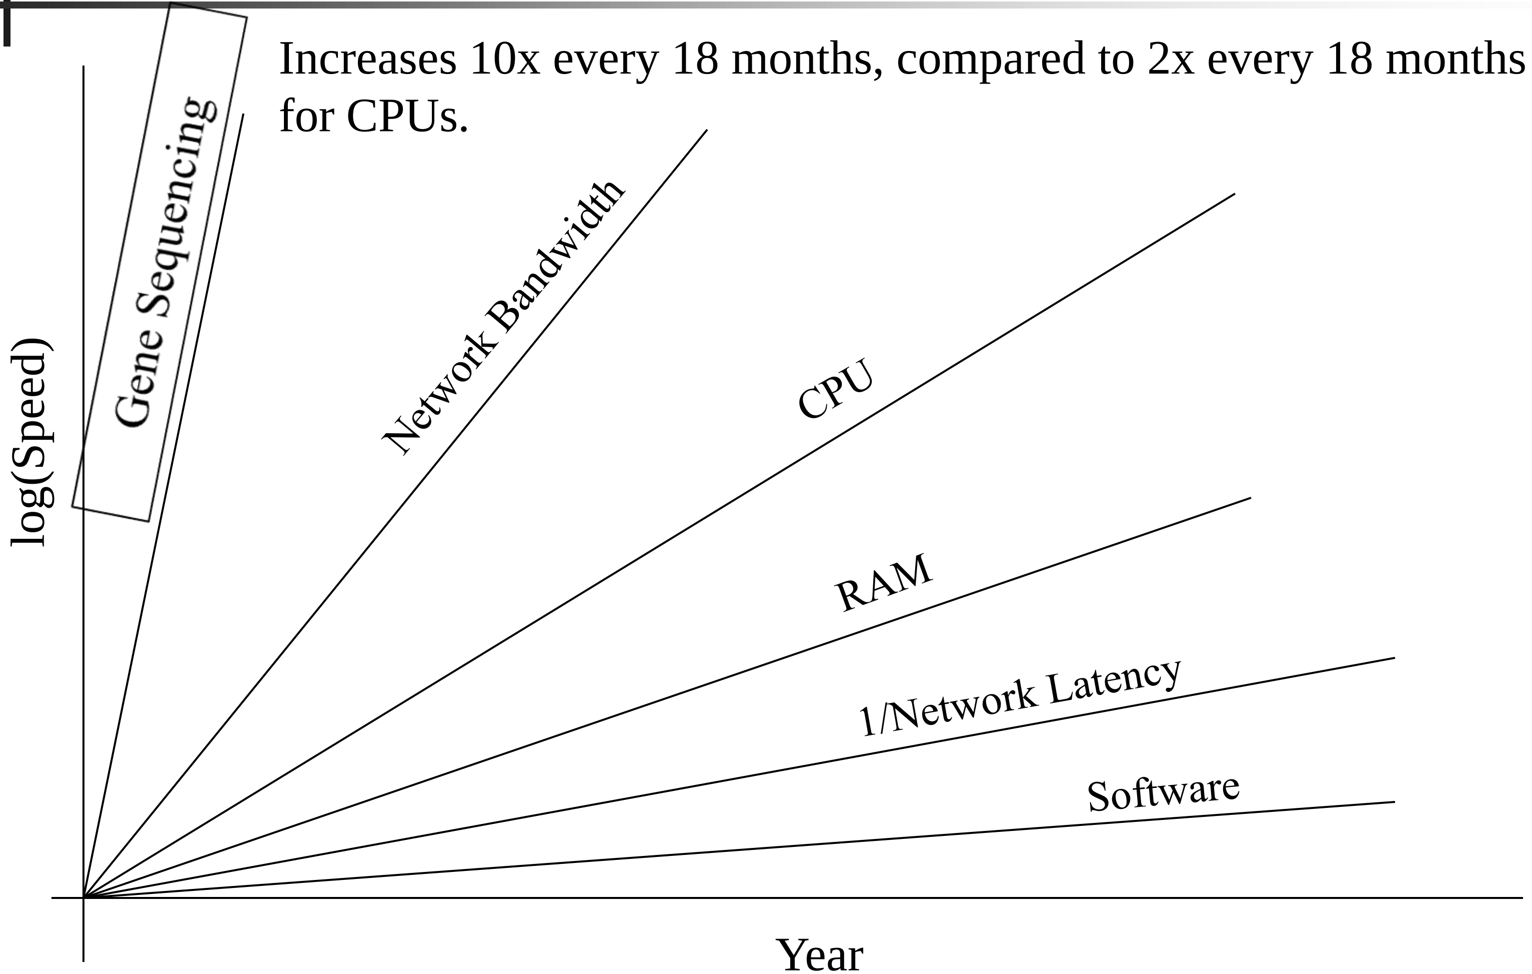

Why Bother?

Why Bother with HPC at All?

It's clear that making effective use of HPC takes quite a bit of effort, both learning how and developing software.

That seems like a lot of trouble to go to just to get your code to run faster.

It's nice to have a code that used to take a day, now run in an hour. But if you can afford to wait a day, what's the point of HPC?

Why go to all that trouble just to get your code to run faster?

Why HPC is Worth the Bother

- What HPC gives you that you won't get elsewhere is the ability to do **bigger, better, more exciting science**. If your code can run faster, that means that you can tackle much bigger problems in the same amount of time that you used to need for smaller problems.
- HPC is important not only for its own sake, but also because what happens in HPC today will be on your desktop in about 10 to 15 years: it puts you **ahead of the curve**.

The Future is Now

Historically, this has always been true:

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

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

OK Supercomputing Symposium 2012

2003 Keynote:
Peter Freeman
NSF
Computer & Information Science & Engineering Assistant Director

2004 Keynote:
Sangtae Kim
NSF Shared Cyberinfrastructure Division Director

2005 Keynote:
Walt Brooks
NASA Advanced Supercomputing Division Director

2006 Keynote:
Dan Atkins
Head of NSF's Office of Cyberinfrastructure

2007 Keynote:
Jay Boisseau
Director
Texas Advanced Computing Center
U. Texas Austin

2008 Keynote:
José Munoz
Deputy Office Director/ Senior Scientific Advisor
NSF Office of Cyberinfrastructure

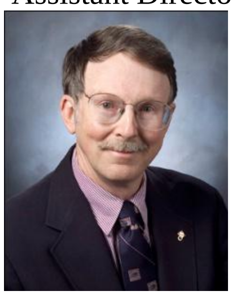

2009 Keynote:
Douglass Post
Chief Scientist
US Dept of Defense HPC Modernization Program

2010 Keynote:
Horst Simon
Deputy Director
Lawrence Berkeley National Laboratory

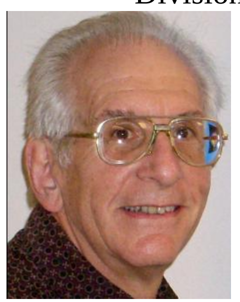

2011 Keynote:
Barry Schneider
Program Manager
National Science Foundation

Thom Dunning, Director
National Center for Supercomputing Applications

FREE! Wed Oct 3 2012 @ OU

http://symposium2012.oscer.ou.edu/

Reception/Poster Session
FREE! Tue Oct 2 2012 @ OU
FREE! Symposium Wed Oct 3 2012 @ OU

Thanks for your attention!

Questions?

www.oscer.ou.edu

References

[1] Image by Greg Bryan, Columbia U.

[2] "Update on the Collaborative Radar Acquisition Field Test (CRAFT): Planning for the Next Steps." Presented to NWS Headquarters August 30 2001.

[3] See http://hneeman.oscer.ou.edu/hamr.html for details.

[4] http://www.dell.com/

[5] http://www.vw.com/newbeetle/

[6] Richard Gerber, The Software Optimization Cookbook: High-performance Recipes for the Intel Architecture. Intel Press, 2002, pp. 161-168.

[7] RightMark Memory Analyzer. http://cpu.rightmark.org/

[8] ftp://download.intel.com/design/Pentium4/papers/24943801.pdf

[9] http://www.samsungssd.com/meetssd/techspecs

[10] http://www.samsung.com/Products/OpticalDiscDrive/SlimDrive/OpticalDiscDrive_SlimDrive_SN_S082D.asp?page=Specifications

[11] ftp://download.intel.com/design/Pentium4/manuals/24896606.pdf

[12] http://www.pricewatch.com/