The Potential of Cloud Computing: Challenges, Opportunities, Impact

Armando Fox, UC Berkeley

Adapted from NAE Gilbreth Lecture, October 2011

Image: John Curley http://flickr.com/photos/jay_que/1834540/

Get Your Own Supercomputer

- 41.8 Teraflops on LINPACK benchmark
  - #451 on current Top500 list of supercomputers
- No batch queues—start your job in minutes
- No grantwriting—just a credit card (from $1300/hour, 1 hour minimum purchase)
- No need to predict usage—pay for what you use, add capacity on demand
- Lease several simultaneously, to run multiple experiments in parallel
- Matlab, Mathematica, other packages available

Warehouse Scale Computers

- Built to support consumer demand of Web services (email, social networking, etc.)
  - “Private clouds” of Google, Microsoft, Amazon, ...
- “Warehouse scale” buying power = 5-7x cheaper hardware, networking, administration cost

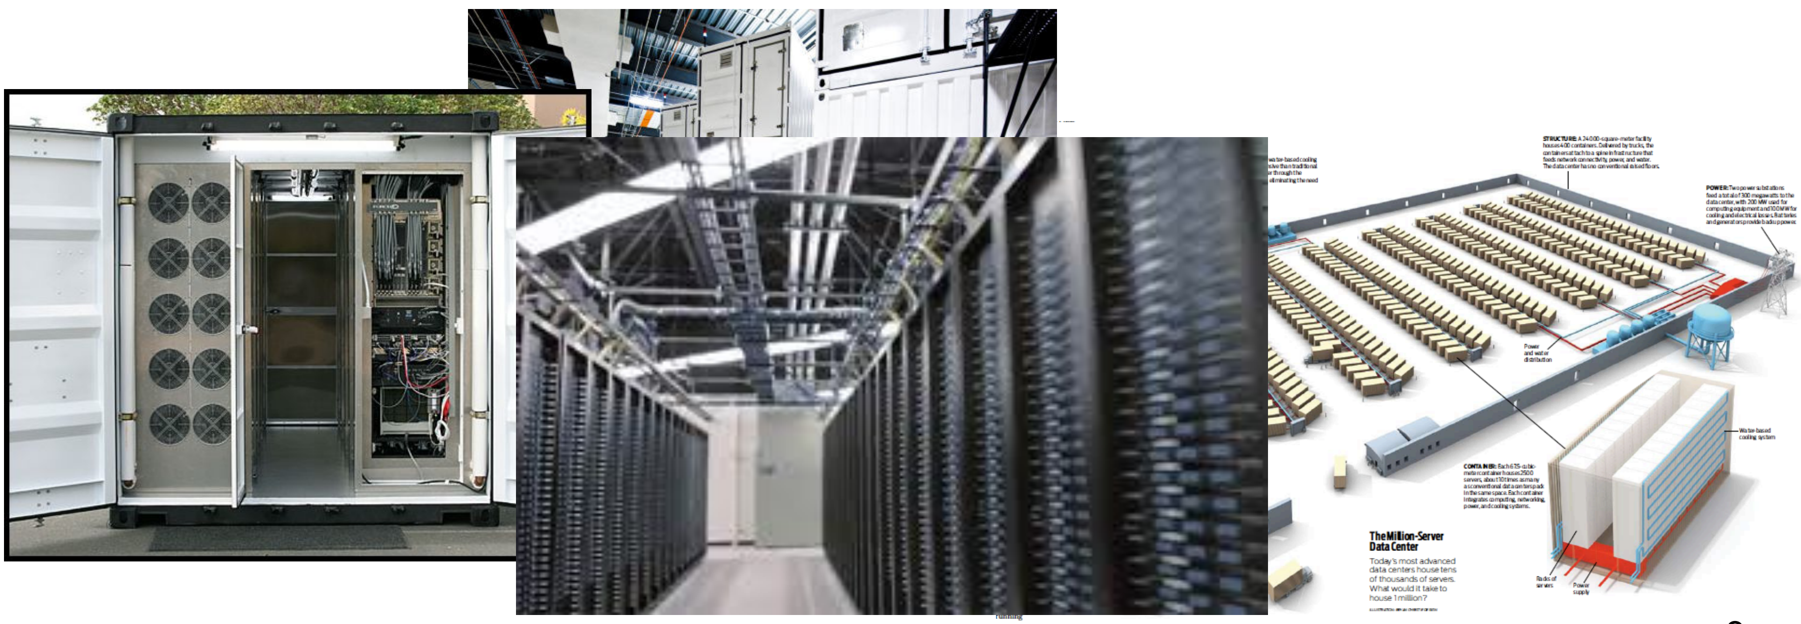

photos: Cnet News, Sun Microsystems (Oracle), datacenterknowledge.com

- The first demonstration of how to build really large Internet sites out of *clusters* of *commodity* computers was done by:

(a) Stanford

(b) Berkeley

(c) Yahoo!

(d) Google

(e) IBM

2008: Public Cloud Computing Arrives

Type & Price/Hour	1GHz core eqv.	RAM	Disk and/or I/O
Small - $0.085	1	1.7 GB	160 GB
Large - $0.34	4	7.5 GB	850 GB, 2 spindles
XLarge - $0.68	8	15 GB	1690 GB, 3 spindles
Cluster ("HPC") - $1.60	32*	23 GB	1690 GB + 10Gig Ethernet + optional 2x NVidia GPU

- Virtual machines from $0.085/hr to $2.10/hr
  - Pay as you go with credit card, 1 hr. minimum
  - Cheaper if willing to share or risk getting kicked off
  - Machines provisioned & booted in a few minutes

1,000 machines for 1 hr = 1 machine × 1,000 hrs

* 2x quad-core Intel Xeon (Nehalem)

What’s In It For You?

Cloud Computing progress 2008-2011

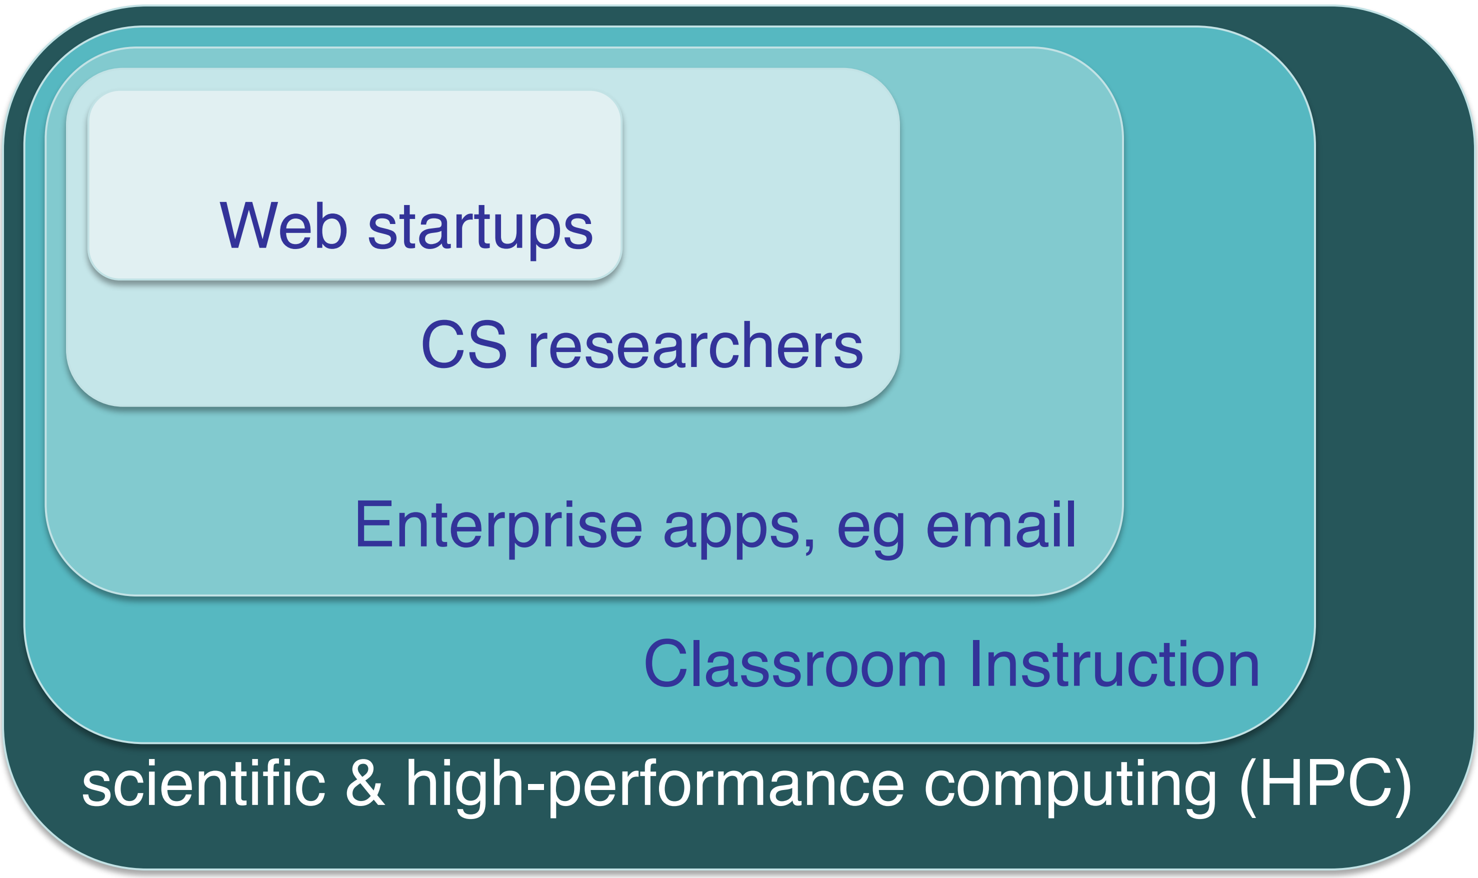

What are the challenges to adoption?

Trivia Fact

- The first full Web browser running on a mobile device was developed by:

(a) Apple
(b) Stanford
(c) Berkeley
(d) Nokia
(e) Motorola

Hiding the Cloud's Complexity

- Sophisticated software can ***hide complexity & mask unreliability*** of cloud hardware
- Example: **MapReduce** (2004)
  - e.g.: "**Apply** f(x) over all these x, **keep** K max values"
  - Google's MapReduce software automatically masks failures and stewards cloud resources
  - Hadoop (open source version) followed soon after

*"Warehouse scale" software engineering issues **hidden** from application programmer*

MapReduce in Practice

- Example: classifying Twitter spam
  - training: 10^7 samples x 64KB data each = 640GB data
  - One heavy-duty server: ~270 hours
  - 100 servers in cloud: ~3 hours (~$250)
- Rapid uptake in other scientific research
  - "top 5 pharma" molecular modeling: 3809 machines on Amazon cloud, **30K** cores, $1279/hr* (Sep 2011)
  - Large-population genetic risk analysis & simulation (Harvard Medical School)
  - Genome sequencing (UNC Chapel Hill Cancer Ctr)
  - Compact Muon Solenoid Expt. (U. Nebraska Lincoln)
- What's the downside?

* ArsTechnica, 9/20/11, **http://arst.ch/qy8**

Challenge: Cloud Programming

- Programmers must (re)express problems to ***expose parallelism***

⚠ parallel software hard to debug & operate, so build on existing successes

  – Pig (Yahoo! Research) & Hive (Apache Foundation) – transform database queries to MapReduce

  – Rhipe—include MapReduce operations in R programs

⚠ one abstraction doesn't fit all

  – e.g., clustering algorithms work poorly on MapReduce

  – **Opportunity**: common front-end to cloud & parallel programming tools (`sejits.eecs.berkeley.edu`)

Challenge: "non-cloudy" scientific codes

- ⚠ Existing scientific codes "supercomputer-centric"
  - reliability, static configuration, exclusive resource use...
- Time-to-answer may still be faster, since no wait!
- **Opportunity**: software that shares cloud among multiple frameworks (`mesos.berkeley.edu`)
- **Opportunity**: Vendors listening to HPC customers
  - Hardware: cloud-based "supercomputer" makes Top500 list (July 2010)
  - Software: MathWorks, Mathematica now support sending (parallel) computations to Amazon cloud

Challenge: Big Data

Application	Data generated per day* (1 TB = 10^{12} bytes)
Genome Sequencing	1 TB (Illumina HiSeq machine) Cancer Genome Atlas: 5PB online by end of 2011
Large Synoptic Survey Telescope	30 TB 400 Mbps sustained xfer between Chile & NCSA
Large Hadron Collider	60 TB

⚠ **How to copy 8 TB (8×10^{12} bytes) from cloud?**

- 2009: Download at ~20 Mbps ➔ 35 days, $800
- 2010: Ship drive to Amazon ➔ 1 day, $150
- 2011: Pay-as-you-go fast network ➔ <1 day, $165
    - (~18 hours at 1 Gbps, 0.30/hr. + 0.02/GByte)

* Source: Ed Lazowska, eScience 2010, Microsoft Cloud Futures Workshop, lazowska.cs.washington.edu/cloud2010.pdf

Algorithms, Machines, People

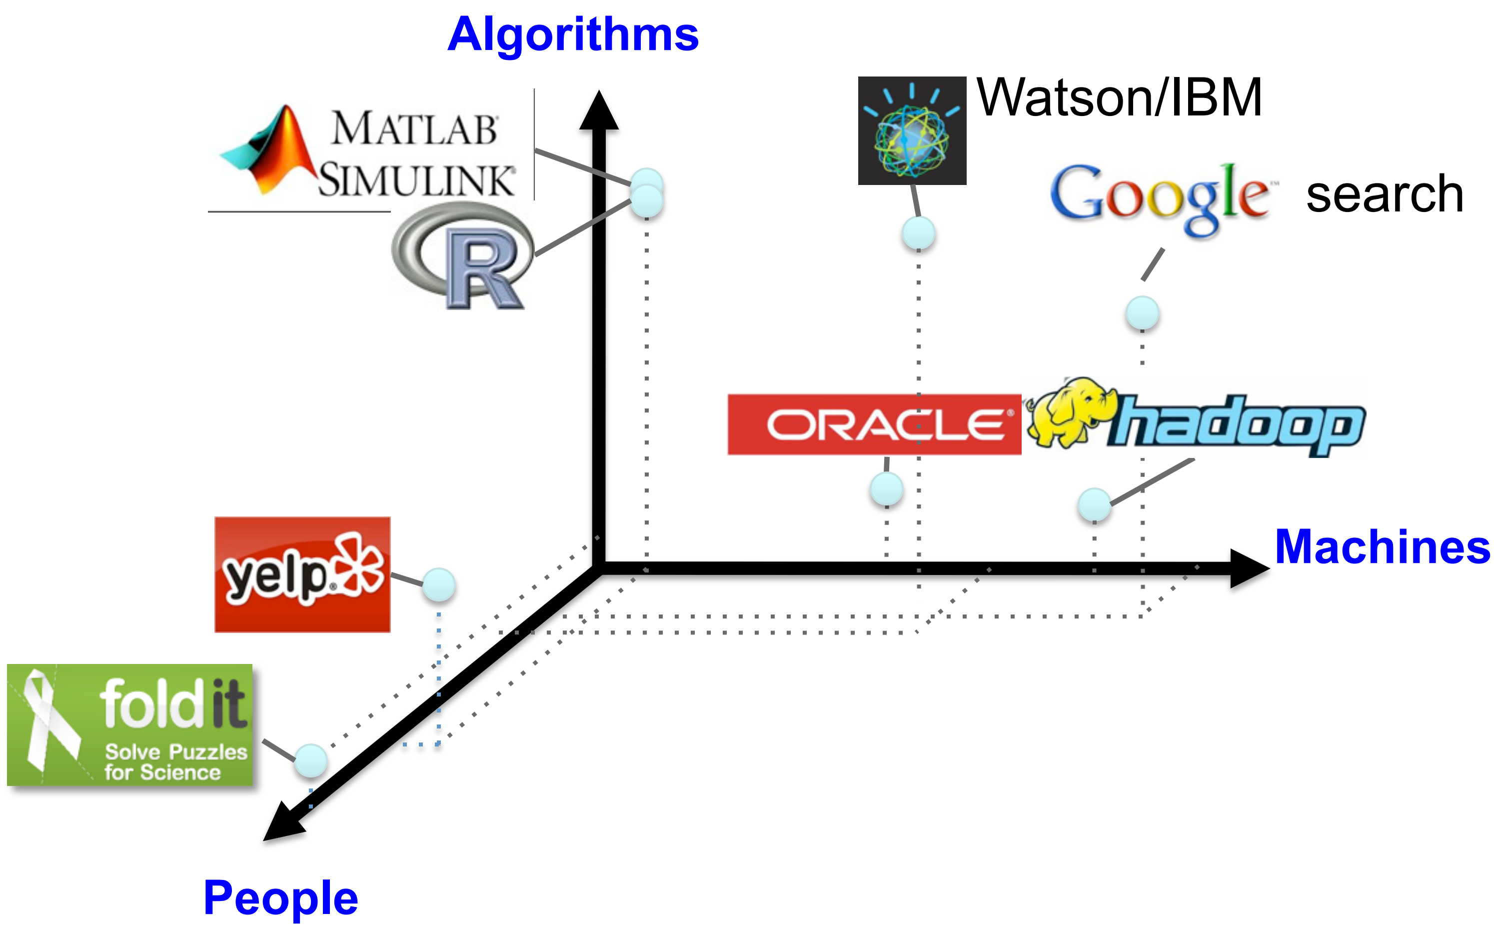

The Berkeley AMP Lab

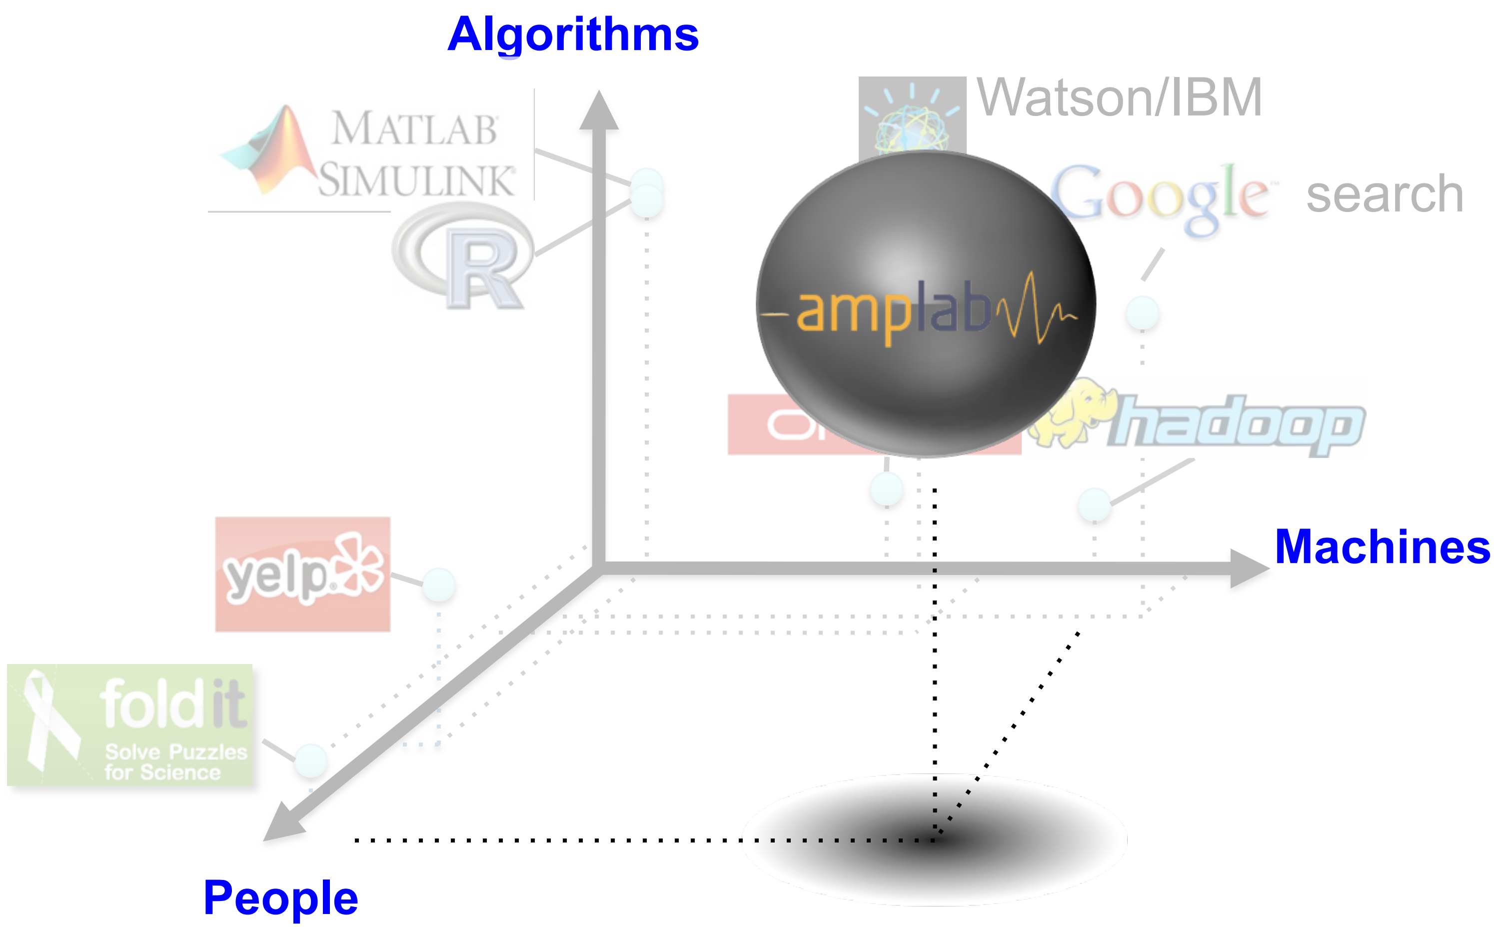

Trivia fact

- The percentage of Americans who will face cancer in their lifetimes is approximately:

(a) 1%
(b) 5%
(c) 10%
(d) 20%
(e) 30%

Trivia fact

- By next year, the "wet lab" cost of sequencing a human genome is expected to be about:

(a) $1,000,000
(b) $100,000
(c) $10,000
(d) $1,000
(e) None of the above

Opportunity: Application-Driven CS research

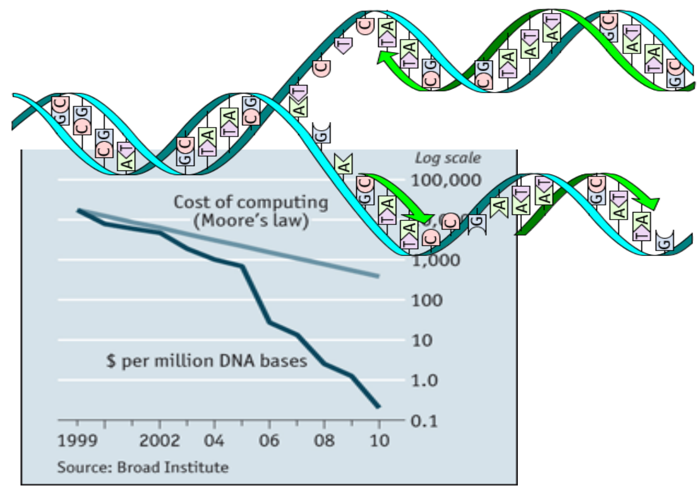

By working with domain experts, we will create better tools to support big-data applications on cloud computing

Cloud & UCB EECS

1. Berkeley research culture: integrate leading research into teaching at all levels
2. RAD Lab need for "killer apps" to show off infrastructure

Current efforts (student counts approximate):

- Great Ideas in Computer Architecture (reinvented Fall 2010): 190 students
- Software Engineering for SaaS (Spr' 12 = 5th iteration): 50+50+50+70+100 students
- Operating Systems: 70 students
- Intro. Data Science: 30
- Adv. topics in HCI: 20 students
- Natural language processing: 20 students
- Cancer Genomics: 10 students

An Analogy: *Networks of Workstations (NOW)*, c.1995

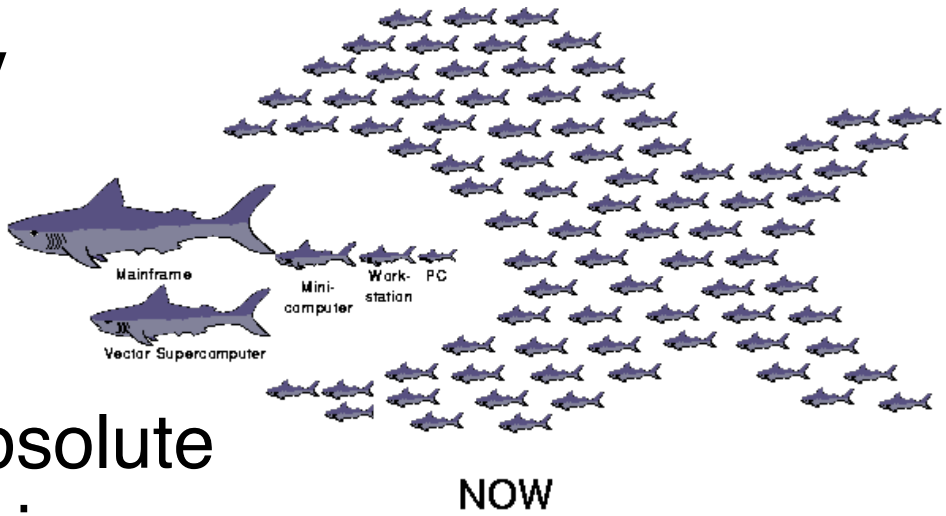

- Clusters of Commodity PC's vs. symmetric multiprocessors
- *Potential* advantages: incremental scaling, absolute capacity, commodity price
- 1995: Web software architecture was SMP-centric
- 2010: Berkeley undergrads prototype Web apps in 6–8 weeks and deploy on cloud computing
- 2015: scientists routinely do 1000-machine e-science experiments using cloud computing

Conclusion

- Democratization of supercomputing capability
  - Time to answer may be faster even if hardware isn't
  - Writing a grant proposal around a supercomputer?
- Software & hardware for science-on-cloud improving at "commodity speed"
- Part of Computer Science "coming of age" with increasingly outward-looking research

Potential democratizing impact comparable to the microprocessor

Acknowledgments

- Colleagues at UC Berkeley RAD Lab
  - Co-authors of *Above The Clouds:* Michael Armbrust, Rean Griffith, Anthony D. Joseph, Randy Katz, Andy Konwinski, Gunho Lee, David Patterson, Ariel Rabkin, Ion Stoica, Matei Zaharia
  - And everyone else in the RAD Lab & AMP Lab
- RAD Lab industrial supporters & collaborators

Thank you!