Soft Computing and Current Trends in Supercomputing

Horst D. Simon

Director, NERSC Center Division, LBNL

Berkeley, December 16, 2003

http://www.nersc.gov/~simon

Overview

- Introducing NERSC
- Current trends in supercomputing (high end computing)
- Computational science today
- Implications for soft computing

NERSC Center Overview

- Funded by DOE, annual budget $28M, about 65 staff
  - Traditional strategy to invest equally in newest compute platform, staff, and other resources
- Supports open, unclassified, basic research
- Close collaborations between university and NERSC in computer science and computational science

Flagship facility for unclassified computing in the DOE Office of Science

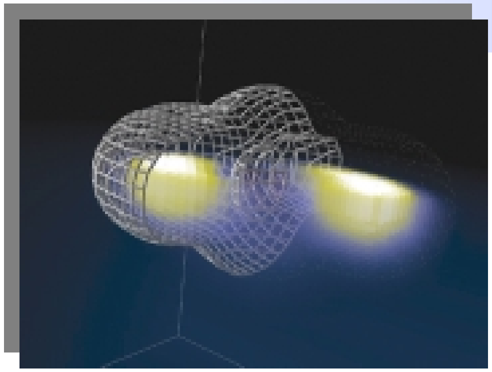

- computational projects in all areas of DOE science
- about 2000 users in ~200 projects

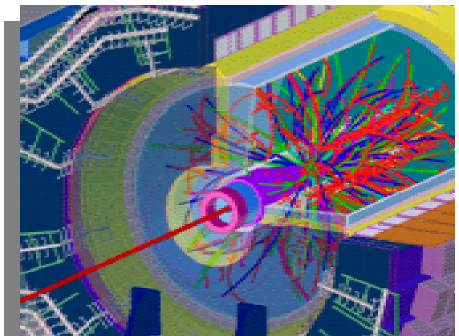

- focus on large-scale computing

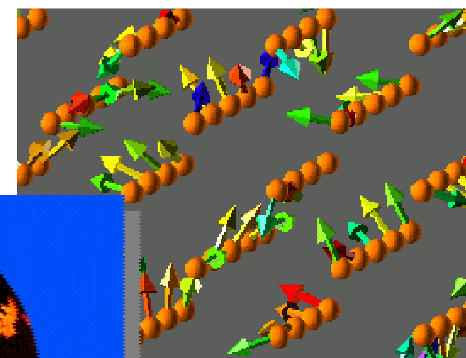

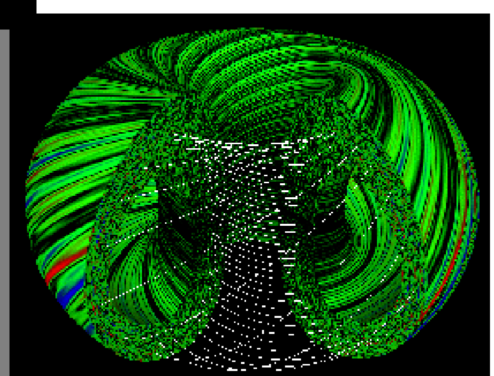

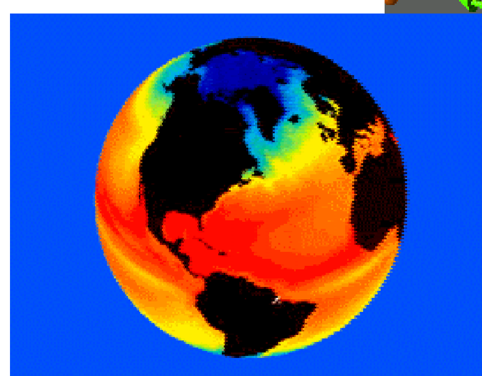

Overview

- Introducing NERSC
- Current trends in supercomputing (high end computing)
- Computational science today
- Implications for soft computing

Technology Trends: Microprocessor Capacity

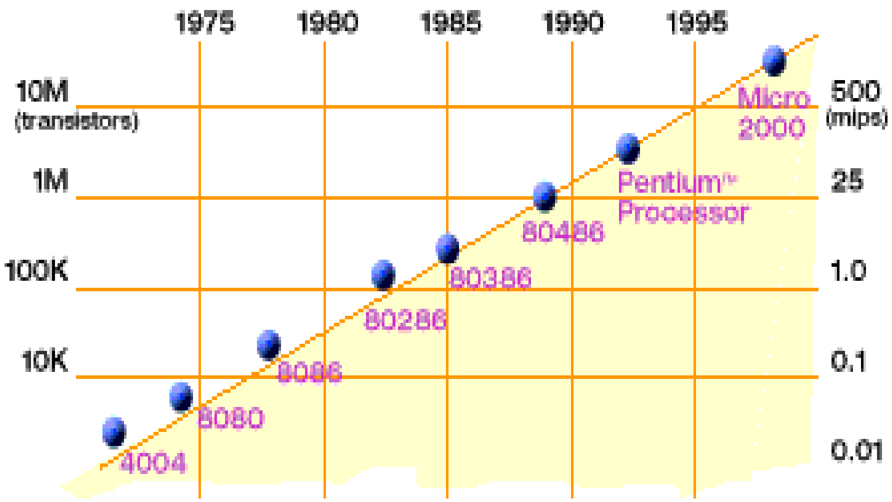

2X transistors/Chip Every 1.5 years

Called "Moore's Law"

Microprocessors have become smaller, denser, and more powerful.

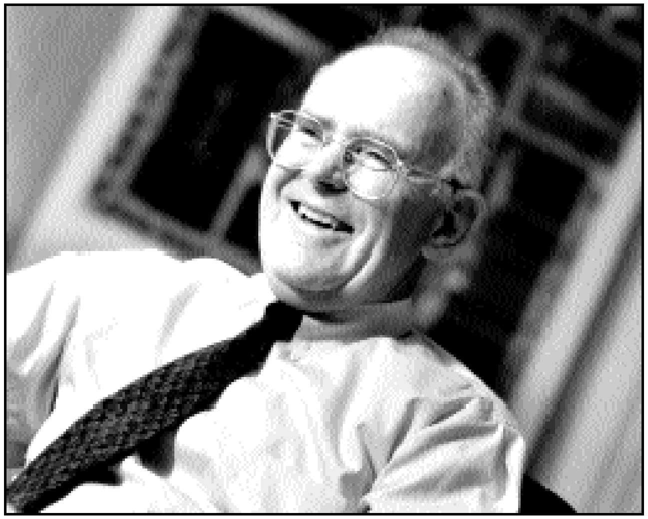

Gordon Moore (co-founder of Intel) predicted in 1965 that the transistor density of semiconductor chips would double roughly every 18 months.

Slide source: Jack Dongarra

TOP500 - Overview

- Listing of the 500 most powerful Computers in the World
- Yardstick: Rmax from Linpack

  Ax=b, dense problem

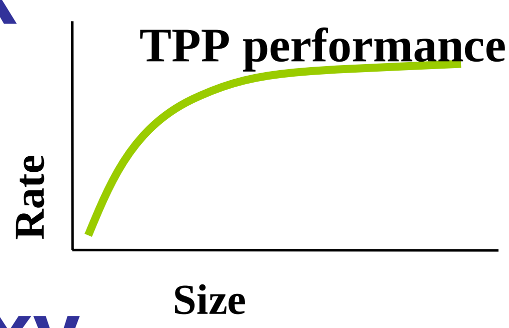

- Updated twice a year:
  - ISC'xy in Germany, June xy
  - SC'xy in USA, November xy
- All data available at www.top500.org

22nd List: The TOP10

Rank	Manufacturer	Computer	R_{max} [TF/s]	Installation Site	Country	Year	Area of Installation	# Proc
1	NEC	Earth-Simulator	35.86	Earth Simulator Center	Japan	2002	Research	5120
2	HP	*ASCI Q* AlphaServer SC	13.88	Los Alamos National Laboratory	USA	2002	Research	8192
3	Self-Made	*X* Apple G5, Mellanox	10.28	Virginia Tech	USA	2003	Academic	2200
4	Dell	*Tungsten* PowerEdge, Myrinet	9.82	NCSA	USA	2003	Academic	2500
5	HP	*Mpp2*, Integrity rx2600 Itanium2, Qadrics	8.63	Pacific Northwest National Laboratory	USA	2003	Research	1936
6	Linux Networx	*Lightning*, Opteron, Myrinet	8.05	Los Alamos National Laboratory	USA	2003	Research	2816
7	Linux Networx/ Quadrics	MCR Cluster	7.63	Lawrence Livermore National Laboratory	USA	2002	Research	2304
8	IBM	*ASCI White* SP Power3	7.3	Lawrence Livermore National Laboratory	USA	2000	Research	8192
9	IBM	*Seaborg* SP Power 3	7.3	NERSC Lawrence Berkeley Nat. Lab.	USA	2002	Research	6656
10	IBM/Quadrics	xSeries Cluster Xeon 2.4 GHz	6.59	Lawrence Livermore National Laboratory	USA	2003	Research	1920

Performance Development

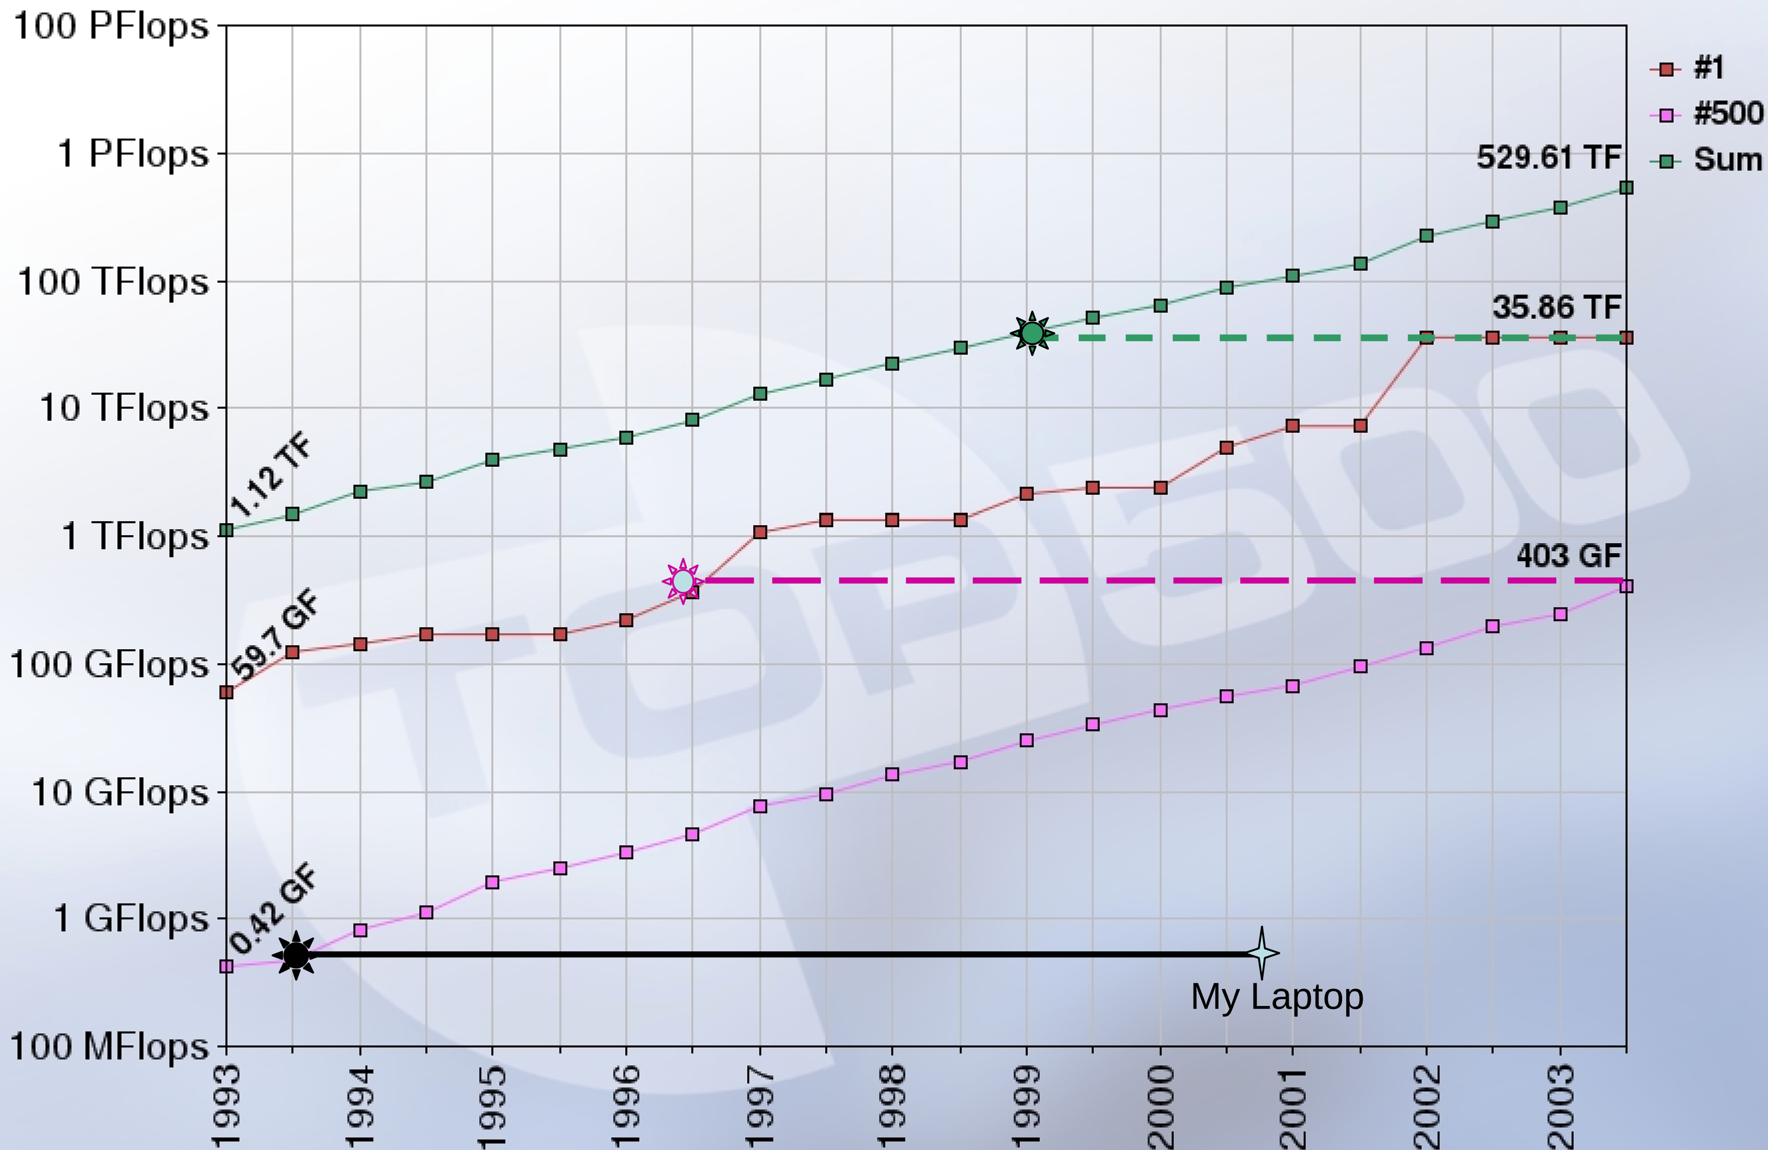

Projected Performance Development

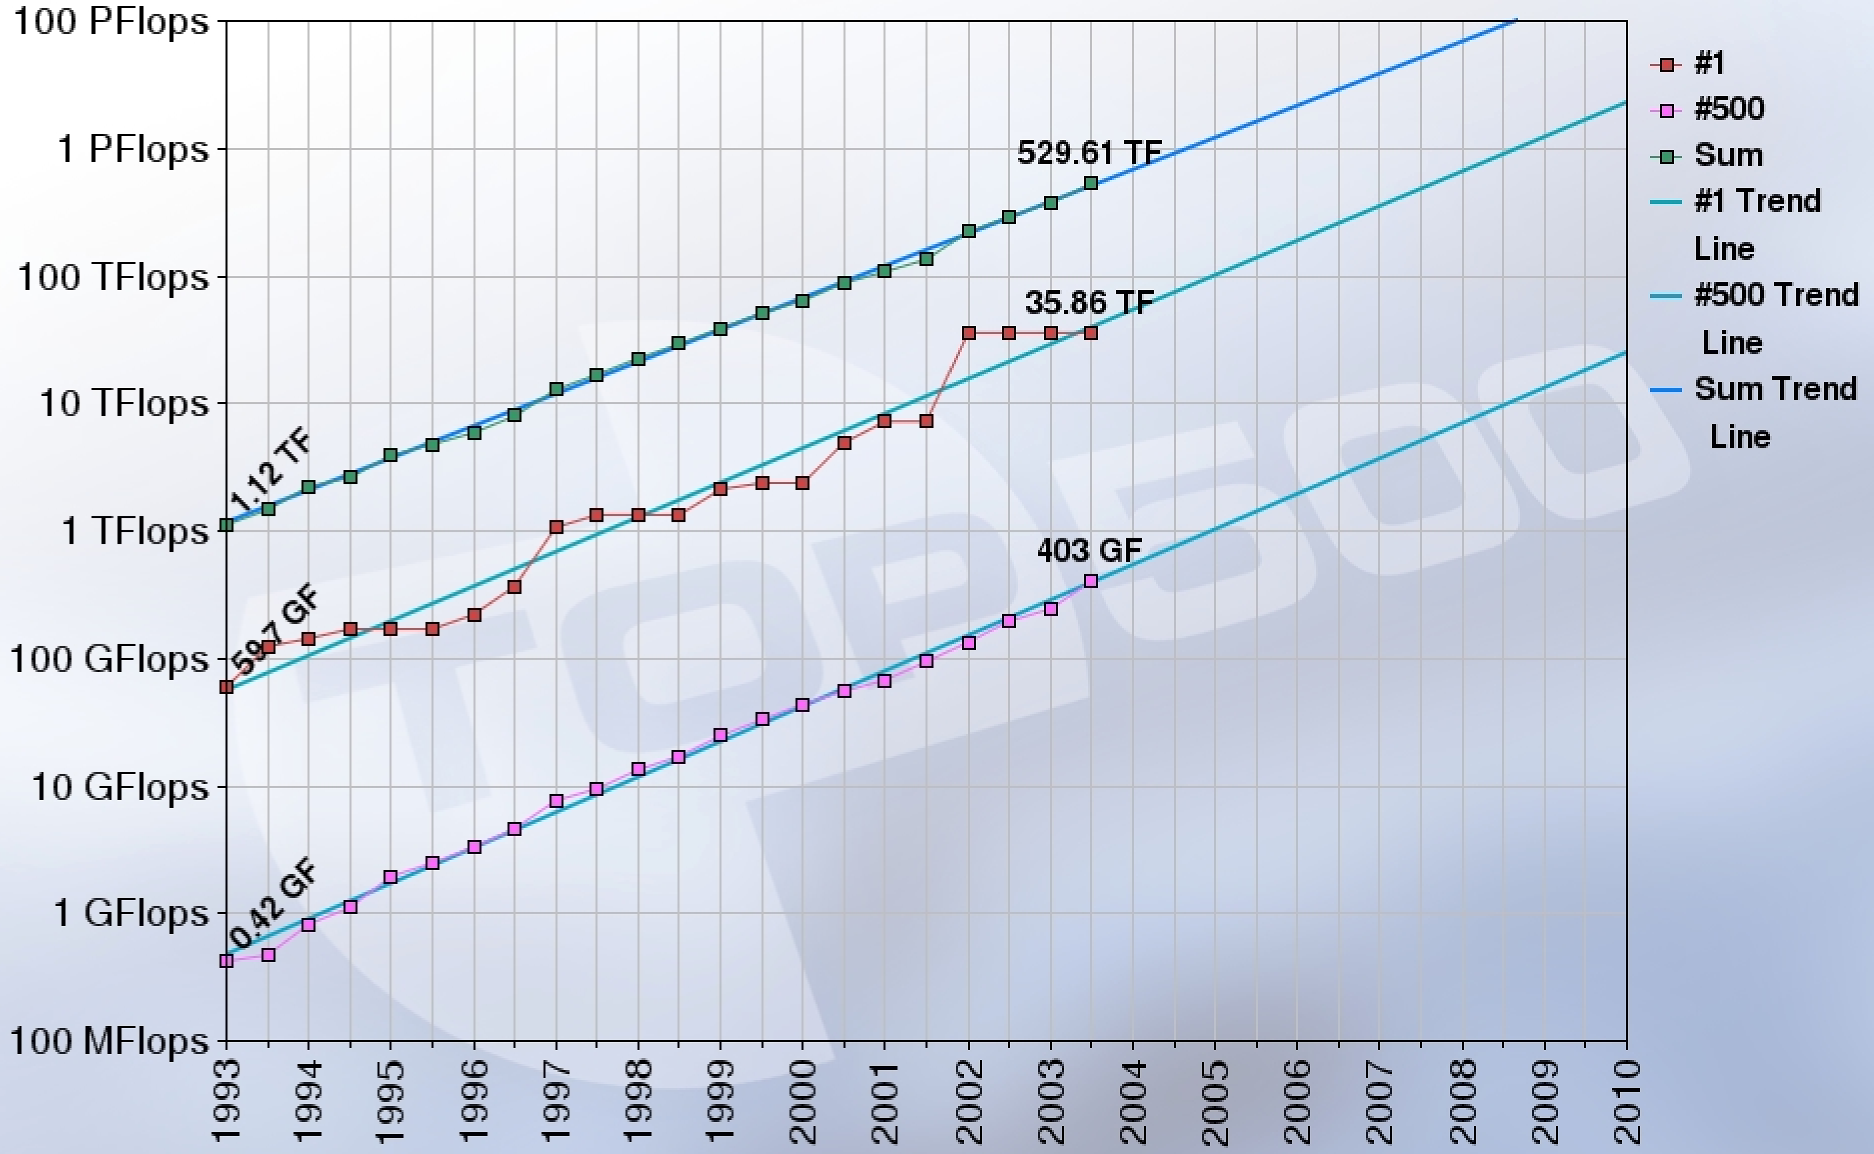

Analysis of TOP500 Data

- Annual performance growth about a factor of 1.82
- Two factors contribute almost equally to the annual total performance growth
- Processor number grows per year on the average by a factor of 1.30 and the
- Processor performance grows by 1.40 compared 1.58 of Moore's Law.
- For more details see paper by Dongarra, Meuer, Simon, and Strohmaier in Parallel Computing

Opportunity: Dramatic Advances in Computing
Terascale Today, Petascale Tomorrow

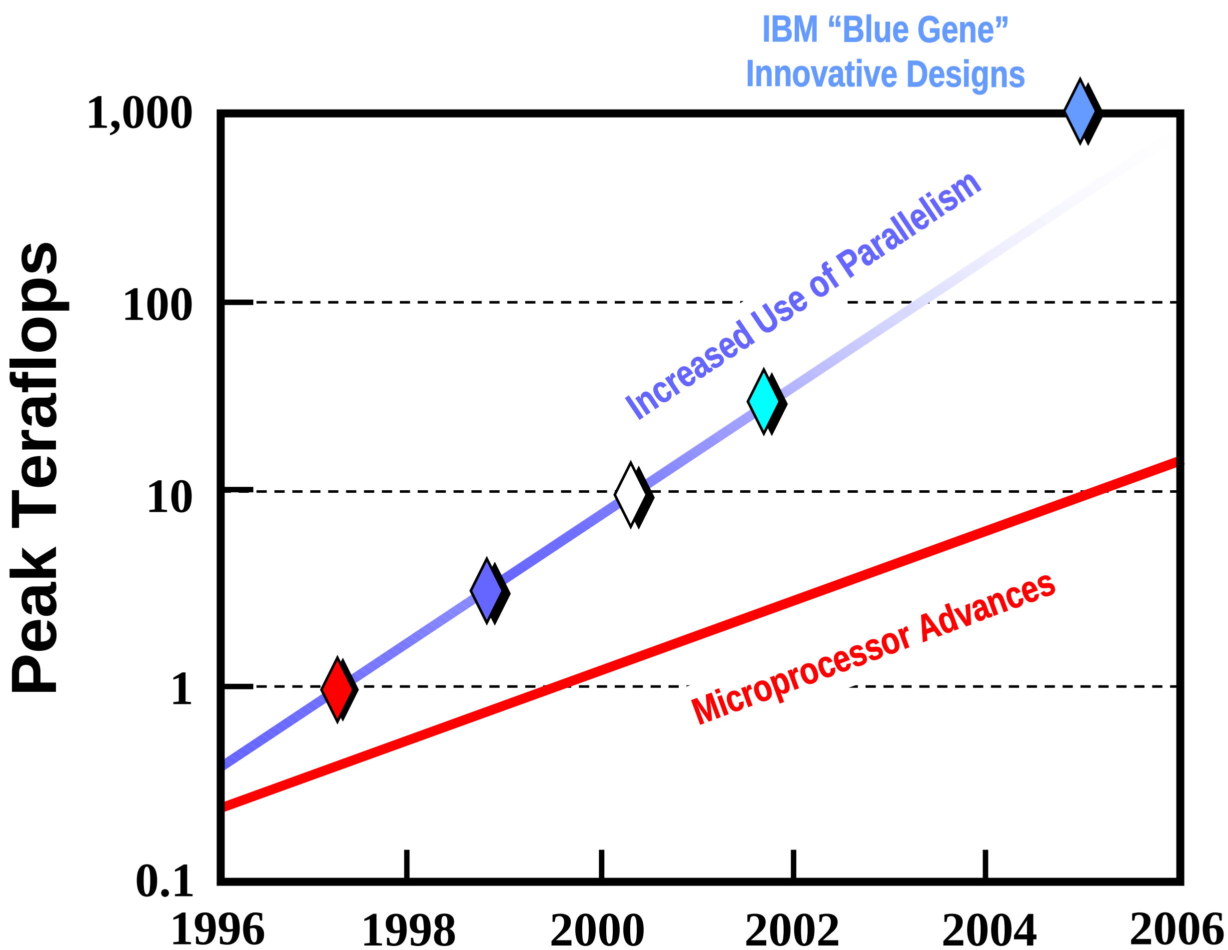

MICROPROCESSORS

2x increase in microprocessor speeds every 18-24 months ("Moore's Law")

PARALLELISM

More and more processors being used on single problem

INNOVATIVE DESIGNS

Processor in memory
Blue Gene

Office of Science
U.S. DEPARTMENT OF ENERGY

Manufacturer / Systems

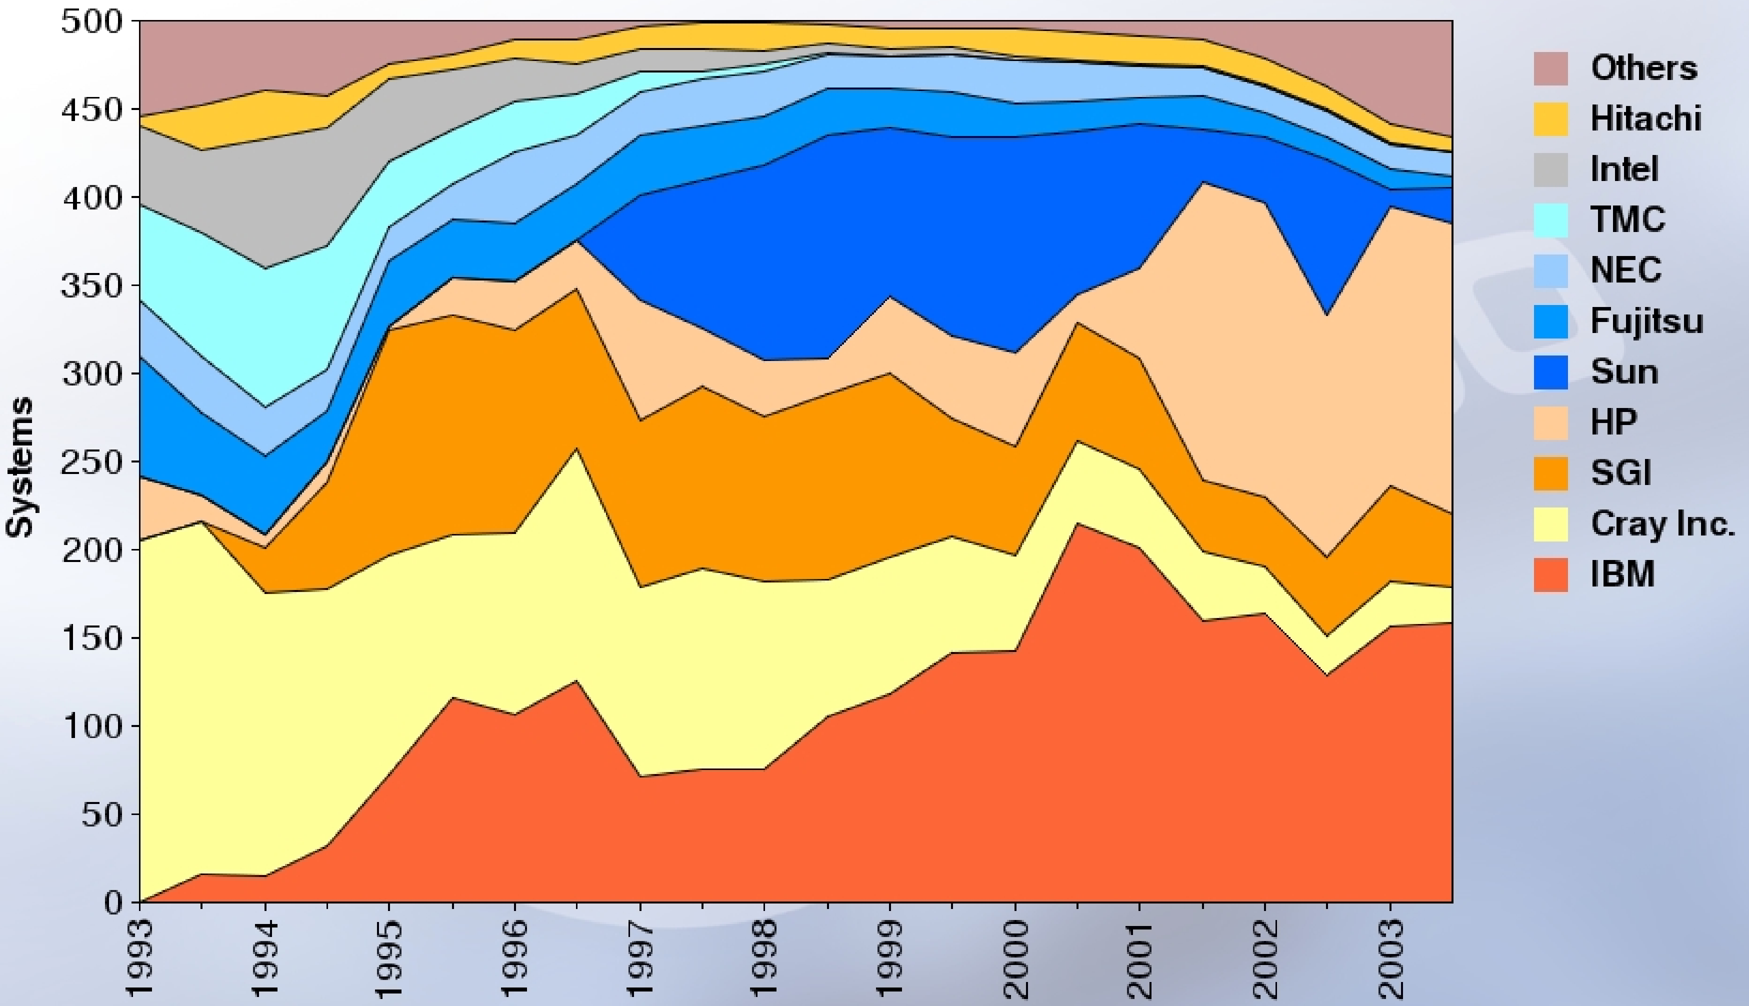

Architectures / Systems

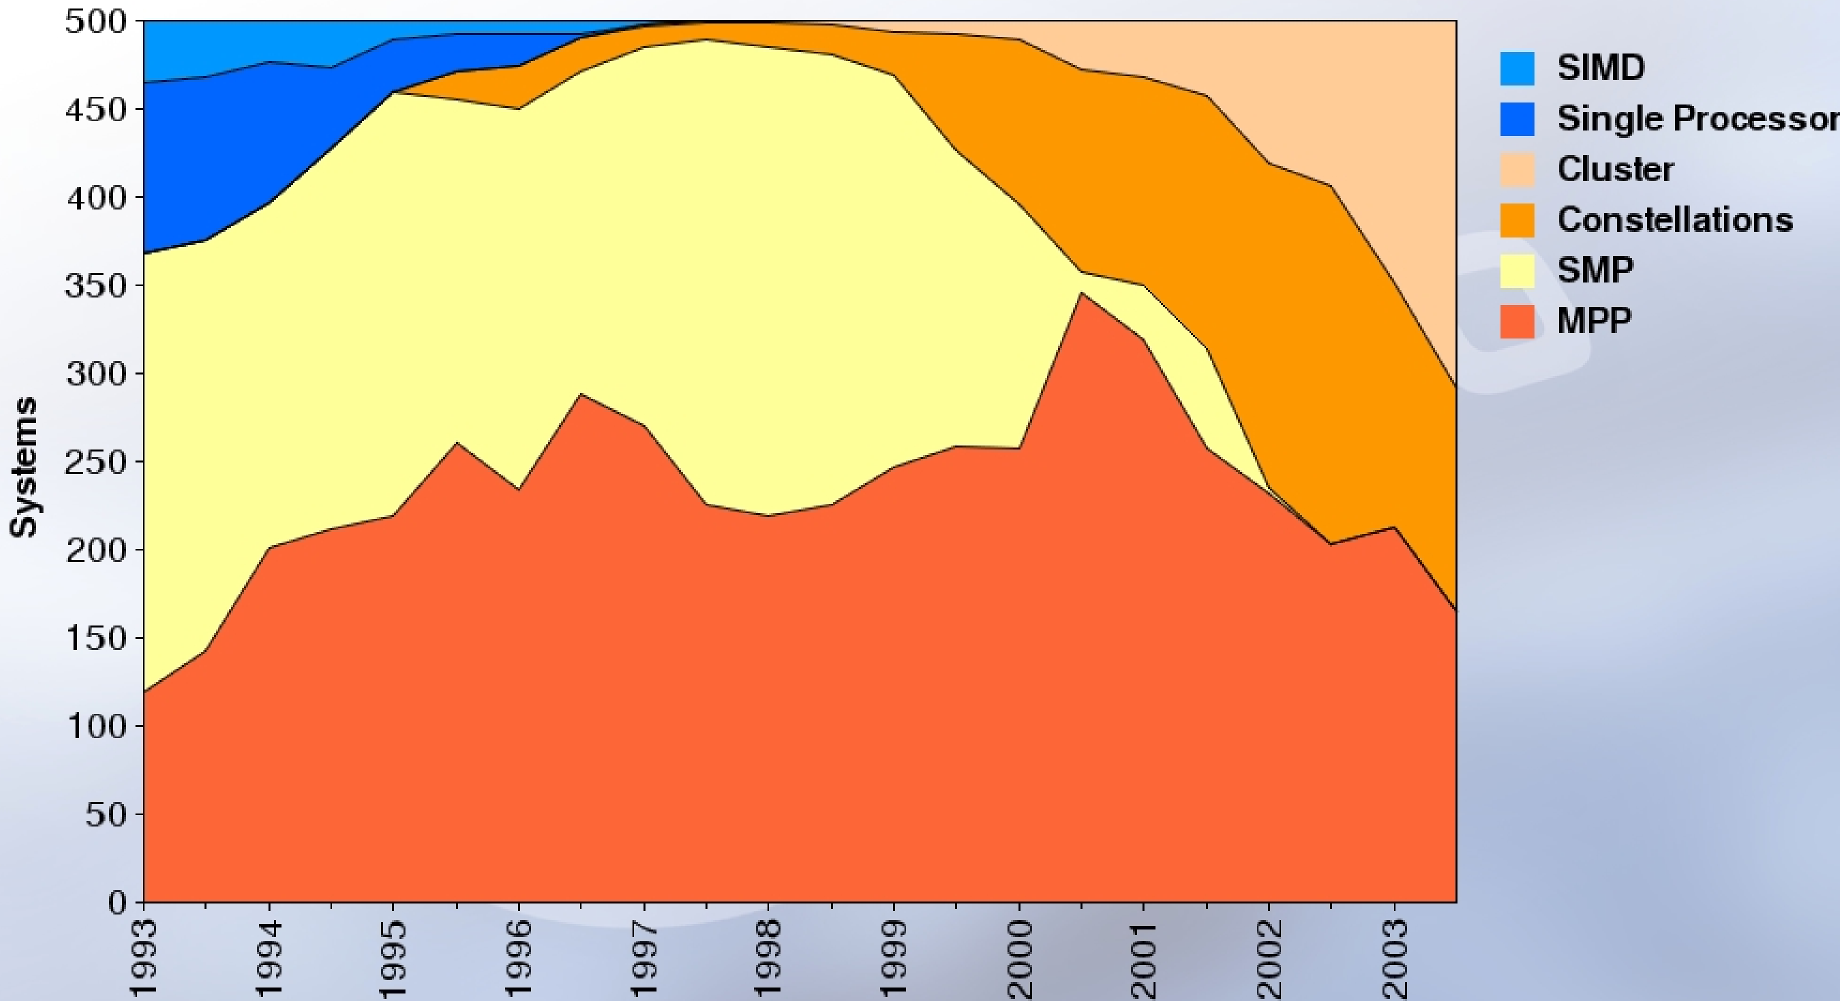

Clusters (NOW) / Systems

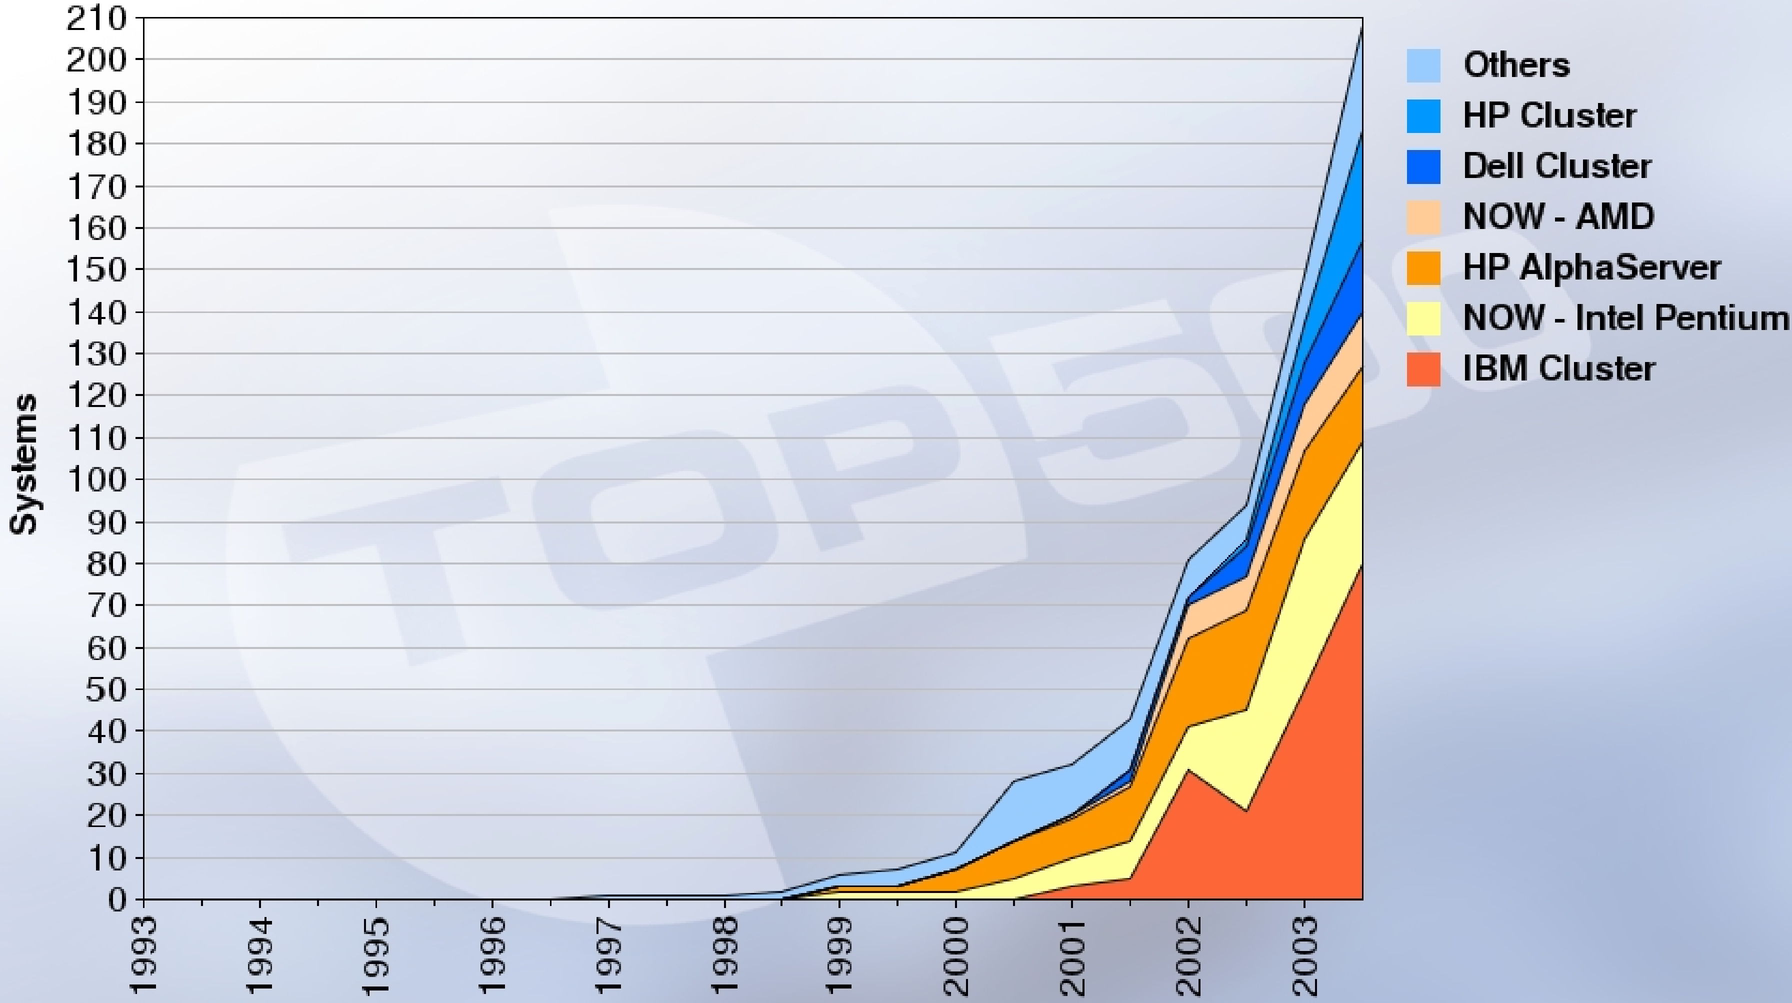

The Big Architecture Debate

- Commodity Clusters

vs

- Custom architecture

22nd List: The TOP10

Rank	Manufacturer	Computer	R_{max} [TF/s]	Installation Site	Country	Year	Area of Installation	# Proc
1	NEC	Earth-Simulator	35.86	Earth Simulator Center	Japan	2002	Research	5120
2	HP	*ASCI Q* AlphaServer SC	13.88	Los Alamos National Laboratory	USA	2002	Research	8192
3	Self-Made	*X* Apple G5, Mellanox	10.28	Virginia Tech	USA	2003	Academic	2200
4	Dell	*Tungsten* PowerEdge, Myrinet	9.82	NCSA	USA	2003	Academic	2500
5	HP	*Mpp2*, Integrity rx2600 Itanium2, Qadrics	8.63	Pacific Northwest National Laboratory	USA	2003	Research	1936
6	Linux Networx	*Lightning*, Opteron, Myrinet	8.05	Los Alamos National Laboratory	USA	2003	Research	2816
7	Linux Networx/ Quadrics	MCR Cluster	7.63	Lawrence Livermore National Laboratory	USA	2002	Research	2304
8	IBM	*ASCI White* SP Power3	7.3	Lawrence Livermore National Laboratory	USA	2000	Research	8192
9	IBM	*Seaborg* SP Power 3	7.3	NERSC Lawrence Berkeley Nat. Lab.	USA	2002	Research	6656
10	IBM/Quadrics	xSeries Cluster Xeon 2.4 GHz	6.59	Lawrence Livermore National Laboratory	USA	2003	Research	1920

Earth Simulator

- 40 Tflop/s system in Japan
- completion April 2002
- driven by climate and earthquake simulation requirements
- built by NEC
- 640 CMOS vector nodes

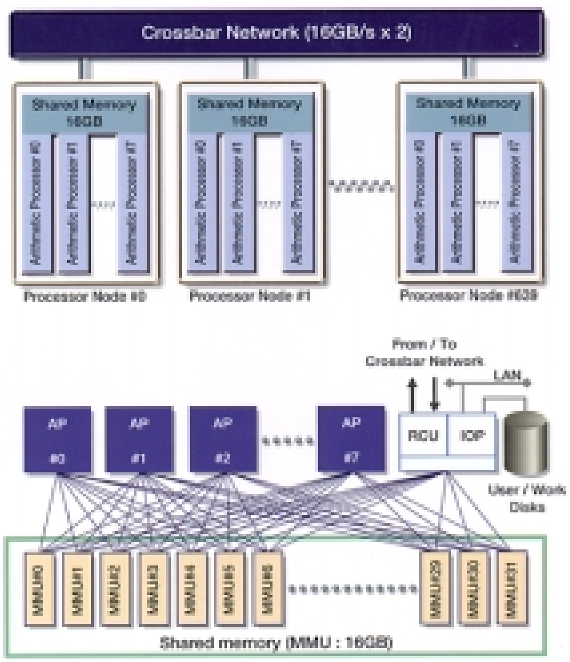

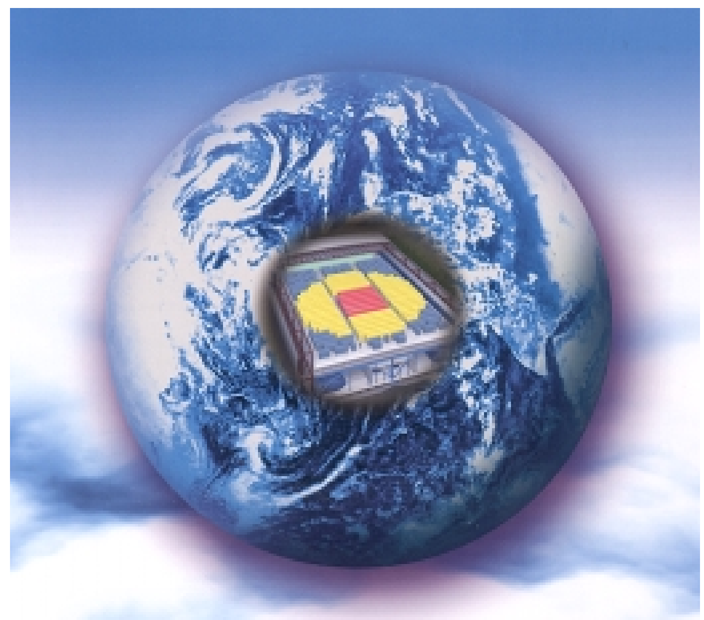

The Earth Simulator in Japan

COMPUTENIK!

- Linpack benchn[…] TF/s = 87% of 4[…]
- Completed Apr[…]
- Driven by clima[…]d earthquake simulation
- Built by NEC

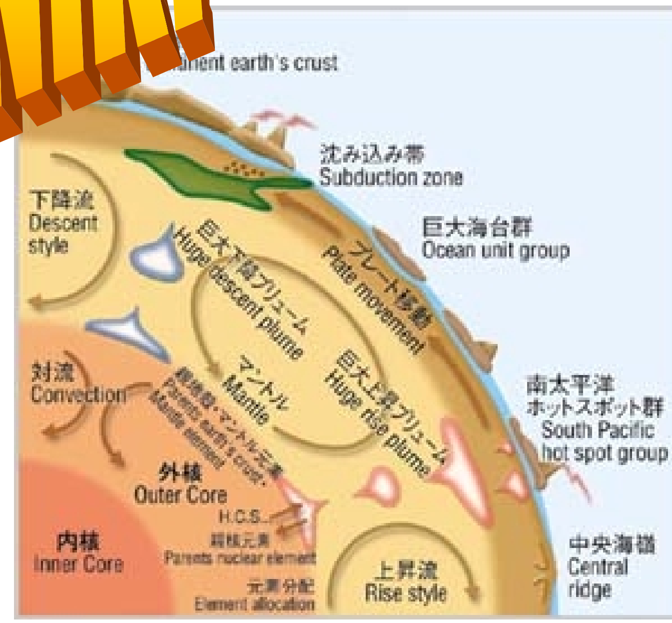

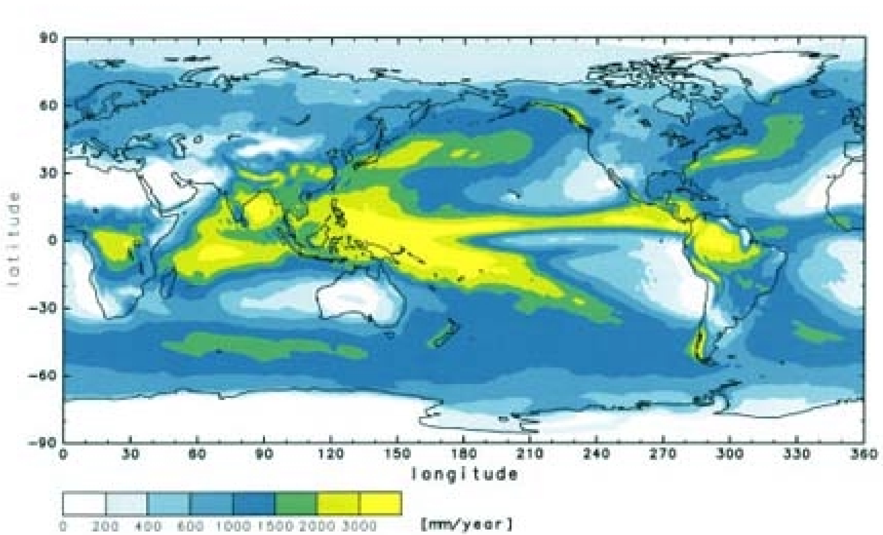

Understanding and Prediction of Global Climate Change	Understanding of Plate Tectonics
Occurrence prediction of meteorological disaster	Understanding of long-range crustal movements
Occurrence prediction of El Niño	Understanding of mechanism of seismicity
Understanding of effect of global warming	Understanding of migration of underground water and materials transfer in strata
Establishment of simulation technology with 1km resolution	

http://www.es.jamstec.go.jp/esrdc/eng/menu.html

PC Clusters: Contributions of Beowulf

- An experiment in parallel computing systems
- Established vision of low cost, high end computing
- Demonstrated effectiveness of PC clusters for some (not all) classes of applications
- Provided networking software
- Conveyed findings to broad community (great PR)
- Tutorials and book
- Design standard to rally community!
- Standards beget: books, trained people, software … virtuous cycle

Adapted from Gordon Bell, presentation at Salishan

Office of Science
U.S. DEPARTMENT OF ENERGY

Linus's Law: Linux Everywhere

- All source code is "open"
- Software is or should be free (Stallman)
- Everyone is a tester
- Everything proceeds a lot faster when everyone works on one code (HPC: nothing gets done if resources are scattered)
- Anyone can support and market the code for any price
- Zero cost software attracts users!
- All the developers write lots of code
- Prevents community from losing HPC software (CM5, T3E)

"BigMac" at VATech

- 1100 dual Apple G5 2GHz CPU based nodes.
  - 8 Gflops/processor peak double precision floating performance (17.6 Tflop/s peak).
  - 4.4 TByte memory, 176 Tbyte disk
- Infiniband Interconnect

Supercomputing Today

- Microprocessors have made desktop computing in 2003 what supercomputing was in 1993.
- Massive Parallelism has changed the “high end” completely.
- Today clusters of Symmetric Multiprocessors are the standard supercomputer architecture
- The microprocessor revolution will continue with little attenuation for at least another 10 years.
- Continued discussion over architecture for HEC (custom versus commodity)

Overview

- Introducing NERSC
- Current trends in supercomputing (high end computing)
- Computational science today
- Implications for soft computing

Simulation: The Third Pillar of Science

- Traditional scientific and engineering paradigm:
  1) Do theory or paper design.
  2) Perform experiments or build system.
- Limitations:
  – Too difficult -- build large wind tunnels.
  – Too expensive -- build a throw-away passenger jet.
  – Too slow -- wait for climate or galactic evolution.
  – Too dangerous -- weapons, drug design, climate experimentation.
- Computational science paradigm:
  3) Use high performance computer systems to simulate the phenomenon
     - Basd on known physical laws and efficient numerical methods.

Major Areas of Computational Simulation

- **Science**
  - Global climate modeling
  - Astrophysical modeling
  - Biology: genomics; protein folding; drug design
  - Computational Chemistry
  - Computational Material Sciences and Nanosciences
- **Engineering**
  - Crash simulation
  - Semiconductor design
  - Earthquake and structural modeling
  - Computational fluid dynamics
  - Combustion
- **Business**
  - Financial and economic modeling
  - Transaction processing, web services and search engines
- **Defense**
  - Nuclear weapons -- test by simulations
  - Cryptography

Global Climate Modeling Problem

- **Problem is to compute:**

  f(latitude, longitude, elevation, time) $\rightarrow$

  temperature, pressure, humidity, wind velocity
- **Approach:**
  - **Discretize the domain, e.g., a measurement point every 10 km**
  - **Devise an algorithm to predict weather at time t+1 given t**

- Uses:
  - Predict major events, e.g., El Nino
  - Use in setting air emissions standards

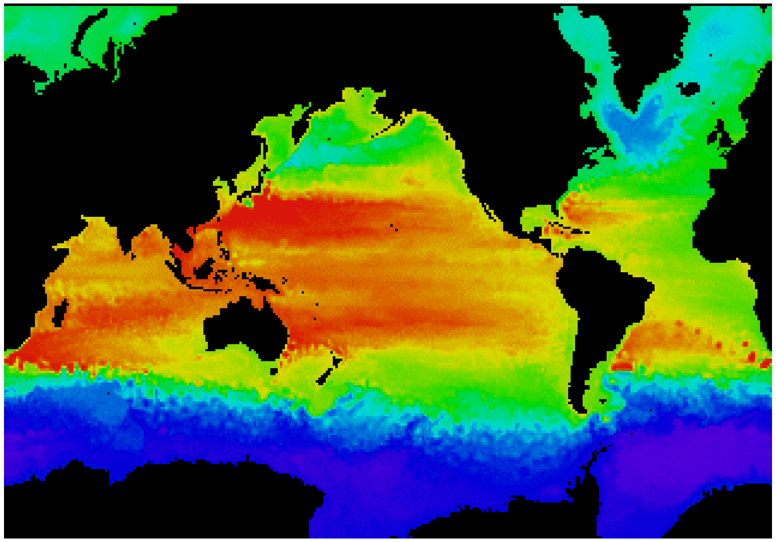

Source: http://www.epm.ornl.gov/chammp/chammp.html

High Resolution Climate Modeling on NERSC-3 – P. Duffy, et al., LLNL

Wintertime Precipitation

As model resolution becomes finer, results converge towards observations

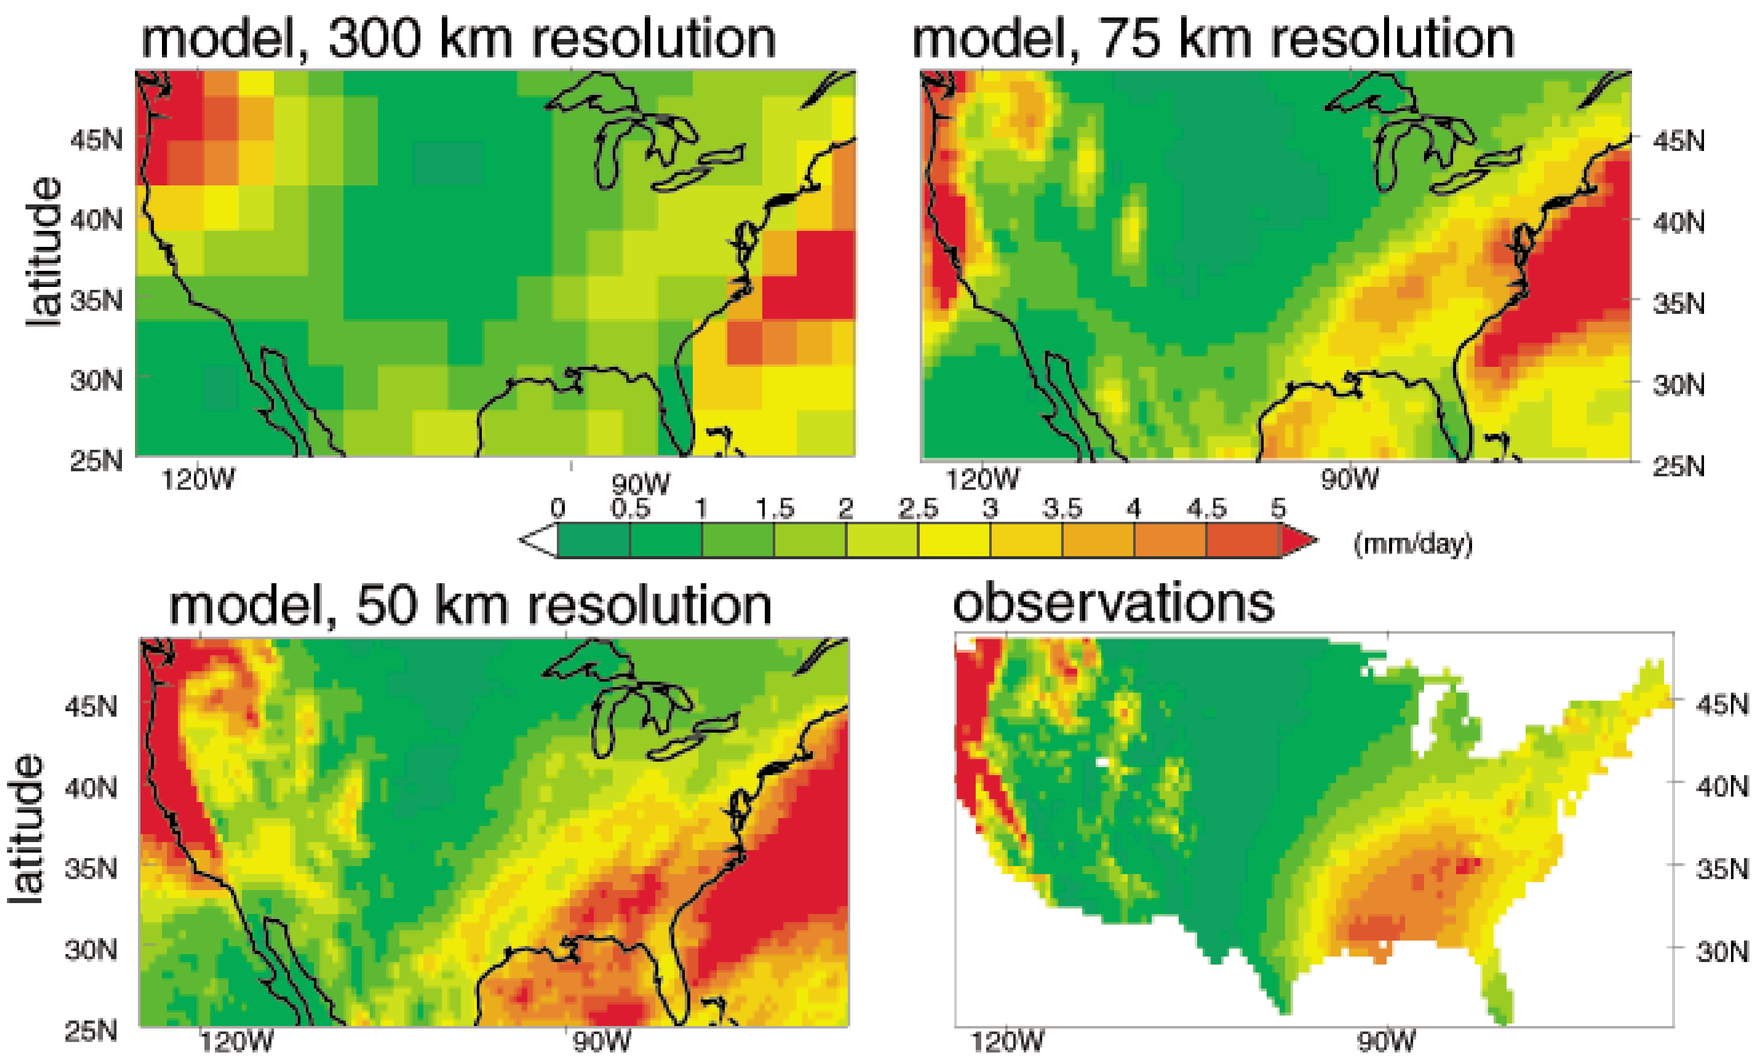

Comp. Science : A 1000 year climate simulation

- ***Warren Washington and Jerry Meehl, National Center for Atmospheric Research; Bert Semtner, Naval Postgraduate School; John Weatherly, U.S. Army Cold Regions Research and Engineering Lab Laboratory et al.***
  - A 1000-year simulation demonstrates the ability of the new Community Climate System Model (CCSM2) to produce a long-term, stable representation of the earth's climate.
  - 760,000 processor hours used

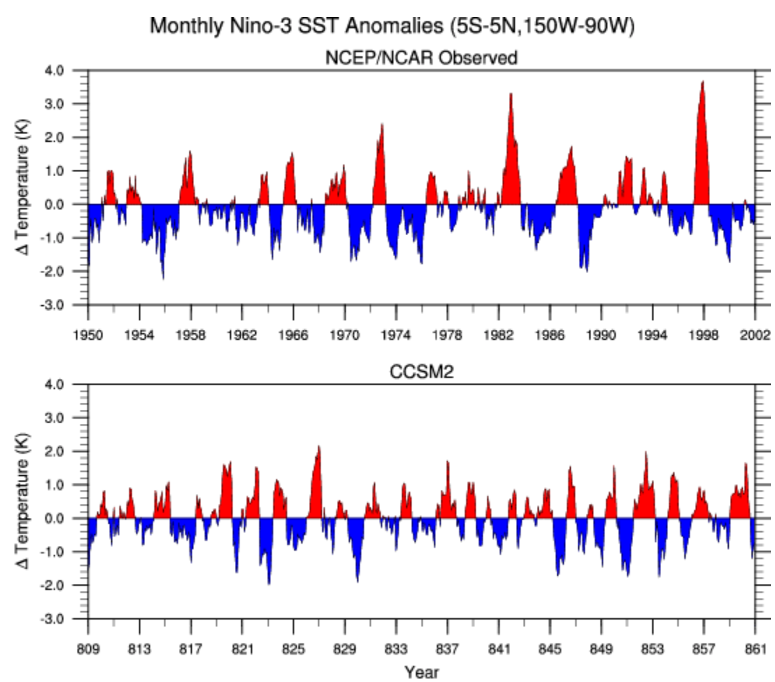

http://www.nersc.gov/aboutnersc/pubs/bigsplash.pdf

Computational Science and Engineering (CSE)

- **CSE (or CSME) is a widely accepted label for an evolving field concerned with the science of and the engineering of systems and methodologies to solve computational problems arising throughout science and engineering**
- **CSE is characterized by**
  - Multi - disciplinary
  - Multi - institutional
  - Requiring high end resources
  - Large teams
  - Focus on community software
- **CSE is not “just programming” (and not CS)**
- **Ref: Petzold, L., *et al*., Graduate Education in CSE, *SIAM Rev*., 43(2001), 163-177**

Overview

- Introducing NERSC
- Current trends in supercomputing (high end computing)
- Computational science today
- Implications for soft computing

Projected Performance Development

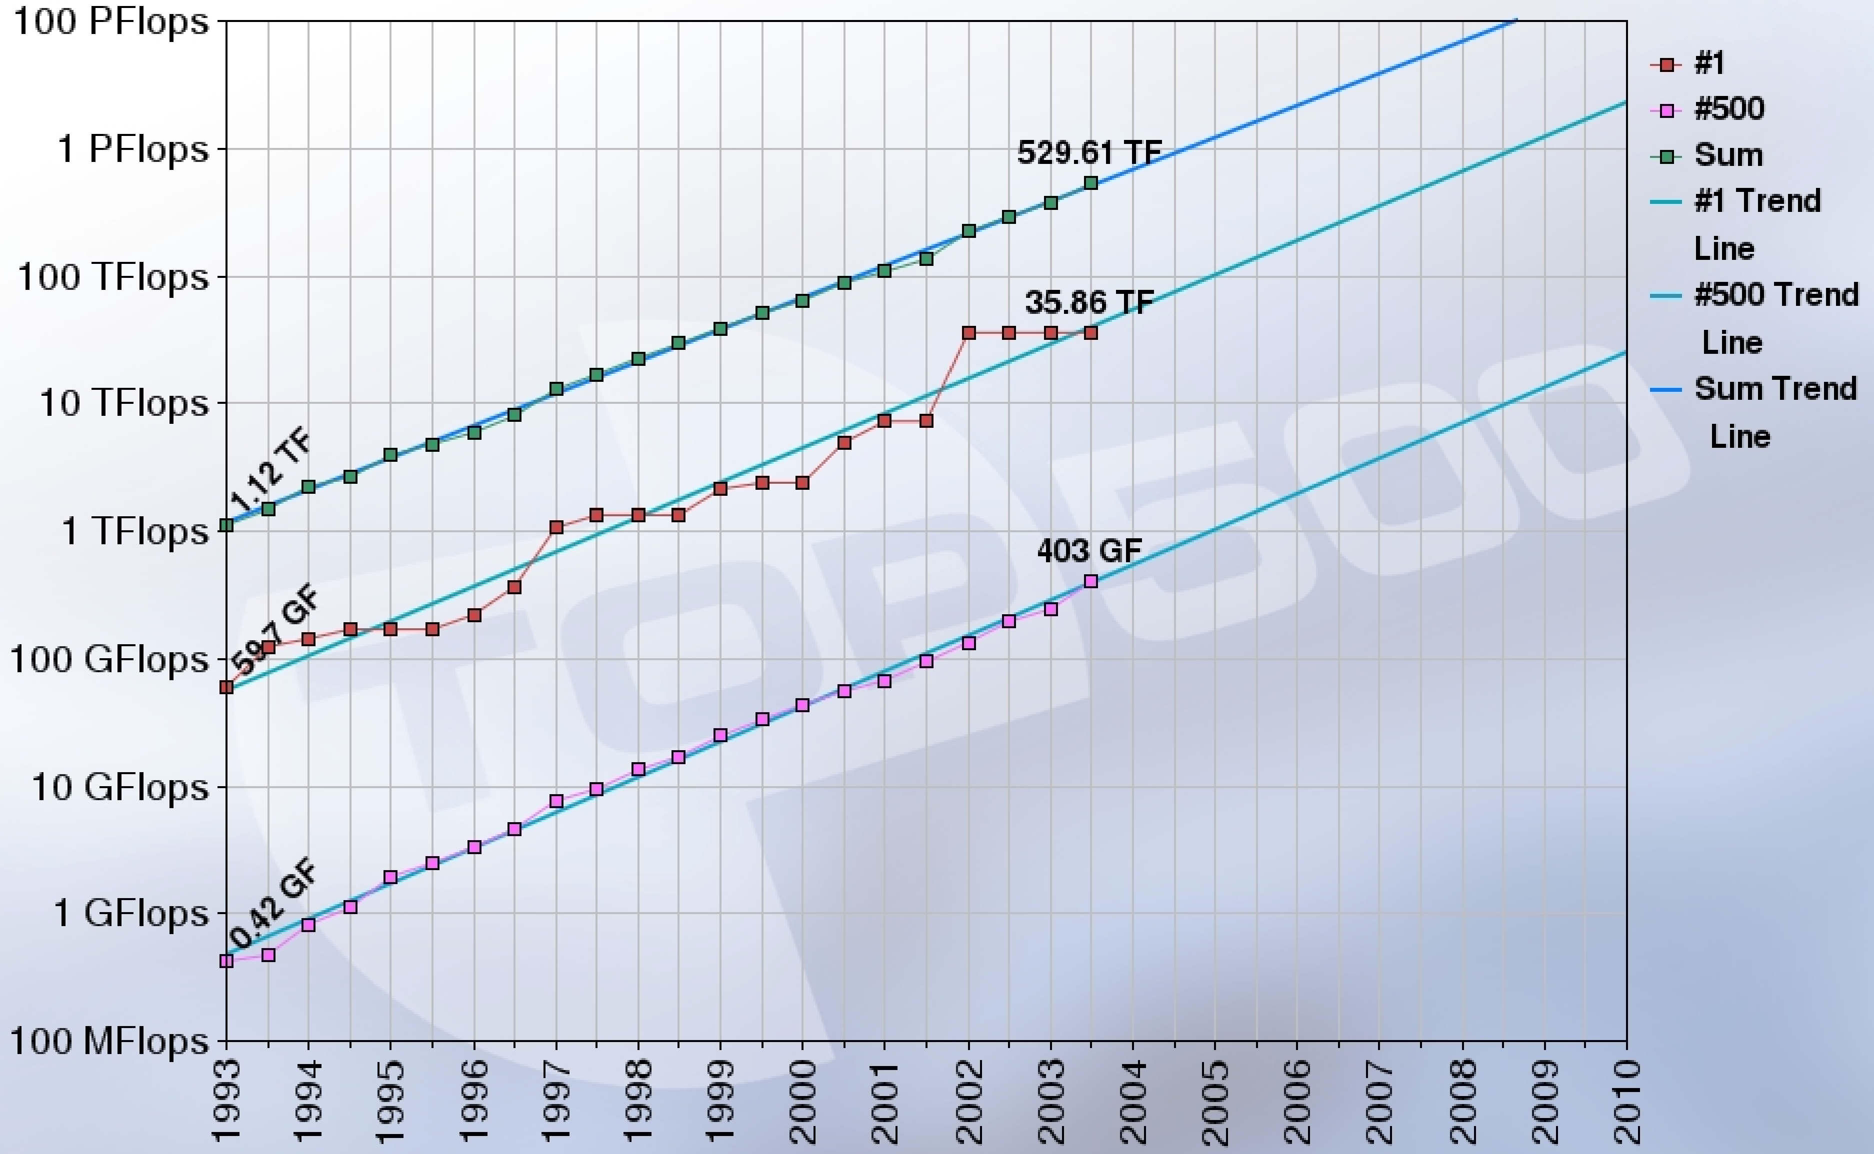

The Exponential Growth of Computing, 1900-1998

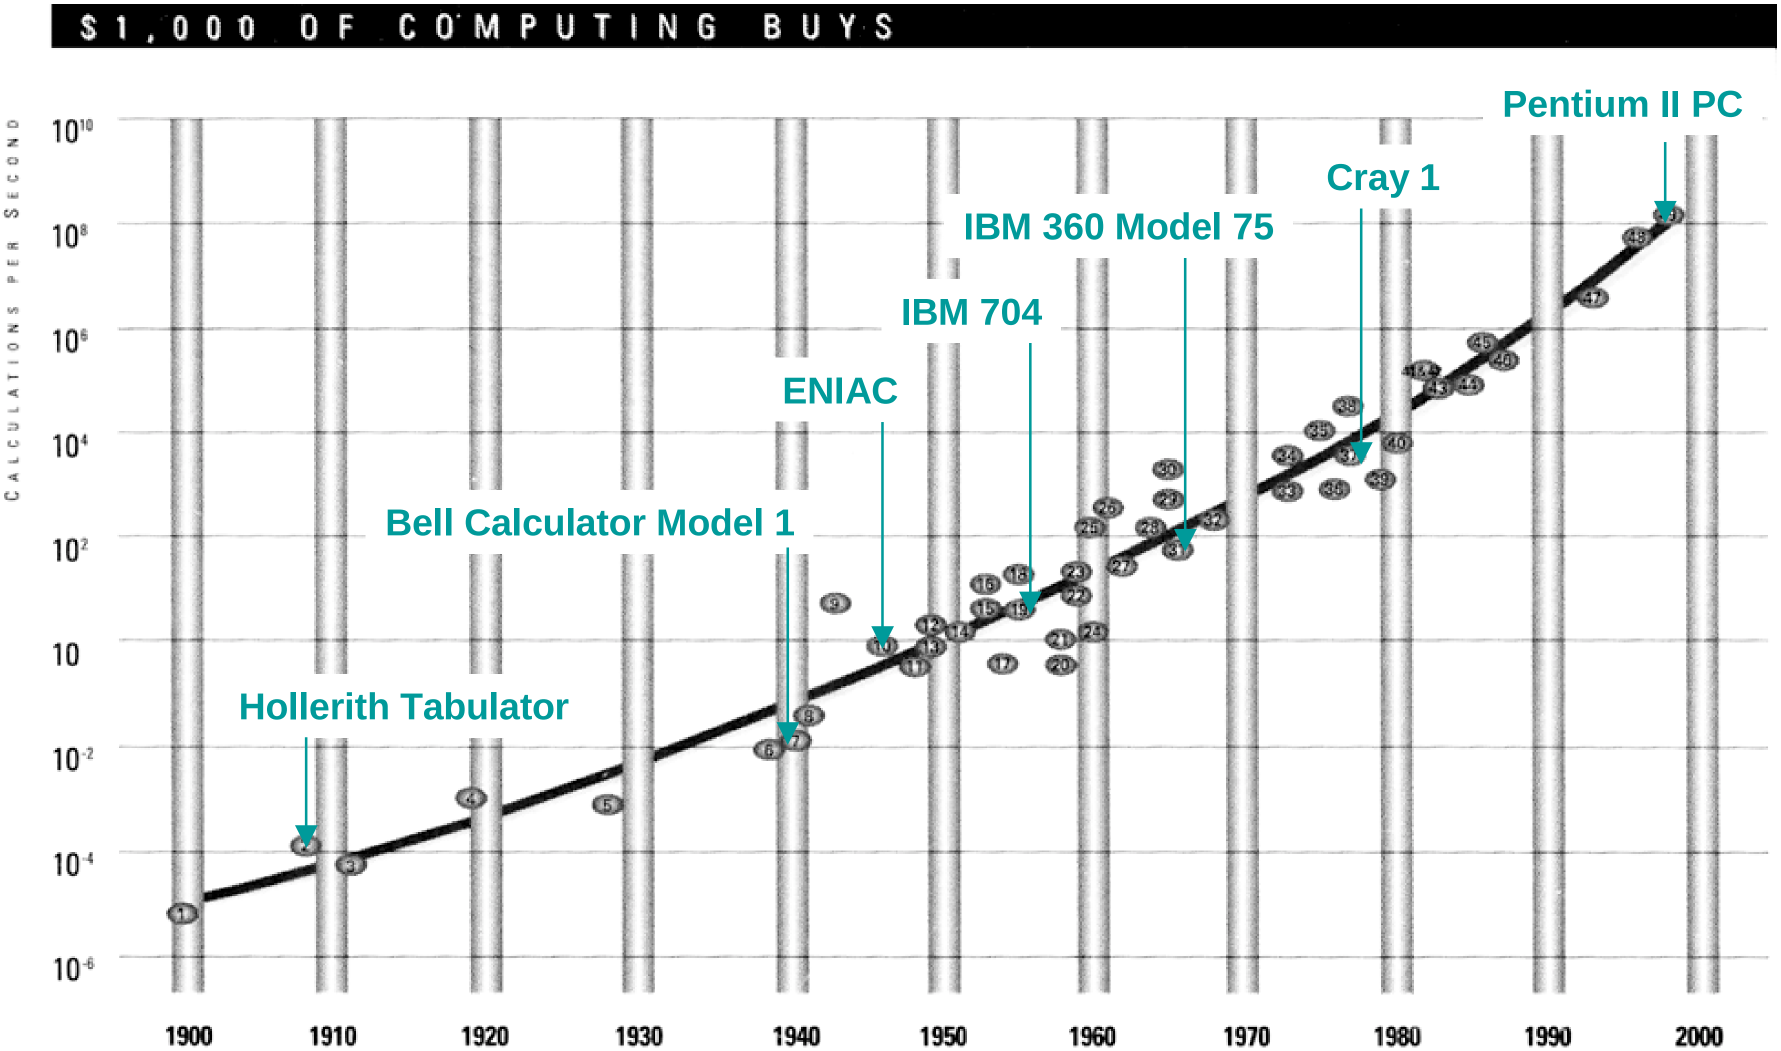

Adapted from Kurzweil, *The Age of Spiritual Machines*

The Exponential Growth of Computing, 1900-2100

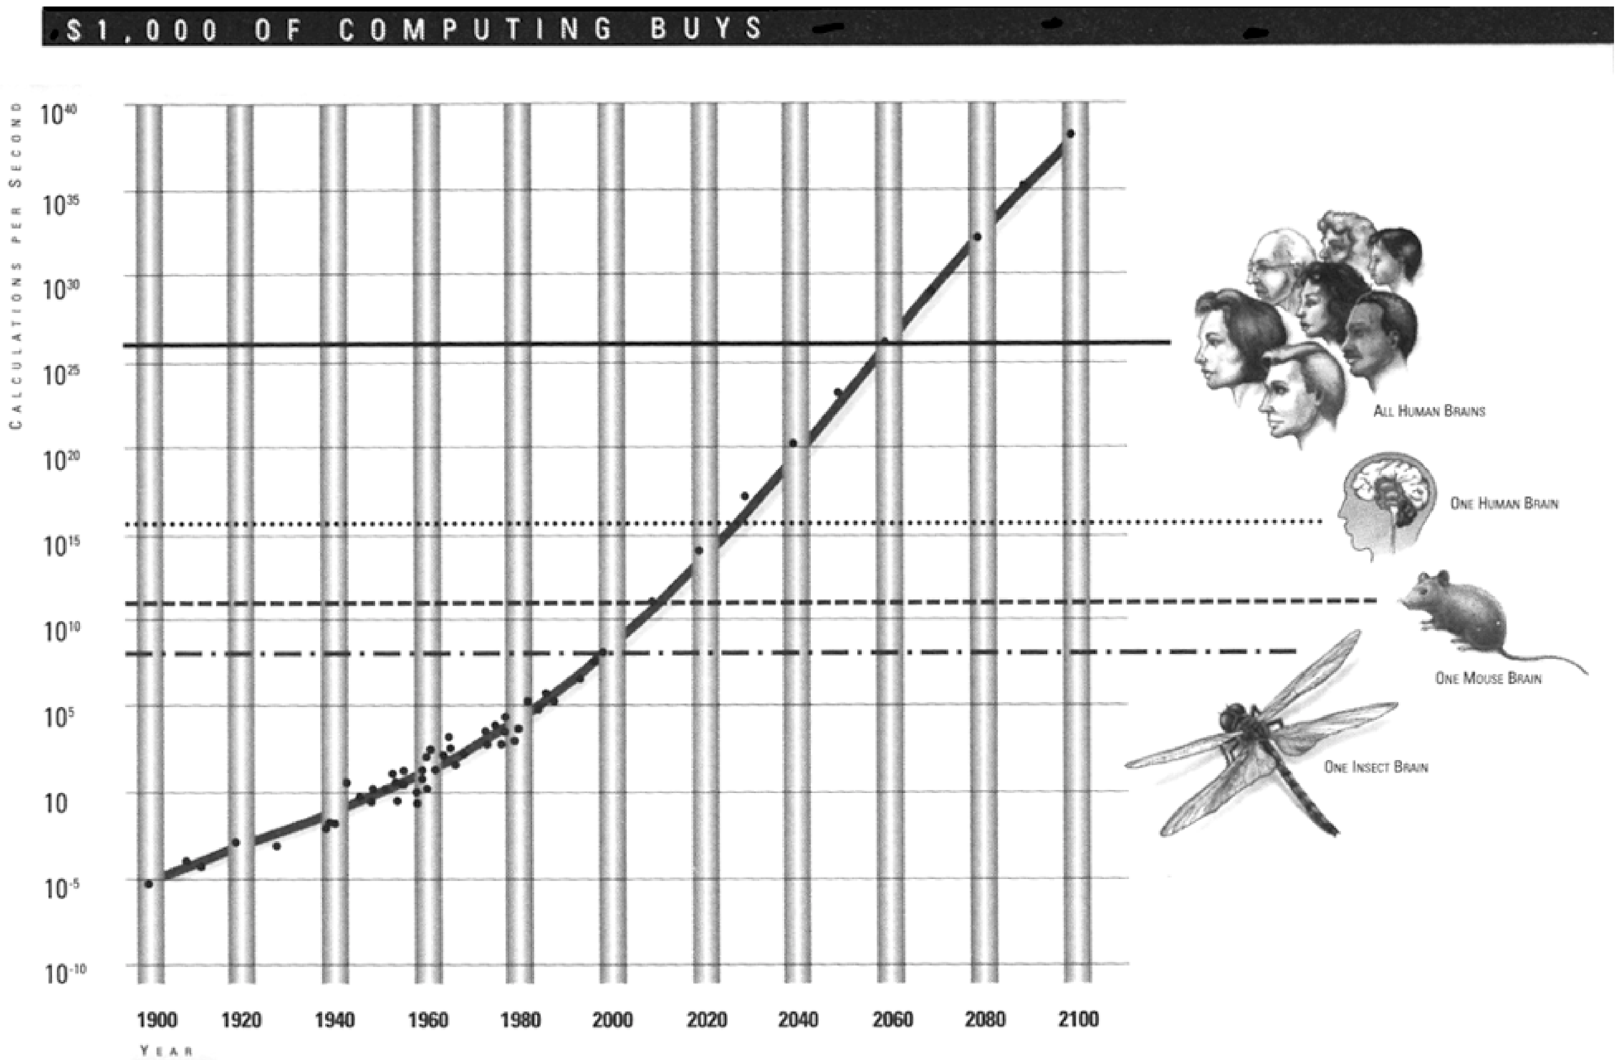

Adapted from Kurzweil, *The Age of Spiritual Machines*

Growth of Computing Power and "Mental Power"

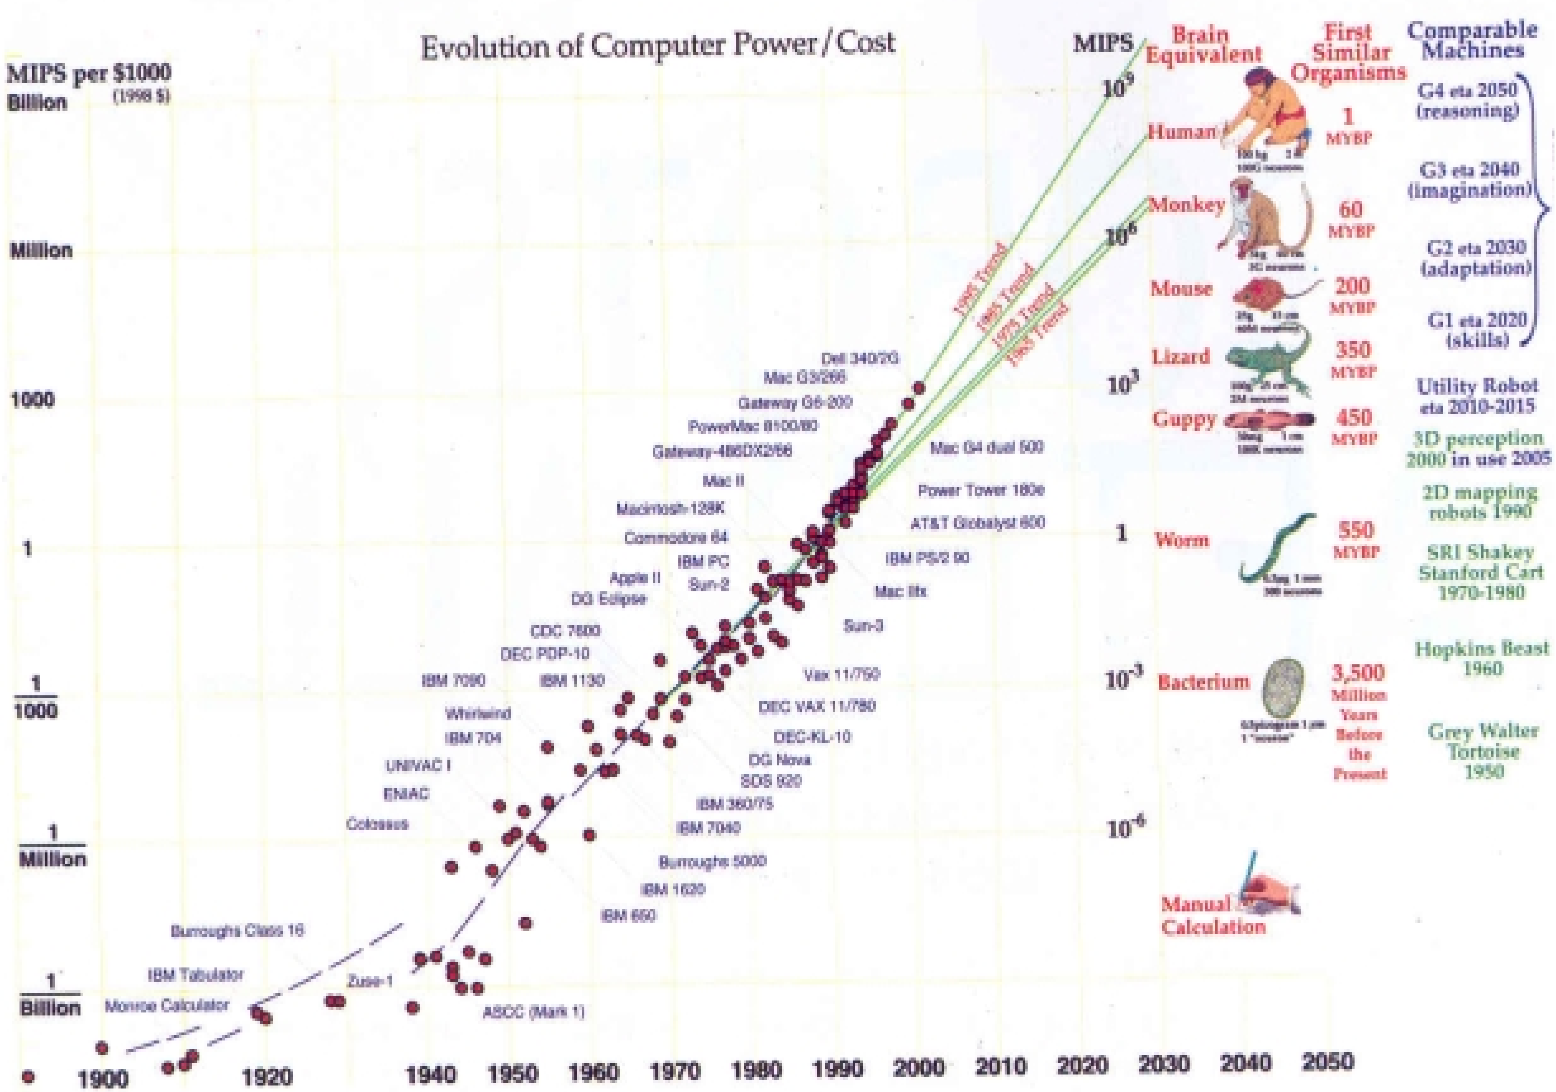

Hans Moravec, CACM 10, 2003, pp 90 - 97

Megabytes-per-MIPS

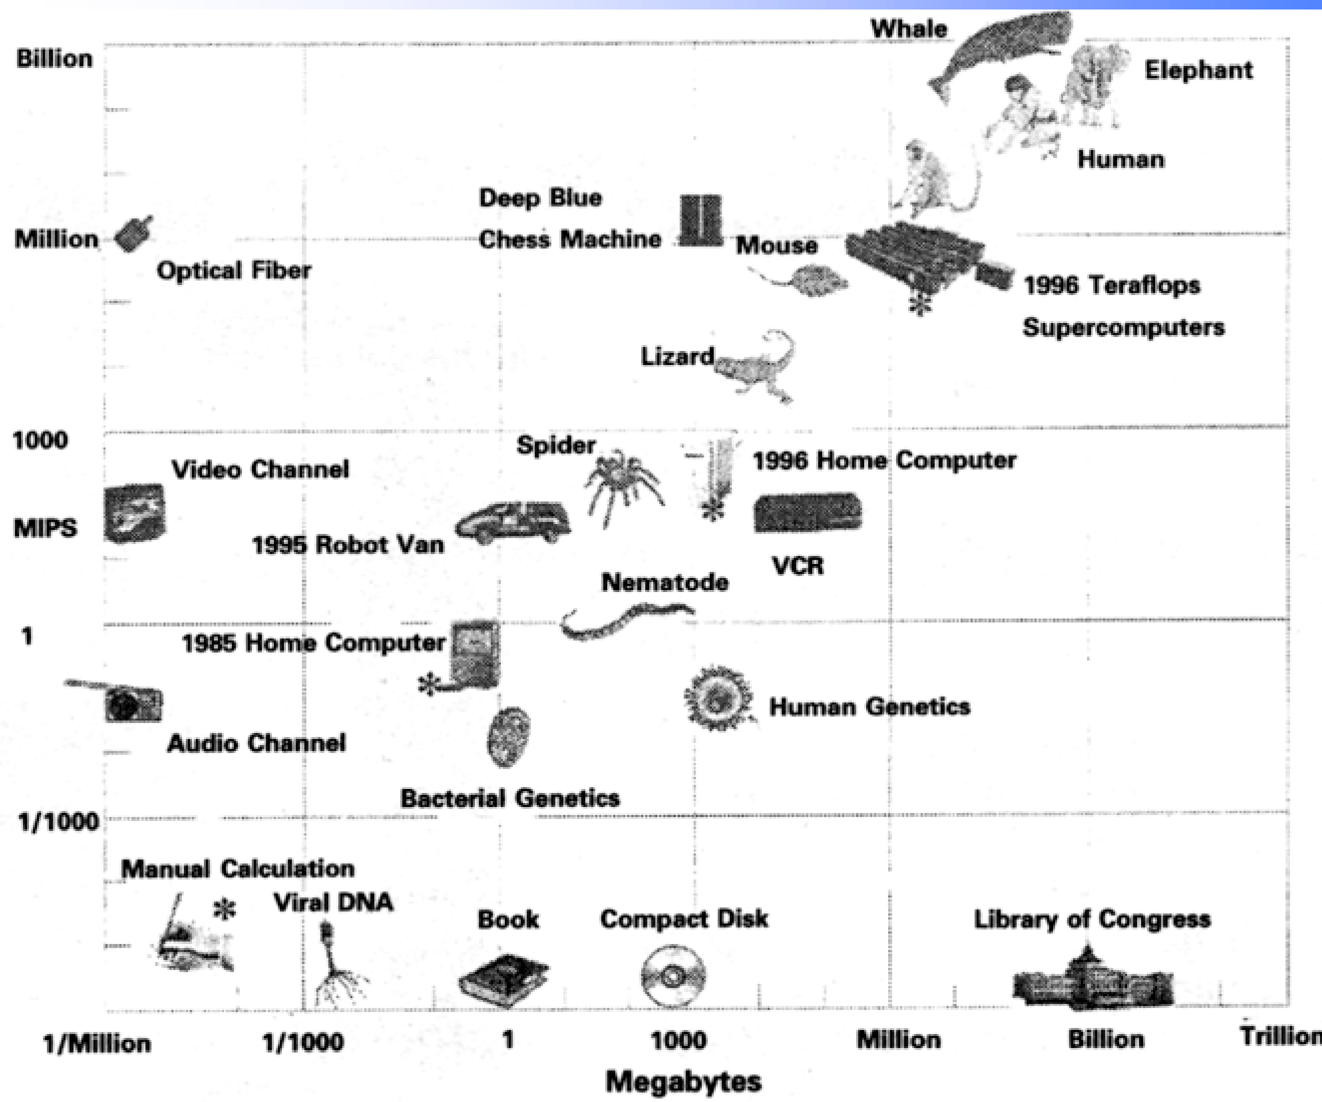

Adapted from Moravec, *Progress in Robotics: An Acclerated Rerun Of Natural History?*

Why I don’t share this vision

- Fundamental lack of mathematical models for cognitive processes
  - that’s why we are not using the most powerful computers today for cognitive tasks
- Complexity limits
  - we don’t even know yet how to model turbulence, how then do we model thought
- past experience: people bring technology into their lives, not vice versa
  - the slow rate of change of human behavior and processes will work against this change

Moore's Law — The Exponential View

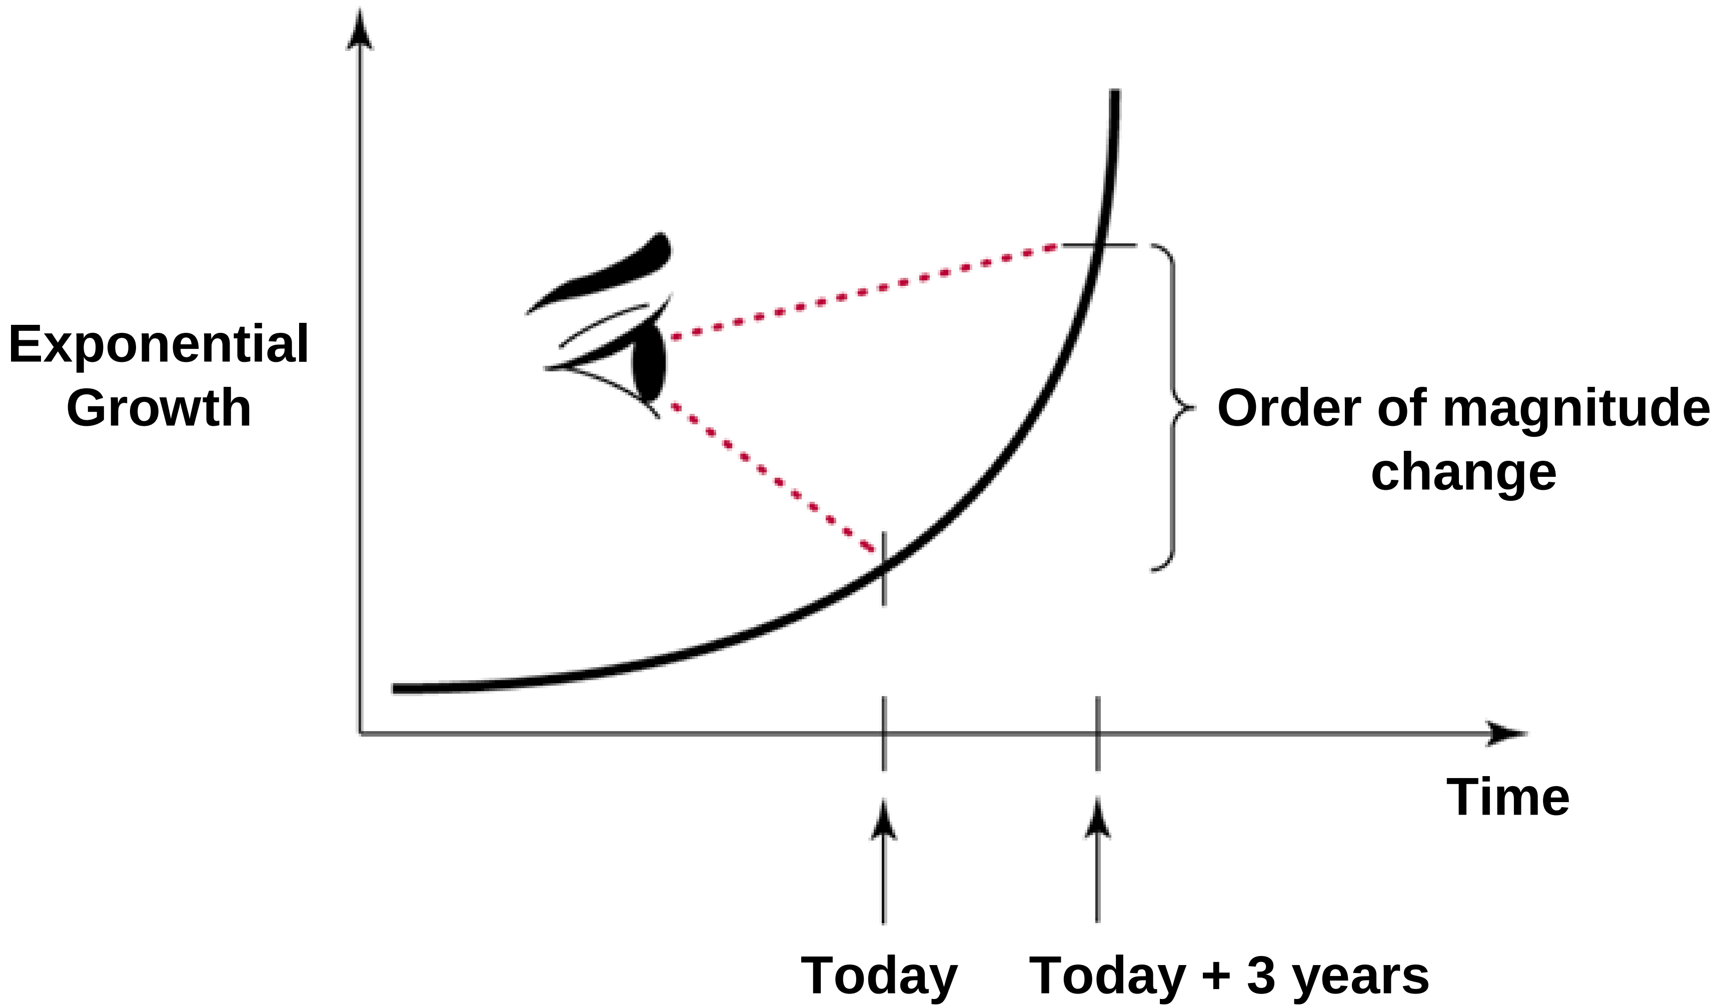

Moore's Wall —
The Real (Exponential) View

Conclusion

- Research opportunities in bringing scientific computing in soft computing together
- Near term:
  - high end parallel computers get rapidly less expensive and easier to use
  - soft computing community is missing a scientific opportunity by not using this technology more aggressively
- Long Term:
  - floating point intensive computational science has driven development of current high end technology
  - in order to make real progress in soft computing, we need research into appropriate high end architectures for the task