

The Supercomputer Class Evolution: A personal perspective

Gordon Bell

Microsoft Research, Silicon Valley Laboratory (1995-)

24 April 2013

Digital Equipment Corp. (DEC) 1960-1983

Abstract

Since my first visit to Livermore in 1961 seeing the LARC, the elegance of the 6600, and just observing computer evolution have been high points of my life as an engineer and computing observer. Throughout their early evolution, supercomputer architecture “trickled down” for use with other computers. In the mid 1990s the flow reversed when large computers became scalable and constructed from clusters of microprocessor-based computers that Livermore’s Eugene Brooks wrote about in 1990. Unlike the two paths of Bell’s Law that account for the birth, evolution, and death of other computer classes e.g. minicomputers, http://ieeeghn.org/wiki/index.php/STARS:Rise_and_Fall_of_Minicomputers supercomputers have doubled in performance every year for the last 50 years.

While computer performance is the first order term to track their high performance, many other technologies e.g. FORTRAN, LINPACK, government funding policy, and applications have contributed to the extraordinary progress. I hope to trace the trajectory and contributors to this exciting class.

Events in Supercomputing

- | | |
|---|--|
| <ol style="list-style-type: none"> 1. 1936 Turing "On Computable Numbers" 2. 1943 Colossus, first programmable computer, not Turing complete, for code breaking... built to a schedule, not to max technology that characterizes supers; many Bombe's 3. 1945 ENIAC built 4. 1945 EDVAC Draft Report 5. 1946 IAS Architecture. Preliminary discussion of the logical design of an electronic computing instrument. Burks, Goldstine, and von Neumann. Over a dozen x-IACS were built from the general design. 6. 1953 IBM 701 Defense Calculator introduced for scientific calculator. 36-bit version of IAS Architecture 7. 1954 IBM 704; Successor was 1958: 709. 8. 1957 Fortran first delivery. Current FORTRAN 2008 9. 1959 7090 and 1108 10. 1960 CDC 1604 (48 bits) 11. 1960 LARC and 1961 Stretch 12. 1962 Atlas commissioned 13. 1964 CDC 6600 (48 MF) no more small computers 14. 1964 IBM System 360 15. 1965 Amdahl's Law arguing difficulty of multiprocessors and vectors in the CDC decree. 16. 1967-68 IBM 360/91-95 introduced as "fighting machine" resulting in the CDC decree. 17. 1969 CDC 7600 (3.3 MF) 18. 1974 CDC STAR 100 (Mflops) 19. 1975 ILLIAC IV connected to ARPAnet 20. 1976 Cray 1 (26 MF) 21. 1977 VAX-11/780 (.14) Became a Personal Supercomputer 22. 1978 Caltech Cosmic Cube 23. 1982 Lax Report of the Panel on Large Scale Computing in Science and Engineering to NSB. 24. 1982 Japanese Fifth Generation Computer Systems Project for AI 25. 1983-93 Billion dollar SCI--Strategic Computing Initiative of DARPA IPTO response to Japanese Fifth Gen. 1990 redirected to supercomputing after failure to achieve AI goals | <ol style="list-style-type: none"> 1. 1983 MatLab establishes a higher level language and programming interface 2. 1982 Cray XMP... (1 GF) 3. 1984 NSF Establishes Office of Scientific Computing. 1985 NCSA awarded 5 year grant of 42.7 M 4. 1985 Convex C-1 (3.0) & Alliant FX-1...8 Mini-supercomputers are born. 5. 1986 Livermore Loops Livermore Fortran Kernels 6. 1986 Ken Wilson declares simulation to be the third paradigm 7. 1987 Bell CISE Interview posits parallelism. 1992: SX, 1997: 100X 8. 1987 nCUBE (1K computers) achieves 400-600 speedup, winning first Bell Prize at Sandia. Gustafson's Law as Amdahl's Law Corollary 9. 1988 Mathematica 10. 1988 Ardent Personal Graphic Supercomputer (6.9-25) 11. 1989 TMC CM2 SIMD with 64K PEs wins Bell Prize 12. 1991 Eugene Brooks, Attack of the Killer Micros report; Supercomputing 90 13. 1992 Intel Touchstone Delta at Sandia? 14. 1993 CM5 (60 GF) 1024 processors 15. 1993 Top500 established using LINPACK Benchmark, 2012 Graph500, 1994 Bell and Gray define Bricks aka SNAP as "the way". 16. 1994 Beowulf standard, and 1994 MPI-1 Standard 17. 1995 Death of o(50) companies in search of "the way" via DARPA SCI 18. 1995 ASCI > Advanced Simulation and Computing (ASC) Program 19. 1995 Numerical Wind Tunnel (124 GF) 20. 1997 ASCI Red (1 TF) at Sandia 21. 1998 Gray Posits 4th Paradigm of Science. 22. 2000 RLX Rocket Logic introduces Blade package 23. 2002 Earth Simulator (40 TF) 24. 2008 IBM BlueGene (1.5 PF) 25. 2010 Graph500 Rating 26. 2012 Cray Titan (17.6) GPU and CUDA 27. 2013 Green500 |
|---|--|

Personal Vignettes

1. Reminiscence: 1st visit to Livermore, 1961
2. Bell's Law for Birth and Death of Computer Classes
3. The Cray Era (1965-1990): Single Memory computers
4. "Killer Micros" transition
5. Massive Parallelism: Clusters
6. Grand Challenges, Benchmarks, Lists, Prizes, Threats
7. HPC aka Supercomputers versus Cloud computing
8. To the ExaFlops-- Continued diversity



PROGRAMMED DATA PROCESSOR - 1 MANUAL

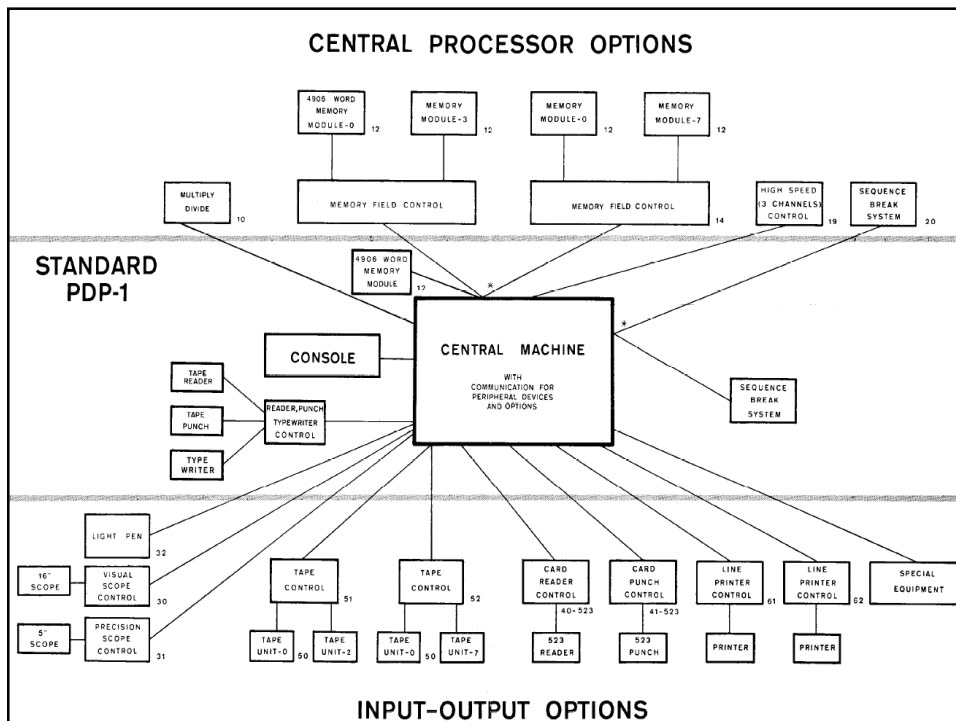
F-15C 6/63

TABLE OF CONTENTS

INTRODUCTION	3
PROGRAMMING PDP-1	7
STANDARD AND OPTIONAL EQUIPMENT	23
PROGRAM LIBRARY	53
APPENDIX	54

Copyright 1960, 1961, 1962 by Digital Equipment Corporation

DIGITAL EQUIPMENT CORPORATION - MAYNARD, MASSACHUSETT




As delivered (George Michael)

4k words x 18 bits/word

Freiden flexowriter + fan folded
Paper Tape reader/punch

Type 30 Direct View CRT +
**Type 31 High Precision 5" CRT +
Camera and light pen**

**Ten frame/sec Mitchell pin-
registered camera ...leaked light**

4 Potter Magnetic tape handlers

**2,000 crd/min Uptime Card
reader folded cards very fast**

600 lpm Analex printer

As Modified ... Called "Romper Room"

8k words x 18 bits/word

Same typewriter + reels for Paper tape Reader + paper
tape punch

Direct View CRT + New Light Pen; 5" CRT + special
phosphor, Mitchell camera, claw pull down, pin
registered, 1,000 ft magazines, special (CRT) dichroic
filters and reference leg.

Quadriphonic sound output + 10 bit D-to-A converters

Telephone handset + microphone for sound input

4 IBM 729VI Magnetic tape handlers

1200 crd/min IBM 1402 card reader/punch

1200 lpm IBM 1403 Printer

Cal Comp Plotters; 12- and 30- inch widths

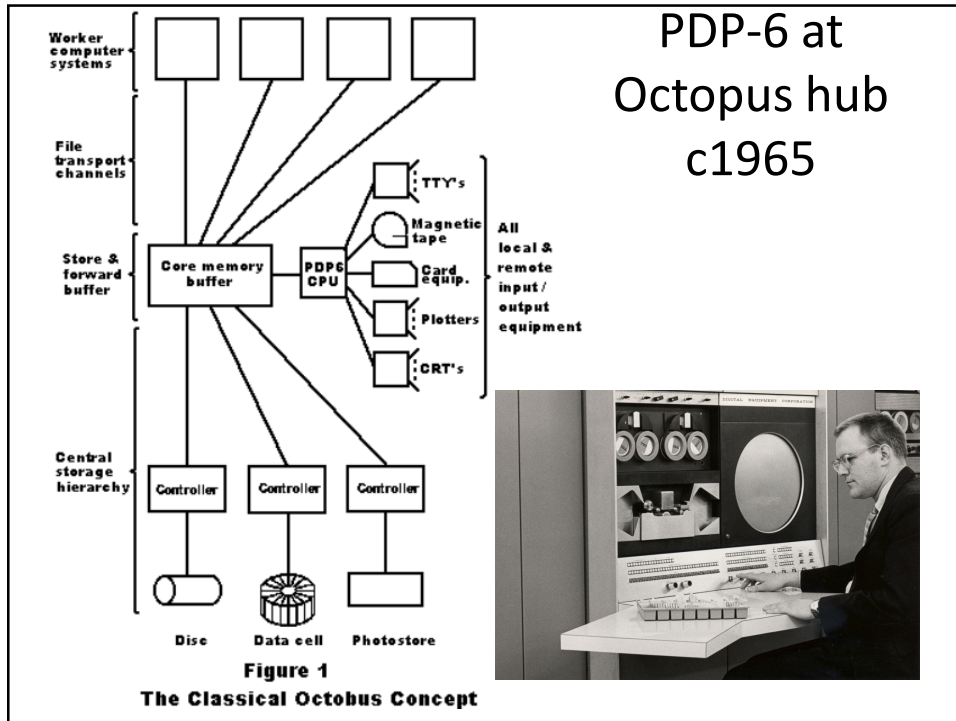
Rand Tablet

EYEBALL image digitizing

Large Dark Room

Mylar Paper Tape Reader/Punch

IBM Selectric Typewriter



LLNL Visit 1999



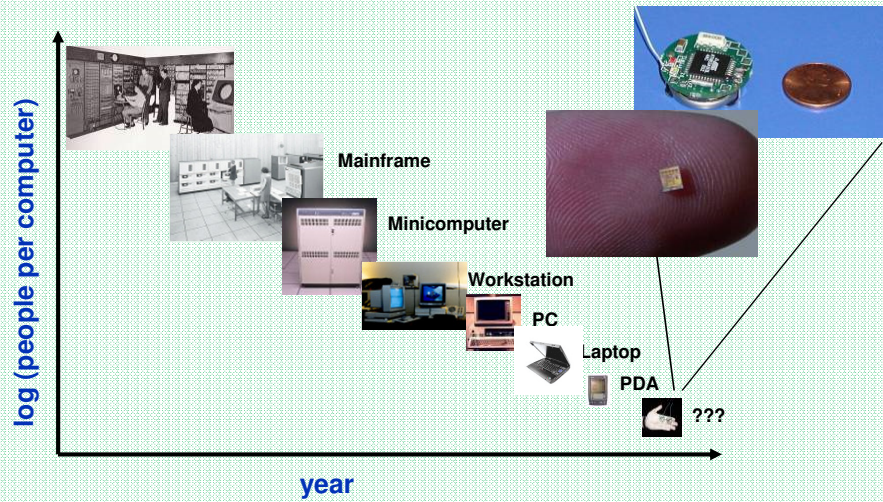
Begin Bell's Law

Bell's Law for the formation of The Birth & Death of Computer Classes

Hardware technology improvements i.e. Moore's Law for semiconductors,... disks, enable two evolutionary paths(t) for computers:

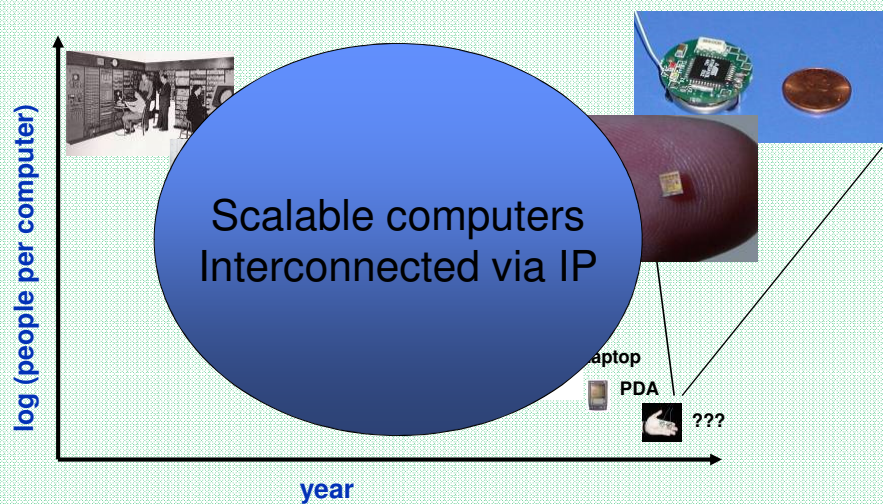
1. constant price, increasing performance
direct consequence of Moore's Law
2. **Constant or decreasing performance,
decreasing cost by a factor $O(10)X$
.. Leads to new structures & new computer class!**
3. **Spend more: build the largest computer that
you can and that customers can afford.**
Class = platform, price, use, market, interface, etc

The classes, sans phones, 2006

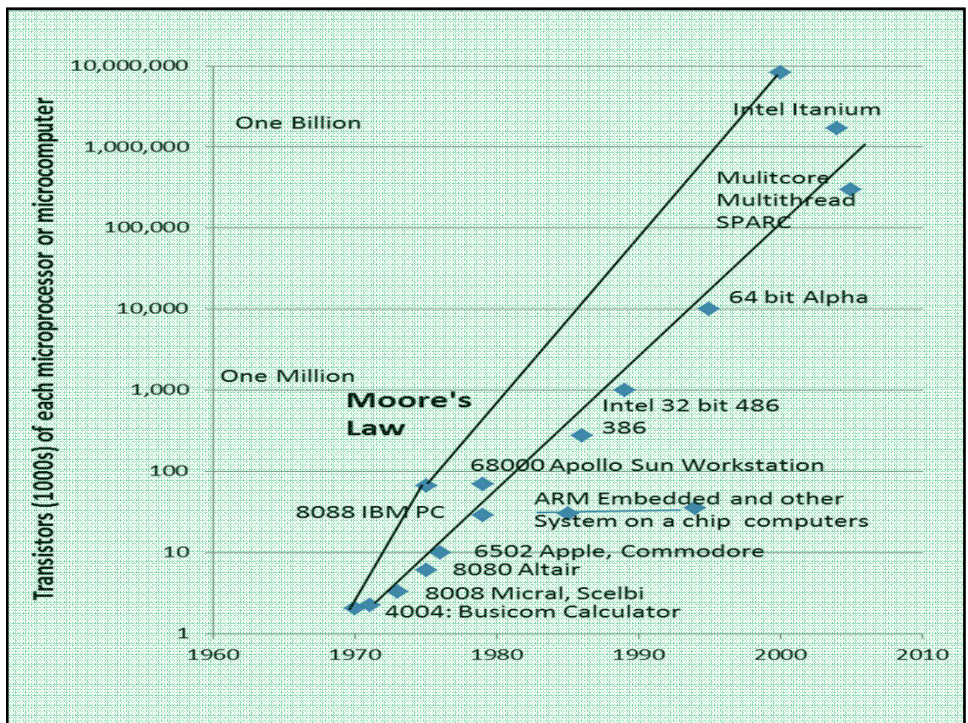
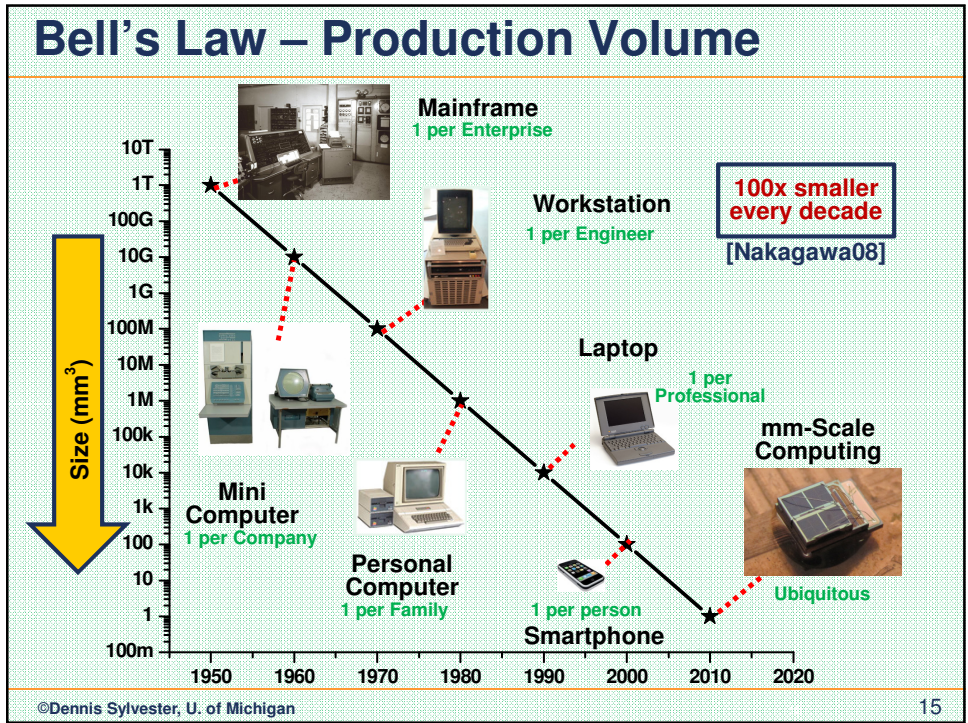


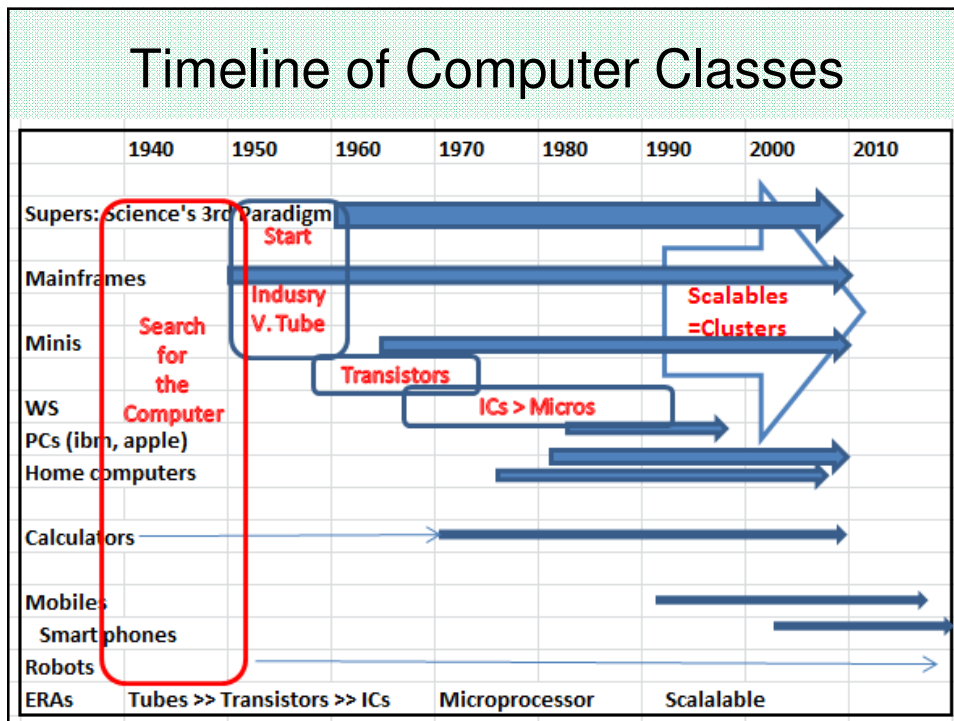
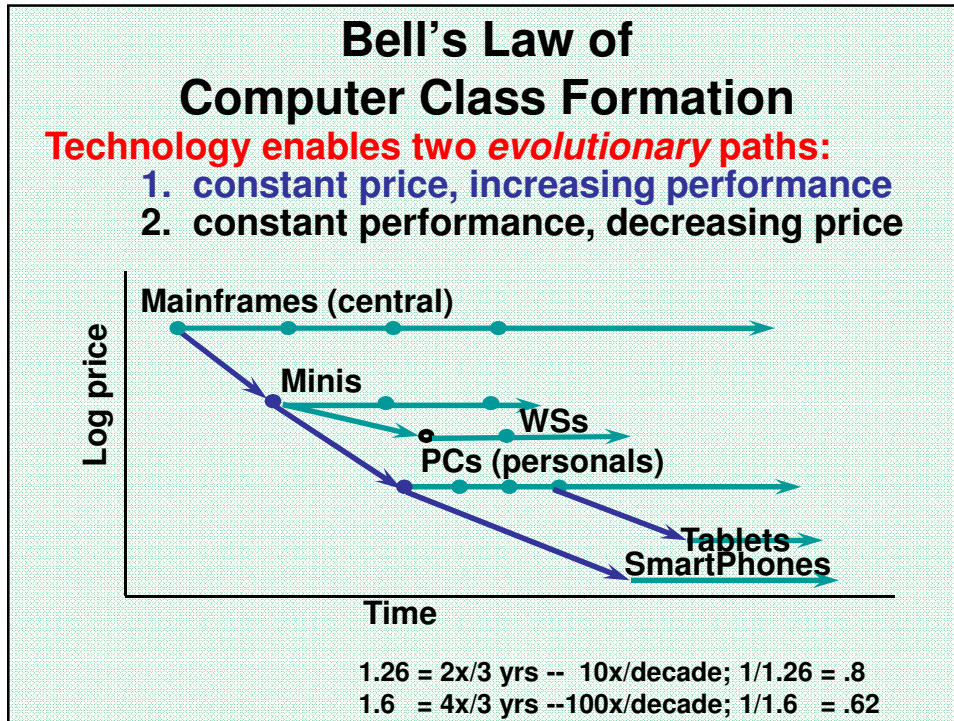
David Culler UC/Berkeley

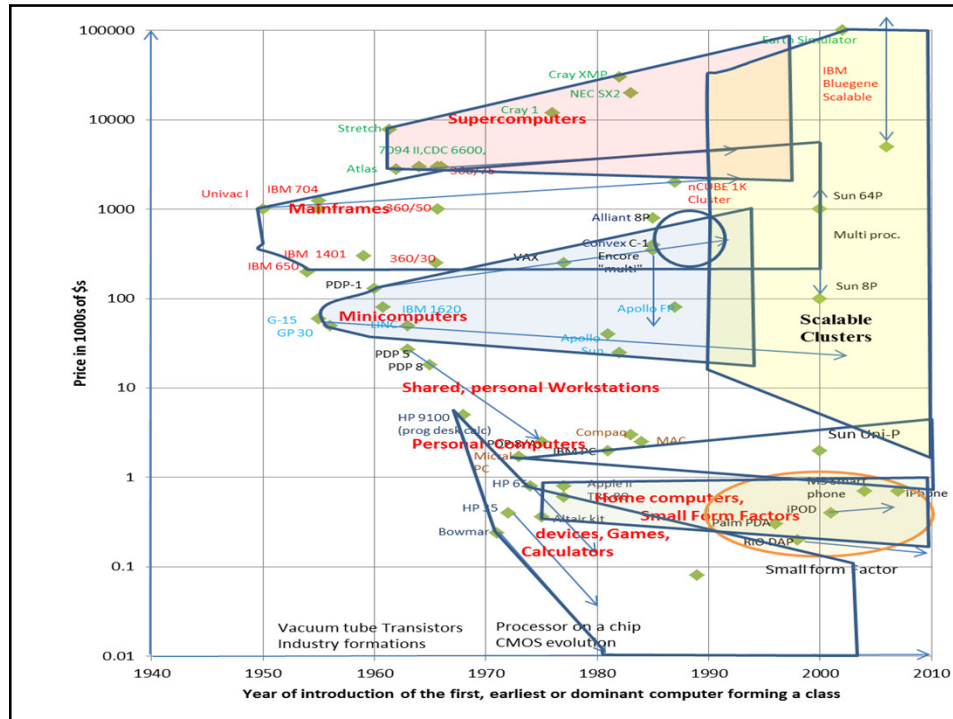
The classes, sans phones, 2006



David Culler UC/Berkeley







In retrospect...by 1971, the next 40+ years of computing was “set in stone”

1. Moore’s Law (1965) transistors/die double every 18 mos.
2. Intel 4004, Processor-on-a-chip (1971)
Clearly, by 1978 16-bit processor-on-a-chip
3. Bell et al 1971 observation...computers evolve:
 1. At constant price (Moore’s Law)
 2. Less costly computers form a new class every decade (Bell’s Law)
New technology, interfaces, manufacturers, uses, and markets create new classes

End Bell's Law

Supercomputer attributes (besides being today's largest computer)

- Function: calculate vs. record processing
- Fortran as the target language
- Quest for performance: Who can build the fastest?
Test limits: size, power, complexity, budget, time to market
- Price: How much do you have to spend?
 - To buy, to build the building, to power, to run
 - To program
- Use (market): science & engineering simulation, design, climate, cryptography ...3d, time varying phenomena
- Scalability *post 1995*: Time, machine generation, problem size, and programming environments
- Programming environment (standards): MPI, Beowulf

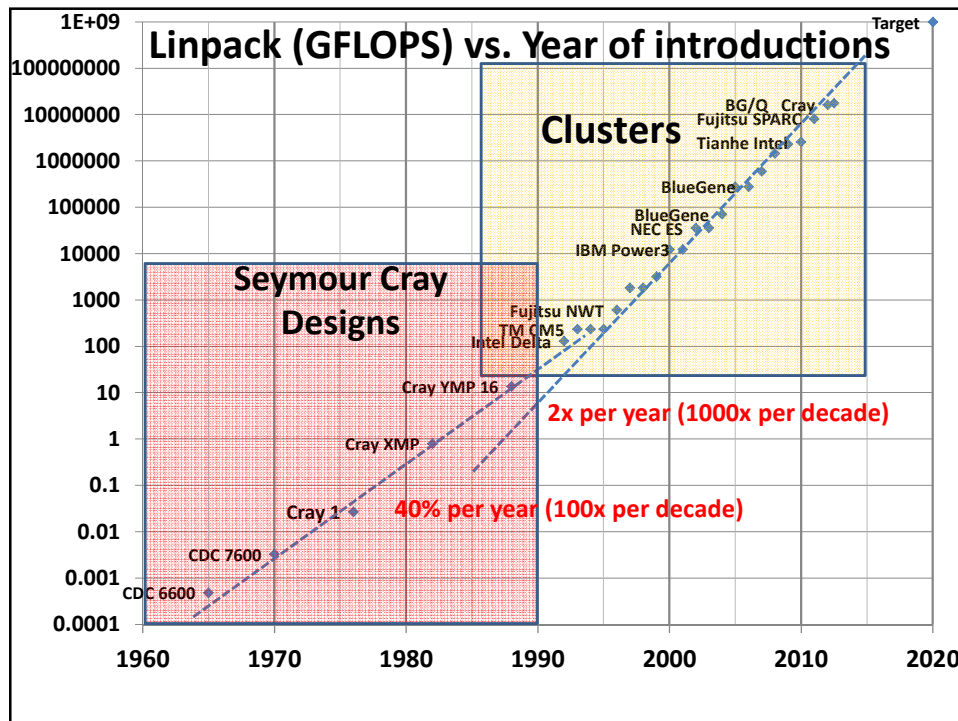
Supercomputer Hardware: Speed & parallelism

- ~~Clock speed~~
- ~~One memory (Scale up)~~
 - Overlap of memory access and instruction execution
 - Parallelism of a single instruction stream
 - VLIW
 - Pipelining
 - Vector processing
 - Multiprocessors—Scale up
 - Multiple streams & multi-threading scalability
- Multiple independent interconnected computers (Scale out) aka Clusters
 - Multiprocessor nodes aka constellations
 - Stream processing using GPUs
 - FPGA ala Convey?

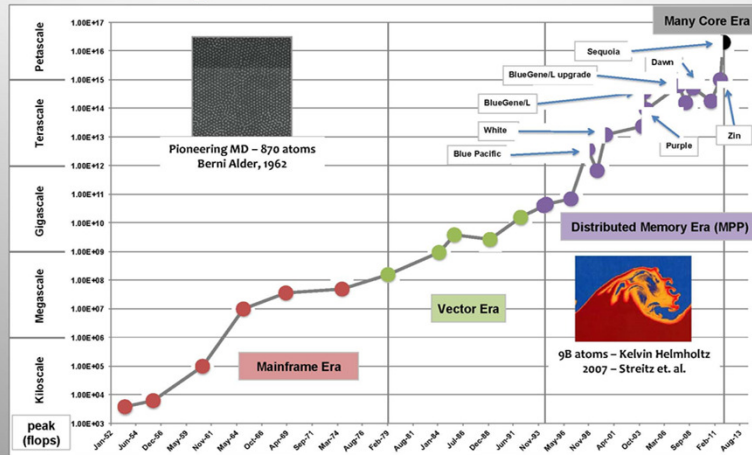
20 Supercomputing events

1. **1946** IAS Architecture. Preliminary discussion of the logical design of an electronic computing instrument. Burks, Goldstone, and von Neumann. Over a dozen x-IACS were built from the general design architecture including IBM's 701
2. **1957** Fortran first delivery. Current FORTRAN 2008
3. **1959** 7090 and 1108, CDC 1604 (scientific mainframe)
4. **1960** LARC and **1961** Stretch; **1962** Atlas commissioned
5. **1964** CDC 6600 (.48 MF)
6. **1964** IBM System 360; Scientific and Record Keeping computer
7. **1965** Amdahl's Law: single processor vs multi-P's or vectors
8. **1969** CDC 7600 (3.3 MF)
9. **1976** Cray 1 (26 MF)
10. **1978** Caltech Cosmic Cube first multicomputer
11. **1982** Cray XMP... (PK: 1 GF) The beginning of the end
12. **1987** nCUBE (1K computers) achieves 400-600 speedup, winning first Bell Prize at Sandia. Gustafson's Law as Amdahl's Law Corollary
13. **1992** Intel Touchstone Delta at Sandia? Reaches 100 GF
14. **1993** CM5 (60 GF Bell Prize) 1024 processors (Linpack 236 GF)
15. **1993** Top500 established using LINPACK Benchmark
16. **1994** Beowulf and MPI-1 Standards established
17. **1997** ASCI Red (1 TF) at Sandia
18. **2008** IBM BlueGene (1.5 PF)
19. **2012** Cray Titan (17.6) GPU and CUDA

Five eras of Scientific Computing		
Period	Technology	Machines (artifacts)
193x-1947	Electromechanical-vacuum tubes; one-of machines Search for "the computer"	Computing with cards at Los Alamos; IBM Multiplying calculator. Atanasoff, Colossus, Harvard Marks, BTL, Zuse, culminating in ENIAC. The EDVAC Report.
1947-1950s	Electronic Computing Era Vacuum Tube Scientific Calculators including von Neumann X-iacs	The Big Bang. First stored program computers that just work (Univac, IBM 701 and ERA); X-Iacs, Illiac, Maniac, etc. Amdahl's WISC
1960s	Discrete transistors. Supercomputer Class forms. Build fast single instruction stream processors; FORTRAN established.	FORTRAN; LARC, STRETCH (61), plus 7090 and CDC 1604 workhorses Seymour Cray wins: CDC 6600 (64) & 7600 (71)
Mid70s-mid 90s	ICs (bipolar) ...CMOS. Vector processor Era	Intro of Cray 1, vector processor 1975 and evolution takes over using multiple processors vector XMP, YMP, C-90, T-90
Mid 80s to the present	Scalables era (commodity killer micros including "game" processors)	Scalable computers using micros: How much money? Seitz Cosmic Cube c1985, move to Intel and others. 45 companies casualties.



There have been multiple eras of programming model stability



This stability has made possible remarkable advances in science and supported national security, but the programming model transitions are tough... and we are initiating one now....

Michel McCoy, LLNL, 2012

Supercomputers evolution


- Evolution of single memory computers

***Colossus: 1943, 1944 10 produced
A Supercomputer?***



Courtesy of Burton Smith, Microsoft

***Harvard Mark I aka IBM ASCC
“I think there is a world market for maybe five
computers.” – Thomas J. Watson, Sr., 1943***




A Supercomputer?

Bletchley Park "Bombe"




ENIAC: Electronic Numerical Integrator and Computer 1946-1955; Cost \$500,000

A Supercomputer?




Courtesy of Burton Smith, Microsoft


Two seminal events: EDVAC recipe & IAS



Other early "supercomputers" Supercomputers?




Zuse Z3 (1941)



Manchester/Ferranti Mark I (1951)




Univac 1 (1951)



The IAS machines (1952)
Courtesy of Burton Smith, Microsoft


Supercomputer Mainframes: LARC

- Begun in 1955 for Livermore and delivered in 1960
- Had dual processors and decimal arithmetic
- Employed surface-barrier transistors and core memory



Courtesy of Burton Smith, Microsoft

Ferranti/Manchester Atlas c1961 (One million instructions per second)



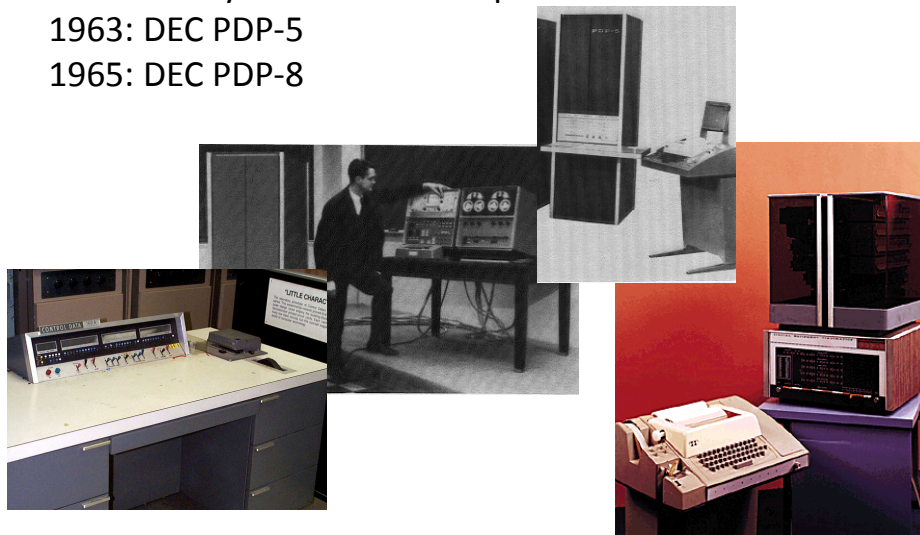
A Supercomputer?

Large transistorized calculators c1959-70
Not Supercomputers

IBM 704 >> IBM 7094
 ERA/Univac 1103 >> 1107, 1108
 CDC 1604 >> CDC 3600
 Philco 212
 Burroughs Datatron >> B5000, etc.

Trickle down

1959: CDC 160 (Offline for CDC 1604)
 1962: LINC (MIT Lincoln Laboratory)
 Laboratory Instrument Computer
 1963: DEC PDP-5
 1965: DEC PDP-8



Supercomputer Mainframes: Stretch, Harvest



- **IBM 7030 (STRETCH)**
- **Delivered to Los Alamos 4/61**
- **Pioneered in both architecture and implementation at IBM**


- **IBM 7950 (HARVEST)**
- **Delivered to NSA 2/62**
- **Was STRETCH + 4 boxes**
 - **IBM 7951 Stream unit**
 - **IBM 7952 Core storage**
 - **IBM 7955 Tape unit**
 - **IBM 7959 I/O Exchange**





Courtesy of Burton Smith, Microsoft



IBM Stretch c1961 & 360/91 c1965 consoles!



Fortran 1957, '60, ... '08



A collage of portraits and a photograph. At the top left, four portraits of men: Erich Bloch, Fred Brooks, Jr., Bob Evans, and Nick Donofrio. Below them are their signatures. To the right, text about the IBM System/360. Below that, the title '360 Revolution' and '40th Anniversary'. At the bottom, a photograph of a woman at a computer terminal, and three more portraits of men.

IBM System \360
April 6, 1964
**Converging
Calculation and
Record Keeping**

Erich Bloch, Fred Brooks, Jr., and Bob Evans
with Nick Donofrio

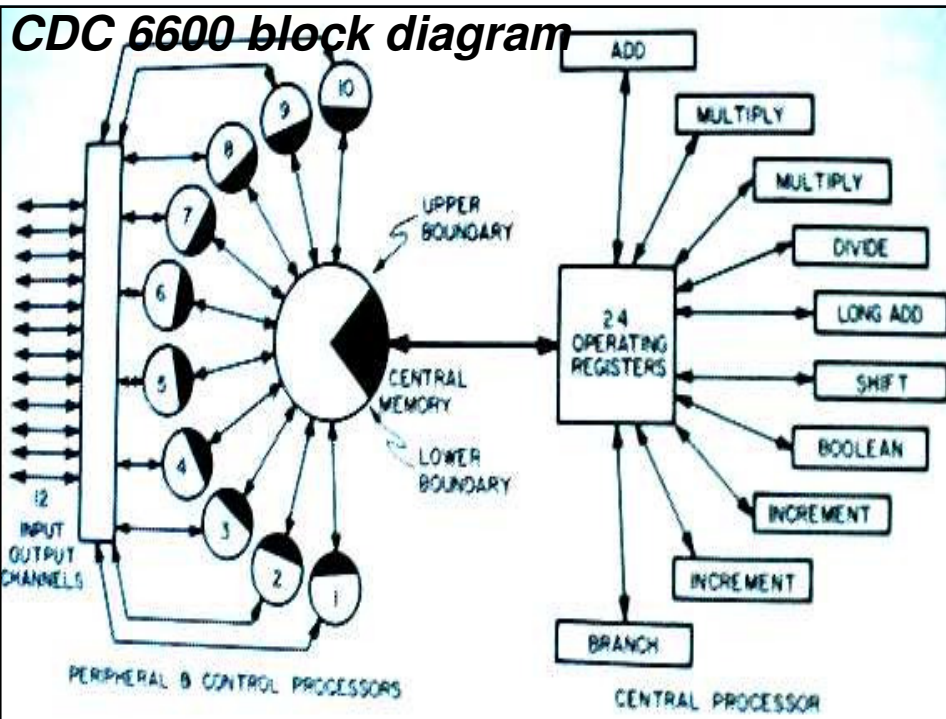
360 Revolution
40th Anniversary

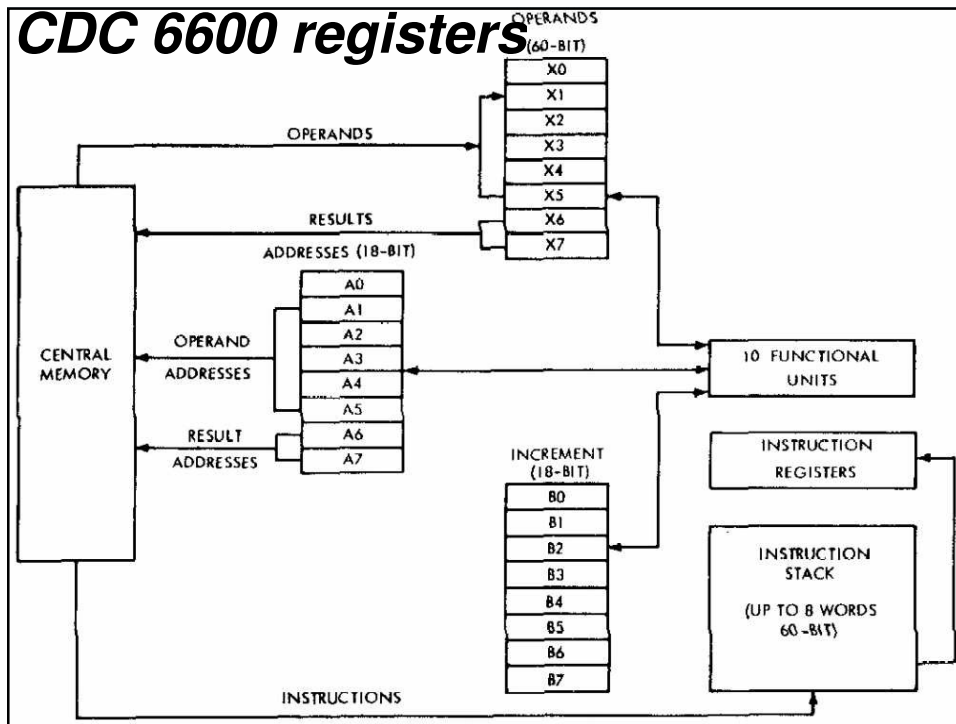
CDC 6600 Console c1964



A Supercomputer?

Courtesy of Burton Smith, Microsoft

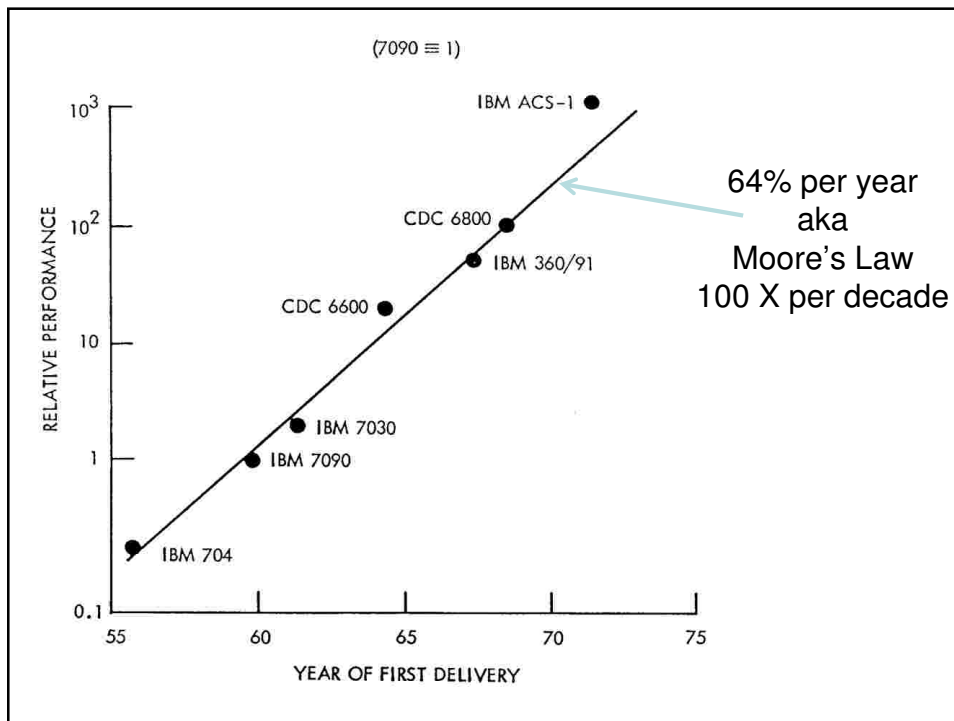
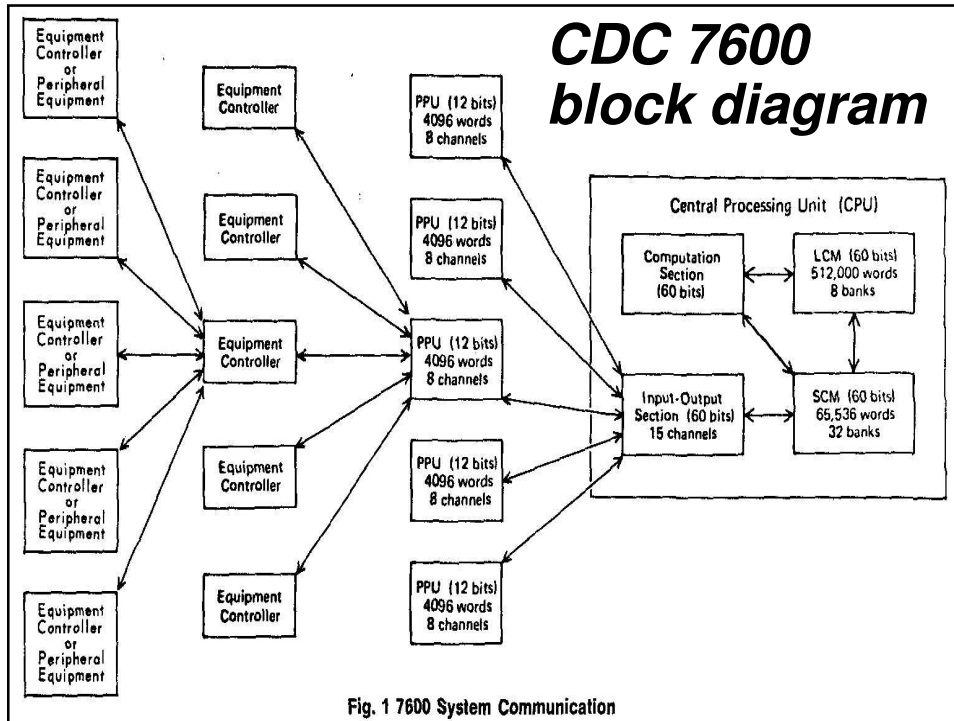




Two CDC 7600s and LLNL c1969




A Supercomputer? Courtesy of Burton Smith, Microsoft



Amdahl's law... the limit of parallelism

- If w_1 work is done at speed s_1 and w_2 at speed s_2 , the average speed s is $(w_1+w_2)/(w_1/s_1 + w_2/s_2)$
 - This is just the total work divided by the total time
- For example, if $w_1=9$, $w_2=1$, $s_1=100$, and $s_2=1$ then $s = 10/1.09 \cong 9$ (speed)
 - This is not the average of s_1 and s_2

Amdahl, Gene M, "Validity of the single processor approach to achieving large scale computing capabilities", Proc. SJCC, AFIPS Press, 1967



Courtesy of Burton Smith, Microsoft

SIMD arrays: Illiac IV

- By the late 60's, it was clear mainframes weren't enough
- To improve performance, SIMD array machines were built or proposed with many arithmetic processing units
 - Solomon was an early Westinghouse SIMD array prototype
- The Illiac IV was a U. of Illinois/Burroughs project
 - Funded by DARPA from 1964 onward, usable in 1975
 - The chief architect, Dan Slotnick, from Westinghouse
- Designed for 256 arithmetic units, cut back to 64
- The thin-film memory system was a major headache
- After student demonstrations at Illinois in May 1970, the project was moved to NASA-Ames
- Languages, especially FORTRAN, aimed to use parallel loops to express parallelism

Courtesy of Burton Smith, Microsoft

ILLIAC IV: Uof IL at NASA in 1971



- 1964 project (U. of IL)
- Burroughs contract
- SIMD 64 PEs
- 10 MB disk/PE
- Moved to NASA
- 1975 on ARPAnet

Courtesy of Burton Smith, Microsoft

CDC STAR 100 at LLNL



CDC Star 100


- Sept 74, 75
 - 25 Mhz
 - 50 Mflops
 - Mem-mem vectors
- Cyber 203, 205

=>


ETA Systems ('83-'89)

- CMOS
- Liquid nitrogen
- SMP 8 proc.

Cray-1 c1976: Supercomputer



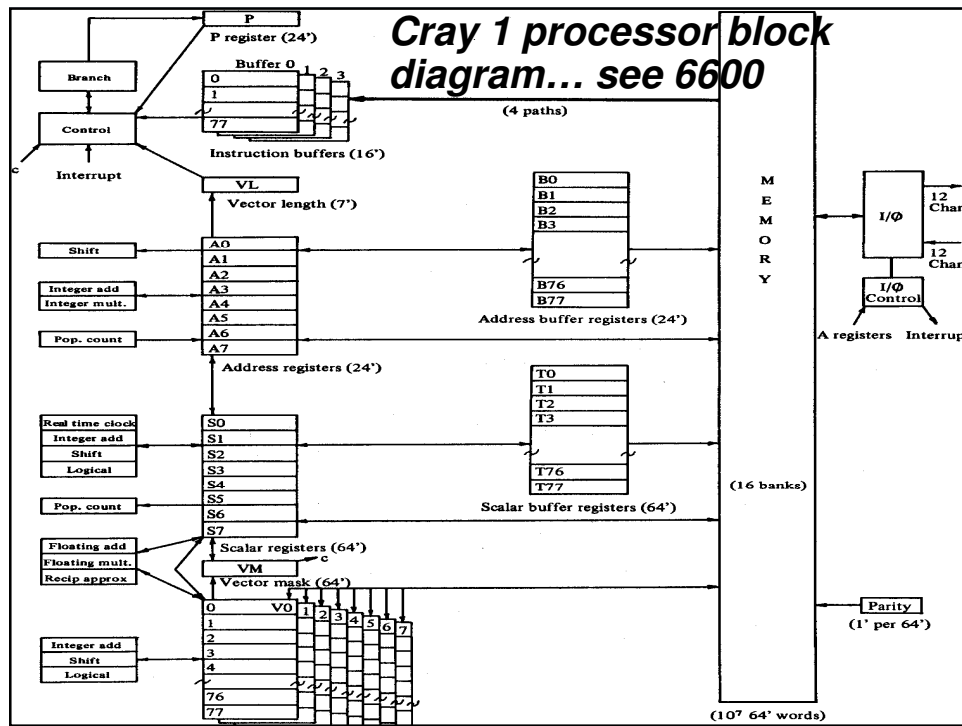
Courtesy of Burton Smith, Microsoft



Cray 1 sans covers The Vector ISA


- Unlike the CDC Star-100, no development contract
- Los Alamos got a one-year free trial. Los Alamos leased the system.
- Los Alamos developed or adapted existing software
- Cray-1 and Amdahl's law
 - Scalar performance 2X the 7600
 - Vector 160 Mflops
 - 80 MHz clock
- Peak floating point ops vs. Instructions per second
- **“Supercomputer” connotes a Cray-1**

Courtesy of Burton Smith, Microsoft




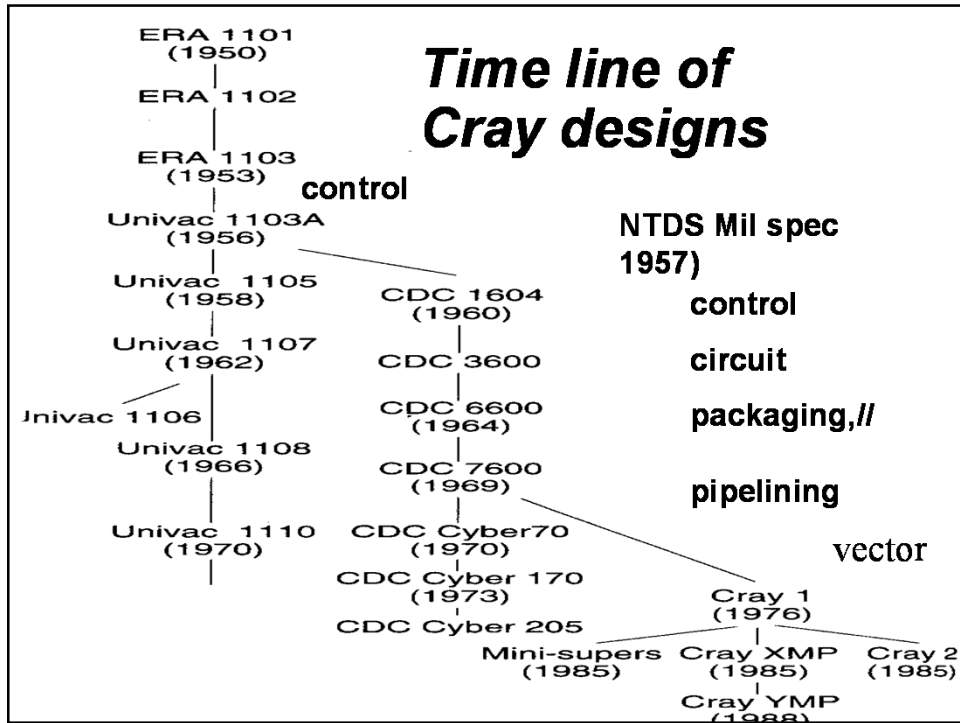
Shared Memory: Cray Vector Systems

- Cray Research, by Seymour Cray
 - Cray-1 (1976): 1 processor
 - Cray-2 (1985): up to 4 processors*
- Cray Research, not by Seymour Cray
 - Cray X-MP (1982): up to 4 procs
 - Cray Y-MP (1988): up to 8 procs
 - Cray C90: (1991?): up to 16 procs
 - Cray T90: (1994): up to 32 procs
 - Cray X1: (2003): up to 8192 procs
- Cray Computer, by Seymour Cray
 - Cray-3 (1993): up to 16 procs
 - Cray-4 (unfinished): up to 64 procs
- All are UMA systems except the X1, which is NUMA Cray-2




*One 8-processor Cray-2 was built

Courtesy of Burton Smith, Microsoft



VAXen (1977) a super-minicomputer became a “personal supercomputer”



Linpack (Mflops):
 VAX 11/780 0.14
 Cray 1 26

Clock (MHz)
 VAX 5
 Cray 1 80

VAX-11/780 c1977 32 bit super minicomputer



32 bit word. VM
 VAX = virtual
 address extension
 \$250,000
 1-2 Mbytes

Serial 3? delivered
 to CMU's John
 Pople for
 Computational
 Chemistry
As a PC.

1982: The Lax Report to NSF's NSB

- Gresham's Law: VAXen are driving out supercomputers
- NSF needs to fund supercomputer access and centers
- 1984: NSF Establishes Office of Scientific Computing
 - NCSA at U IL (1985)
 - SDSC at UCSF
 - Cornell
 - Pittsburgh
 - Von Neumann at Princeton etc.

Trickle Down: Small scale computers

- **Mini-supercomputers (Convex and Alliant)**
- **Personal supercomputers (Ardent)**


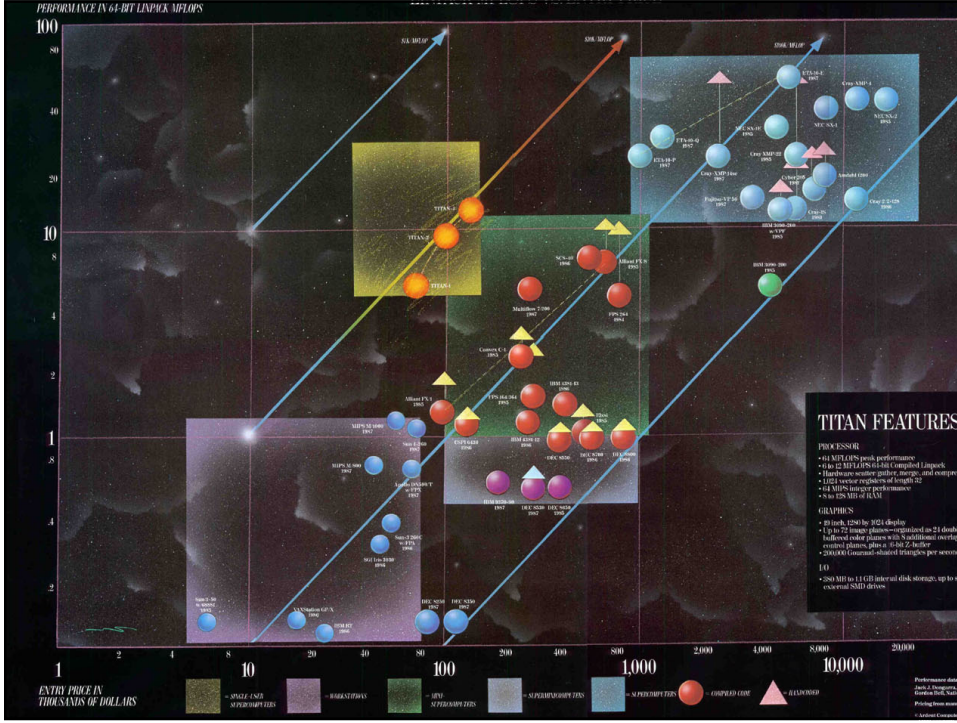


Ardent Titan Personal Graphics Supercomputer c1988



1987: The personal graphics supercomputer

Peak MIPS	32-128	
Peak Scalar MFLOPS	16-64	
Peak Vector MFLOPS	32-128	
Memory MB	32-512	
I/O MB/s	23-46	
Disk Capacity MB (internal)	380-2280	
Disk Capacity GB (external)	1-50	
Bus BW MB/s	256	
Dhrystones (KD)	51-204	
VAX MIPS	29-117	
Whetstones (MW DP)	28-112	
Linpack (MF DP)	10-12.4	
Livermore Loops (MF HM DP)	4.9-18	
Vectors/s*	100,000-300,000	
Polygons/s**	50,000-150,000	
Spheres/s	2,000	
Pixels/s	5.5-11M	
		Linpack (Mflops):
		Ardent (4) 24
		Cray 1 26
		Clock (MHz)
		Ardent 32
		Cray 1 80
		Cost \$120,000



The Transition


MIT Whirlwind Computer c1952
Back to when computers are buildings



TSF computer room power is being scaled
from 15MW to 30MW



- Capitalize on the computational efficiencies (TF/MW and SF/TF)
- Capitalize on the electrical/mechanical system efficiencies
- Adding an additional 15MW into the TSF




8

Supercomputers evolution: Supercomputers become multicomputers


- Break from single memory computers
- Factors
 - Japanese threat
 - Clusters forming independently
 - CMOS
 - Powerful microprocessors cross-over TTL and ECL
 - CS research interconnecting computers
 - Parallel compiler efforts

1982 The Japanese are Coming!



1982: Threat or opportunity

- Threat = Funding opportunity



Kickoff of DARPA's SCI program c1984...

Steve Squires, DARPA & Gordon Bell, Encore seated at our "Cray".


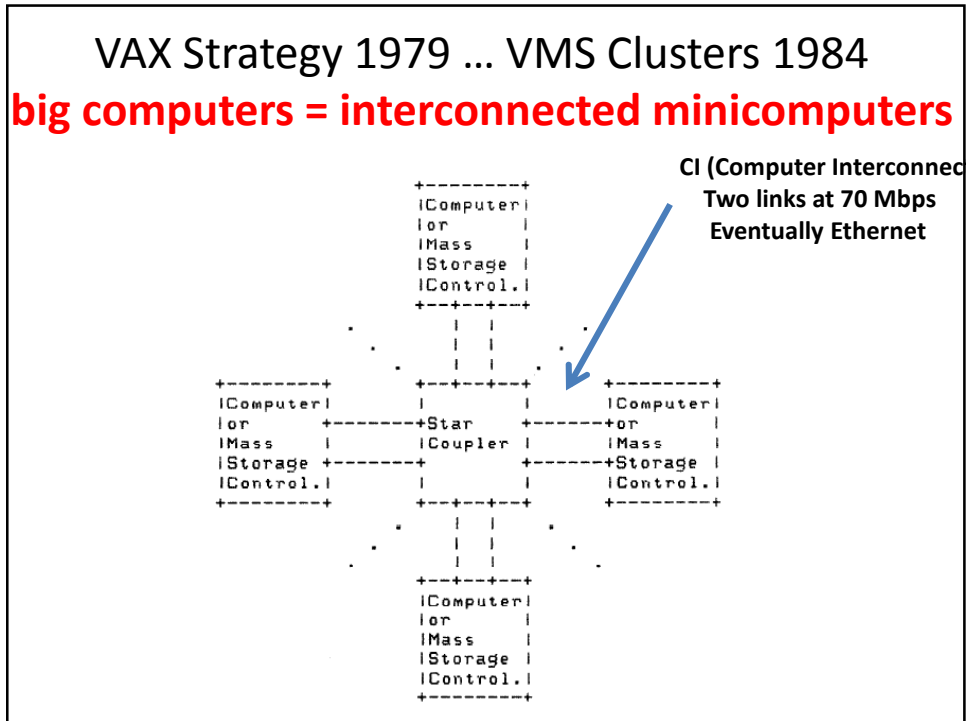
10+ years later:
"Killer micros" Clusters
begin to be standard

Lost: The search for parallelism c1983-1997

DOE and DARPA Adv. Sci Comp. Initiative

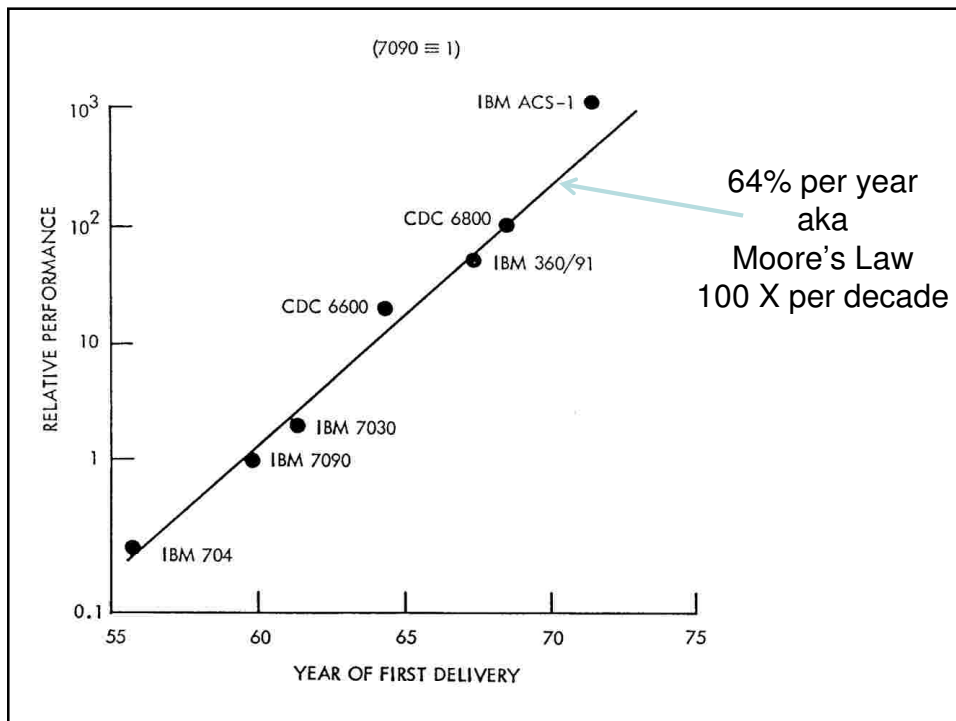
- ACRI *French-Italian program*
- Alliant *Proprietary Crayette*
- American Supercomputer
- Ametek
- Applied Dynamics
- Astronautics
- BBN
- CDC >ETA *ECL transition*
- Cogent
- Convex > HP
- Cray Computer > SRC *GaAs flaw*
- Cray Research > SGI > Cray *Manage*
- Culler-Harris
- Culler Scientific *Vapor...*
- Cydrome *VLIW*
- Dana/Ardent/Stellar/Stardent
- Denelcor
- Encore
- Eleksi
- ETA Systems *aka CDC;Amdahl flaw*
- Evans and Sutherland Computer
- Exa
- Flexible
- Floating Point Systems *SUN savior*
- Galaxy YH-1
- Goodyear Aerospace *MPP SIMD*
- Gould NPL
- Guiltech
- Intel Scientific Computers
- International Parallel Machines
- Kendall Square Research
- Key Computer Laboratories *searching again*
- MasPar
- Meiko
- Multiflow
- Myrias
- Numerix
- Pixar
- Parsytec
- nCUBE
- Prisma
- Pyramid *Early RISC*
- Ridge
- Saxpy
- Scientific Computer Systems (SCS)
- Soviet Supercomputers
- Supertek
- Supercomputer Systems
- Suprenum
- Tera > Cray Company
- Thinking Machines
- Vitesse Electronics
- Wavetracer *SIMD*

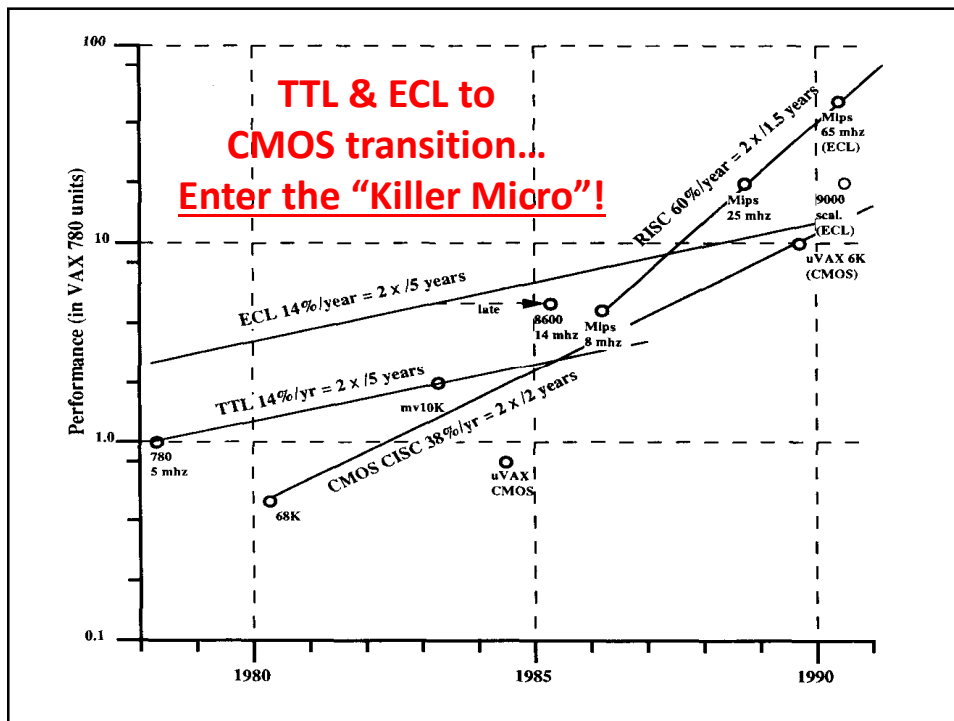






1989: The “killer micros” –Eugene Brooks, LLNL

Challenge: how do you utilize (program)
a large number of interconnected,
independent computers?






Clock speed and Moore's law



Courtesy, Jonathan Koomey


When the micro-processor's clock stopped



Amdahl's law... the limit of parallelism

- If w_1 work is done at speed s_1 and w_2 at speed s_2 , the average speed s is $(w_1+w_2)/(w_1/s_1 + w_2/s_2)$
 - This is just the total work divided by the total time
- For example, if $w_1=9$, $w_2=1$, $s_1=100$, and $s_2=1$ then $s = 10/1.09 \cong 9$ (speed)
 - This is obviously not the average of s_1 and s_2

Amdahl, Gene M, "Validity of the single processor approach to achieving large scale computing capabilities", Proc. S.



Courtesy of Burton Smith, Microsoft

Bell Prize for Parallelism, July 1987

IEEE Software launches annual Gordon Bell Award

Editor-in-Chief Ted Lewis has announced the First Annual Gordon Bell Award for the most improved speedup for parallel-processing applications. The two \$1000 awards will be presented to the person or team that demonstrates the greatest speedup on a multiple-instruction, multiple-data parallel processor.

One award will be for most speedup on a general-purpose (multiapplication) MIMD processor, the other for most speedup on a special-purpose MIMD processor. Speedup can be accomplished by hardware or software improvements, or by a combination of the two.

To qualify for the 1987 awards, candidates must submit documentation of their results by Dec. 1. The winners will be announced in the March 1988 issue. This year's judges are Alan Karp of IBM's Palo Alto Scientific Center, Jack Dongarra of Argonne National Laboratory, and Ken Kennedy of Rice University.

For a complete set of rules, definitions, and submission guidelines, write to the Gordon Bell Award, *IEEE Software*, 10662 Los Vaqueros Cir., Los Alamitos, CA 90720.

Alan Karp:
Offers \$100 for a program with 200 X parallelism by 1995.

Bell, 1987 goals:
10 X by 1992
100 X by 1997

Researcher claims:
1 million X by 2002

Development of Parallel Methods For a 1024-Processor Hypercube

John L. GUSTAFSON, Gary R. MONTRY, and Robert E. BENNER
Sandia National Laboratories, Albuquerque, New Mexico

March 1988

As printed in SIAM Journal on Scientific and Statistical Computing
Vol. 9, No. 4, July 1988, pp. 609–638.
(Minor revisions have been made for the Web page presentation of this paper. JLG 1995)

EDITOR'S NOTE

[This paper] reports on the research that was recognized by two awards, the Gordon Bell Award and the Karp Prize, at IEEE's COMPCON 1988 meeting in San Francisco on March 2.

The Gordon Bell Award recognizes the best contributions to parallel processing, either speedup or throughput, for practical, full-scale problems. Two awards were proposed by Dr. Bell: one for the best speedup on a general-purpose computer and a second for the best speedup on a special-purpose architecture. This year the two awards were restructured into first through fourth place awards because of the nature of the eleven December 1987 submissions. Bell presented the first place award of \$1,000 to the authors of [this paper].

Gustafson's Law


Benner, Gustafson, Montry winners of first Gordon Bell Prize


$$S(P) = P - \alpha \times (P-1)$$

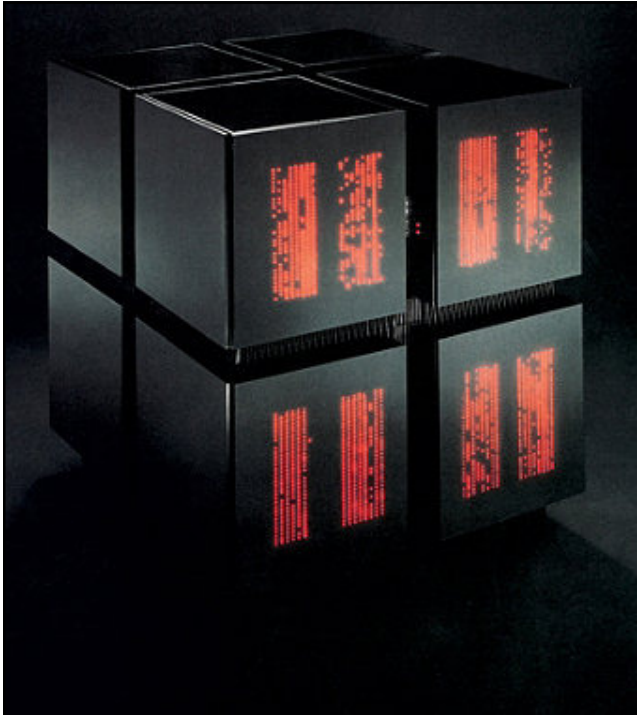
P is the number of processors,

S is the **speedup**, and

α , the non-parallelizable fraction of any parallel process








**Thinking
Machines
CM-2
1989, 1990**

**64K PE
SIMD
6, 14 Gflops**

Intel Touchstone Delta



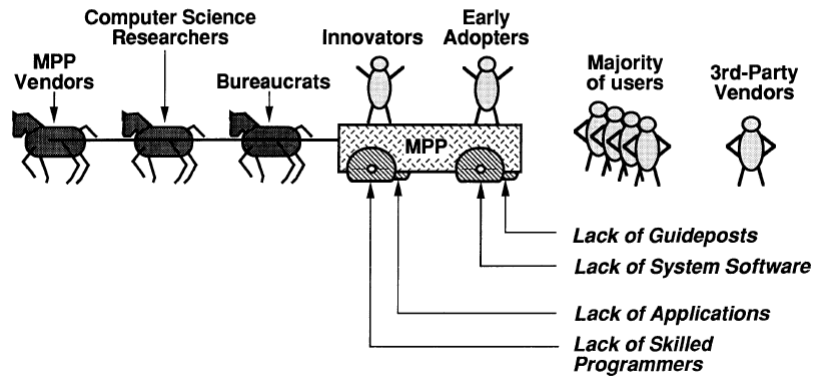
Intel Paragon and
Touchstone Delta
\$10.8 Million
5 GF (1992 Bell Prize)

30^{'90}-131^{'12/91} Gflops
8.2^{'89}- Gbytes
512^{'90}-2048^{'92} nodes




Worlton view c 1991

THE MPP BANDWAGON



CM-5 1992 1K Sparc computers Bell Prize winner



1994: Computers will All be Scalables

Thesis: **SNAP: Scalable Networks as Platforms**

- upsize from desktop to world-scale computer
- based on a few standard components

Because:

- Moore's law: exponential progress
- standards & commodities
- stratification and competition


When: Sooner than you think!

- massive standardization gives massive use
- economic forces are enormous

1994 Meeting with Jim Gray

“the day I gave up on shared memory computers”


Copyright G Bell and J Gray 1996




93


HPC & Cloud: Twins, Separated at Birth (Computation versus Storage Centric)

HPC: Separate Storage Area Network, two switches



Cloud: Attached Storage, single switch





Beowulf:
Computer Cluster
by Don Becker &
Tom Sterling,
NASA 1994

BSD, LINUX, Solaris,
and Windows Support
for MPI and PVM


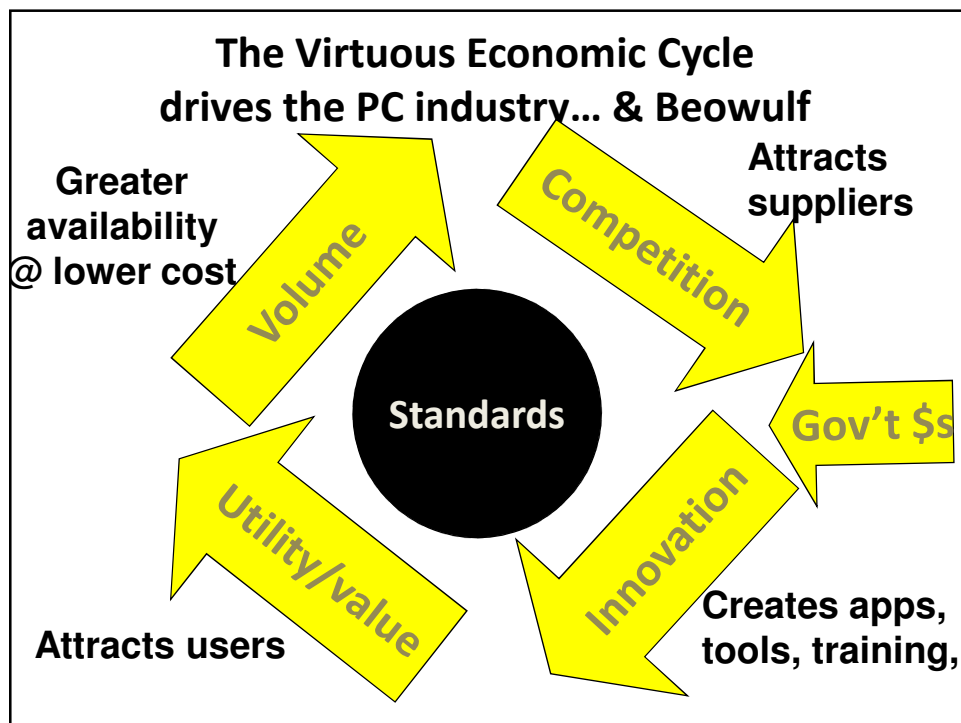


Figure 4-6. Typical Beowulf cluster configuration.

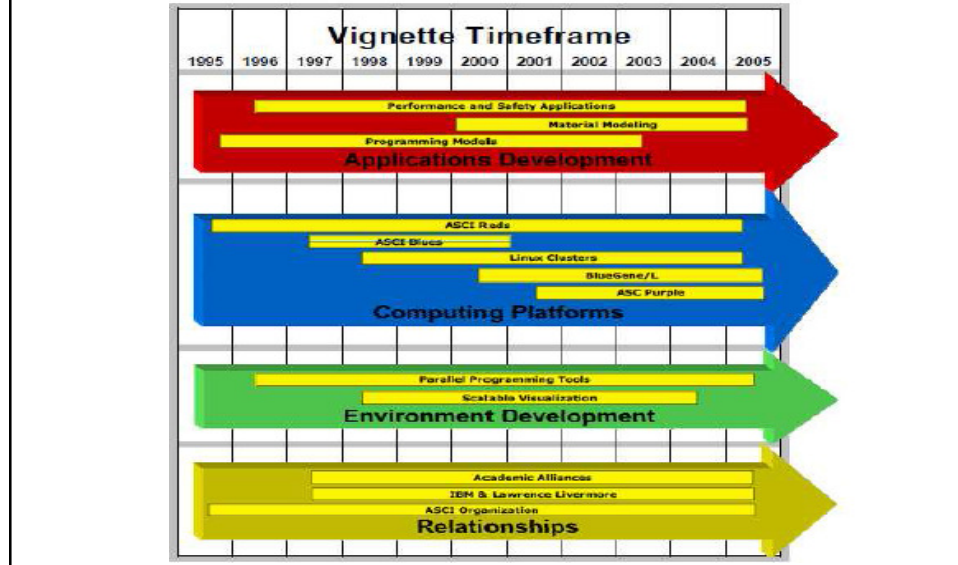
Lessons from Beowulf

- An experiment in parallel computing systems '92
- Established vision- low cost high end computing
- Demonstrated effectiveness of PC clusters for some (not all) classes of applications
- Provided networking software
- Provided cluster management tools
- Conveyed findings to broad community
- Tutorials and the book
- Provided design standard to rally community!
- Standards beget: books, trained people, software ... virtuous cycle that allowed apps to form
- Industry began to form beyond a research project



ASCI: Accelerated Strategic Computing Initiative =>ASC: Advanced Simulation and Computing

from Alex R. Larzelere II, History of ASCI, 1995-2005




ASCI Red 1997-2005




Figure 4-1. The ASCI Red system at Sandia.

- January 1997
- 1.338 Tflops
- 6/1997
- Sandia NL
- 9,216 proc, 640 disks, 1,540 PS, 616 intercon

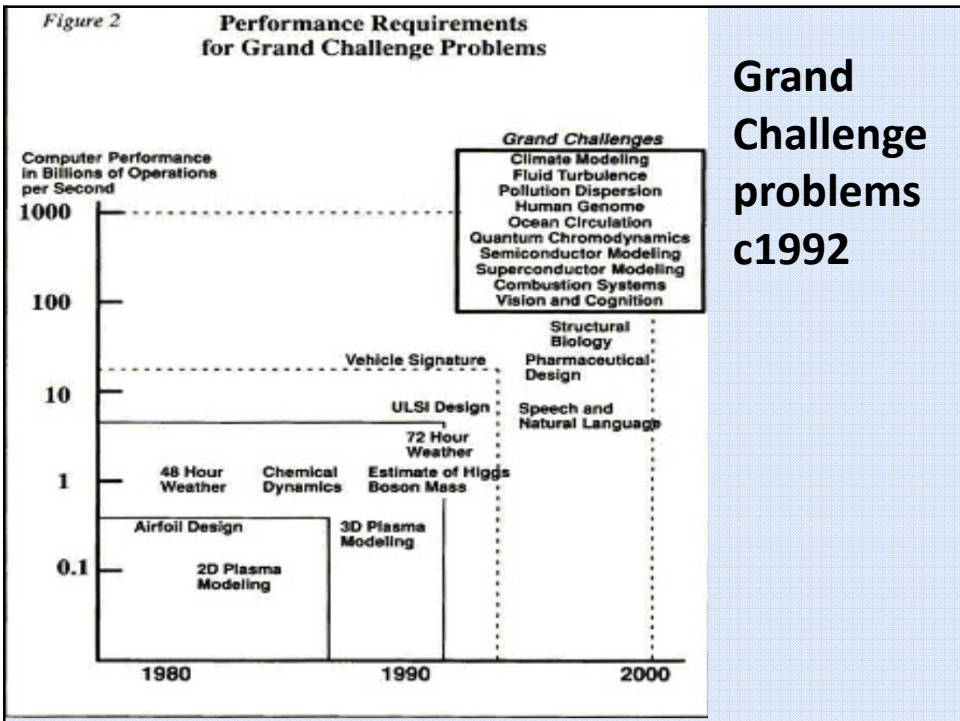
*First Clusters RLX Startup c2002
Defines blade...*



Japanese Earth Simulator (NEC)
2002 35 Teraflops 5,000 vector processor

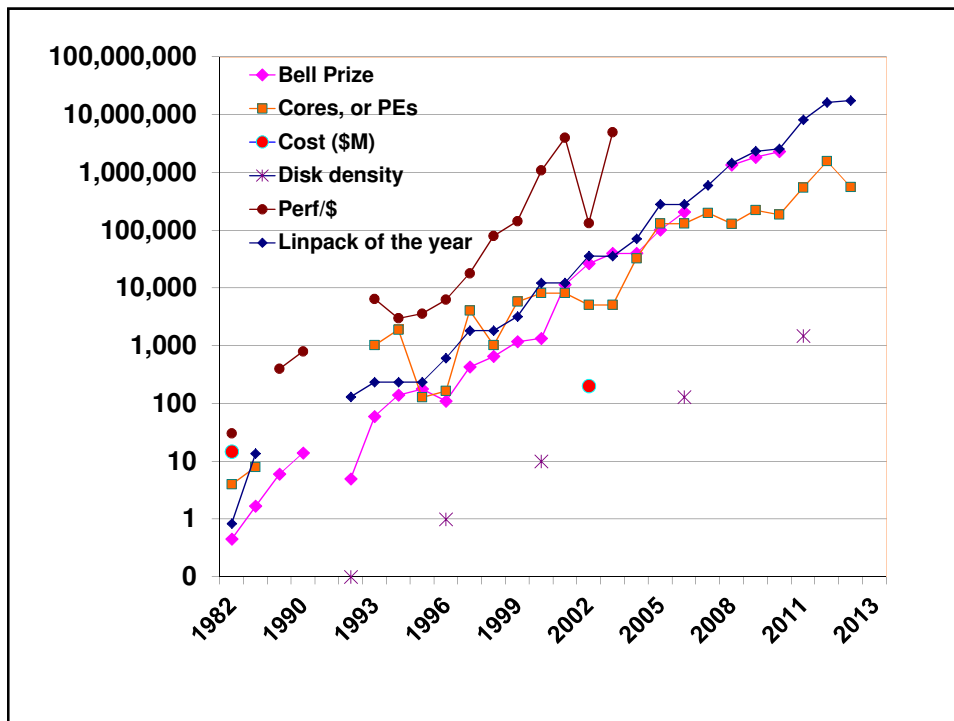


Grand Challenges, Top(x), Prizes, Benchmarks, Kernels, and Threats



Top 500 1993-2012


- Cray [Titan](#) [US](#), November 2012 - present)
- IBM [Sequoia Blue Gene/Q](#) [US](#), June 2012 – November 2012)
- Fujitsu [K computer](#) [Japan](#), June 2011 – June 2012)
- NUDT [Tianhe-1A](#) [China](#), November 2010 – June 2011)
- Cray [Jaguar](#) [US](#), November 2009 – November 2010)
- IBM [Roadrunner](#) [US](#), June 2008 – November 2009)
- IBM [Blue Gene/L](#) [US](#), November 2004 – June 2008)
- NEC [Earth Simulator](#) v [Japan](#), June 2002 – November 2004)
- IBM [ASCI White](#) [US](#), November 2000 – June 2002)
- Intel [ASCI Red](#) [US](#), June 1997 – November 2000)
- Hitachi [CP-PACS](#) v [Japan](#), November 1996 – June 1997)
- Hitachi [SR2201](#) v [Japan](#), June 1996 – November 1996)
- Fujitsu [Numerical Wind Tunnel](#) v [Japan](#), November 1994 – June 1996)
- Intel [Paragon XP/S140](#) [US](#), June 1994 – November 1994)
- Fujitsu [Numerical Wind Tunnel](#) v [Japan](#), November 1993 – June 1994)
- TMC [CM-5](#) [US](#), June 1993 – November 1993)
- Intel Touchstone Delta



Some Benchmarks & Kernels

- Whetstones (1972) & Dhrystones
- Spec, Specint, Specfp (1988)
- Livermore Loops (1986)
- TPM (1993)
- Linpack (1993)
- NAS parallel
- Graph (2011)
- Berkeley Dwarf Kernel/App (2011)

Livermore Loops c1986 (e.g.13) Average & Harmonic Mean



Some rules for supercomputer design

1. **Performance, performance, and performance are the three objective criteria for a supercomputer design.**
2. **Amdahl's law generalized implies that everything matters... a variant of "no chain is stronger than its weakest link",** especially when measuring links by harmonic mean of a set of benchmark kernels.
3. The scalar speed matters most and a super must be the fastest of comparable computers in its class... *otherwise the harmonic mean measurement kills it as a super.*
4. The vector speed can be arbitrarily high as costs allow. This is the advertised speed of the computer. The past rule of thumb is to have a vector unit which will produce two results per clock tick. Large increases over the scalar speed beyond a hundred provide a small benefit except for selected applications, making the computer, special purpose (e.g. a Connection Machine). The vector (peak) or advertised speed is the speed which the manufacturer guarantees the computer will not exceed on any application.
5. Allow no holes in the performance space (e.g. arithmetic function, input-output, mass storage) into which a benchmark can step, resulting in large performance losses.
6. Provide peaks in the performance space in order that extraordinary performance for a benchmark will result. Use this single number to advertise (characterize) the machine and to challenge other machines.
7. See Law #1: Provide enough address bits for a decade of constant architecture implementation.
8. Build at least two generations of the architecture. No first design supercomputer has ever been perfect. Do it again after the first one.
9. Build on the work of others. Designing a super is hard. Understand exactly why and how every machine works and move forward using this knowledge and any residual software.
10. Make it easy to use. Have a great compiler and diagnostic tools to aid users in vectorization and parallelization. Training in academe is nil since computer science departments are not oriented to training people to use computers or deal with computers that produce numbers. Texts lacking re. programming a parallelism.
11. **Have lots of resources when embarking on a supercomputer design.** The fatality rate for companies making machines is at least 50%, and even though a design may be good, it has to be re-iterated.

Dwarf Popularity (Red Hot → Blue Cool)




	Embed	SPEC	DB	Games	ML	HPC
1 Finite State Mach.	Red	Red	Red	Yellow	Yellow	Yellow
2 Combinational	Red	Green	Green	Green	Green	Green
3 Graph Traversal	Red	Yellow	Yellow	Yellow	Yellow	Yellow
4 Structured Grid	Red	Red	Red	Red	Red	Red
5 Dense Matrix	Red	Red	Red	Red	Red	Red
6 Sparse Matrix	Red	Red	Red	Red	Red	Red
7 Spectral (FFT)	Yellow	Yellow	Yellow	Yellow	Yellow	Yellow
8 Dynamic Prog	Yellow	Yellow	Red	Red	Red	Red
9 N-Body	Yellow	Yellow	Yellow	Yellow	Yellow	Yellow
10 MapReduce	Green	Green	Red	Red	Red	Red
11 Backtrack/ B&B	Yellow	Yellow	Yellow	Yellow	Yellow	Yellow
12 Graphical Models	Yellow	Yellow	Yellow	Yellow	Yellow	Yellow
13 Unstructured Grid	Yellow	Yellow	Yellow	Yellow	Yellow	Yellow

End Challenges... Threats


HPC and Clouds:
Separated at birth

HPC & Cloud: Twins, Separated at Birth (Computation versus Storage Centric)

HPC: Separate Storage Area Network, two switches



Cloud: Attached Storage, single switch



Traditional Datacenter Builds



Monolithic design & construction


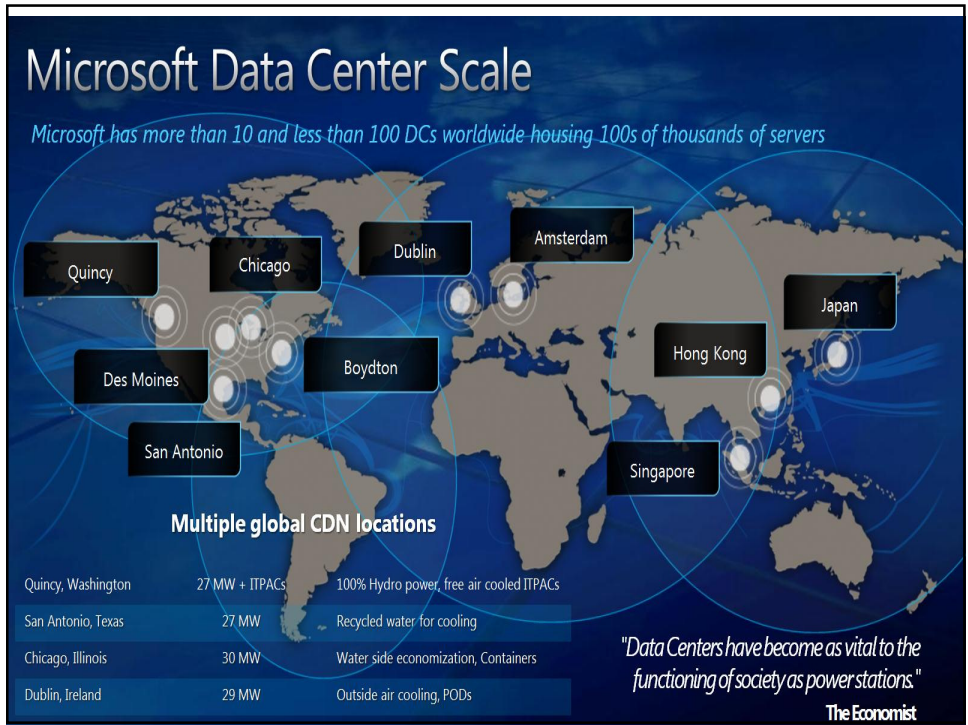
Typical large datacenter = 11 football fields
20 to 50 Megawatts

Huge \$\$\$

Typical construction costs = \$10M to \$20M
per Megawatt

Long lead time

18 to 24 months from design to online



2015: 100 petaflops; 2018: 1 exaflops

- October 31, 2012: the Chinese government announced they are building [Tianhe-2](#) a 100 petaflops supercomputer for completion in 2015.
- It also plans to have a one exaflop supercomputer online by 2018 using the [Intel MIC](#) multi-core processor architecture
- Given the current speed of progress, ... supercomputers will reach one exaflops (10^{18}) (one quintillion FLOPS) by 2018.
- SGI plans to achieve a 500 fold increase in performance by 2018, and achieve one exaflops.

The End