

---

---

Harvey Mudd College

Parallel Computing  
Introduction and Timeline

Compiled by Robert Keller

CS 156  
Spring 2011

# Sources

---

---

- <http://www.computerhistory.org/timeline>
- wikipedia
- other web sites

# 1953: Simultaneous I/O

---

---

- Introduced to absorb latency in I/O, by allowing processor to run in parallel with I/O.
- Gave rise the concept of an I/O channel.

# 1958: Parallel Programming

---

---

- S. Gill, "Parallel Programming,"  
Computer Journal, vol. 1, (1958) p. 2.
- Introduced the idea of *fork & join*
- Related concepts such as cobegin & coend and parallel-for followed.

# 1960: Instruction Pipelining, Memory Interleaving

---

---

- IBM 7030 “Stretch” 7030
- 1.2 MIPS
- about 250,000 64-bit words of memory



Stretch maintenance console

# 1964, CDC 6600: Parallelism in Instruction Streams

---

---



CDC 6600

- CDC's 6600 supercomputer, designed by Seymour Cray, performed up to 3 million instructions per second — a processing speed three times faster than that of its closest competitor, the IBM Stretch. The 6600 retained the distinction of being the fastest computer in the world until surpassed by its successor, the CDC 7600, in 1968. Part of the speed came from the computer's design, which had 10 small computers, known as peripheral processors, funneling data to a large central processing unit.

# 1964: Multics OS

---

---

- Processes
- Memory segmentation
- Memory protection

# 1966-72: SIMD Architecture Illiac IV

(originally called the “Solomon” machine)



ILLIAC IV

- The Department of Defense Advanced Research Projects Agency contracted with the University of Illinois to build a large parallel processing computer, the ILLIAC IV, which did not operate until 1972 at NASA's Ames Research Center. The first large-scale array computer, the ILLIAC IV achieved a computation speed of 200 million instructions per second, about 300 million operations per second, and 1 billion bits per second of I/O transfer via a unique combination of parallel architecture and the overlapping or "pipe-lining" structure of its 64 processing elements.

This photograph shows one of the ILLIAC's 13 Burroughs disks, the debugging computer, the central unit, and the processing unit cabinet with a processing element.

# 1966: Flynn's Taxonomy

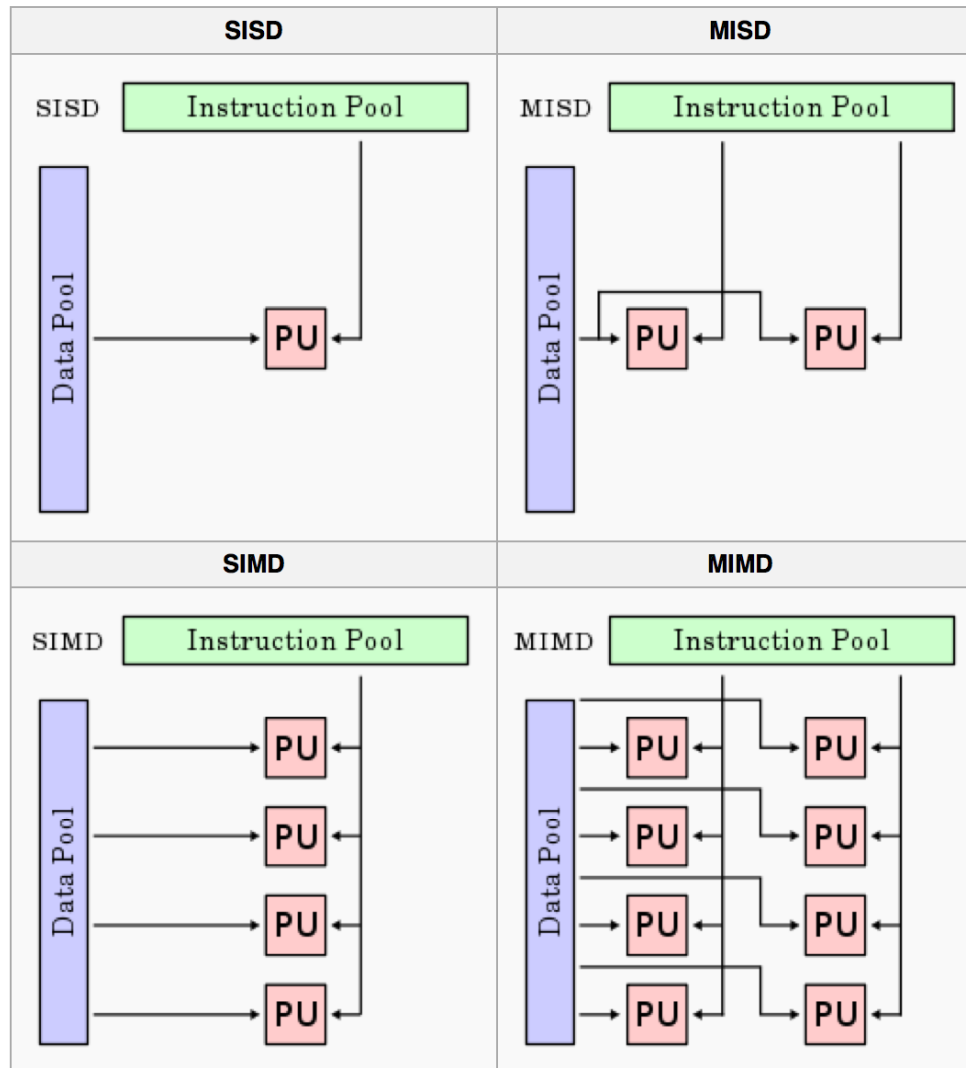
---

---

## Flynn's taxonomy

	<b>Single Instruction</b>	<b>Multiple Instruction</b>
<b>Single Data</b>	SISD	MISD
<b>Multiple Data</b>	SIMD	MIMD

# 1966: Flynn's Taxonomy



# 1967, IBM 360 model 91: Parallelism in Instruction Streams

---

---



IBM System/360

- IBM announced the System/360, a family of six mutually compatible computers and 40 peripherals that could work together. The initial investment of \$5 billion was quickly returned as orders for the system climbed to 1,000 per month within two years. At the time IBM released the System/360, the company was making a transition from discrete transistors to integrated circuits, and its major source of revenue moved from punched-card equipment to electronic computer systems.

# 1971: Unix

---

---

- process fork
- wait rather than join
- pipes on command line

# 1974-79: ICL DAP


---

---

The Distributed Array Processor (DAP) produced by International Computers Limited (ICL) was the world's first commercial massively parallel computer. The original paper study was complete in 1972 and building of the prototype began in 1974.

The ICL DAP had 64x64 single bit processing elements (PEs) with 4096 bits of storage per PE. It was attached to an ICL mainframe and could be used as normal memory. Programs for the DAP were written in DAP FORTRAN which was FORTRAN extended with 64x64 matrix and 64 element vector primitives. It had a Single Instruction Multiple Data<sub>01-13</sub> (SIMD) architecture.


# 1975: Tandem Fault Tolerance



Tandem-16

- Tandem computers tailored its Tandem-16, the first fault-tolerant computer, for online transaction processing. The banking industry rushed to adopt the machine, built to run during repair or expansion.

# 1976: Cray Vector Processor



Cray I

- The Cray I made its name as the first commercially successful vector processor. The fastest machine of its day, its speed came partly from its shape, a C, which reduced the length of wires and thus the time signals needed to travel across them.

Project started: 1972

Project 1976  
completed:

Speed: 166 million floating-point  
operations per second

Size: 58 cubic feet

Weight: 5,300 lbs.

Technology: Integrated circuit

Clock rate: 83 million cycles per second

Word length: 64-bit words

Instruction set: 128 instructions

# 1982: Parallel Vector Processor

---

---

- The Cray XMP, first produced in this year, almost doubled the operating speed of competing machines with a parallel processing system that ran at 420 million floating-point operations per second, or megaflops. Arranging two Crays to work together on different parts of the same problem achieved the faster speed. Defense and scientific research institutes also heavily used Crays.

# Cray / Apple

---

---

- 'Apple Computer purchased a big Cray supercomputer in the mid-1980s. In fact, Steve Jobs was Cray's first and only walk-in customer. He arrived unannounced at Cray headquarters in Mendota Heights, Minnesota and asked to speak to someone about buying a Cray. They nearly threw him out.
- 'Later, Cray president John Rollwagen phoned Seymour and told him that Apple had just purchased a Cray that would be used in designing the next Macintosh. Seymour thought for a bit, and replied that it seemed reasonable, since he was using a Macintosh to design the next Cray!'

# 1979-1983: Goodyear MPP

## (Massively Parallel Processor)

---

---

Based on Goodyear's earlier STARAN array processor, a 4x256 1-bit processing element (PE) computer.

The MPP was a 128x128 2-dimensional array of 1-bit wide PEs.

# 1980's: Sequent SMP

---

---

## Balance

Sequent's first computer systems were the Balance 8000 and Balance 21000 released in 1984.<sup>[5]</sup> The Balance 21000 included up to twenty 8 MHz **National Semiconductor NS32016** processors (in multiples of 2), each with a small write-through cache connected to a common memory to form a **shared memory** system with **SCSI** and **Ethernet**. The systems ran a modified version of **4.2BSD Unix** the company called **DYNIX**, for DYNAmic unIX. The machines were designed to compete with the **DEC VAX-11/780**, with all of their inexpensive processors available to run any process. In addition the system included a series of libraries that could be used by programmers to develop applications that could use more than one processor at a time. The Balance systems were originally intended to be sold to OEMs as computing

# 1980's: Sequent SMP

(example: HMC's jarthur)

---

---

## Symmetry

Their next series was the [Intel 80386](#)-based Symmetry, released in 1987. Various models supported between 2 and 30 processors, using a new [copy-back cache](#) and a wider [64-bit memory bus](#). 1991's Symmetry 2000 models added multiple [SCSI](#) boards, and were offered in versions with from one to six [Intel 80486](#) processors. The next year they added the [VMEbus](#) based Symmetry 2000/x50 with faster CPUs.

The late 1980s and early 1990s saw big changes on the software side for Sequent. DYNIX was replaced by DYNIX/ptx, which was based on a merger of [AT&T's UNIX System V](#) and [4.2BSD](#). And this was during a period when Sequent's high-end systems became particularly successful due to a close working relationship with [Oracle](#), specifically [their high-end database servers](#). In 1993 they added the

# 1983: Hypercube Interconnect

(example: HMC's Henry, 2nd floor Sprague)

---

---

- Cosmic Cube at Caltech
- Used 8086 chips
  
- iPSC at Intel
- Used 8086 chips, multiple ethernet cards to interconnect
  
- nCUBE corp.: 1024 custom architecture

# 1986: Connection machine

---

---



Connection Machine

- Daniel Hillis of Thinking Machines Corp. moved artificial intelligence a step forward when he developed the controversial concept of massive parallelism in the Connection Machine. The machine used 16,000 processors and could complete several billion operations per second. Each processor had its own small memory linked with others through a flexible network that users could alter by reprogramming rather than rewiring.

The machine's system of connections and switches let processors broadcast information and requests for help to other processors in a simulation of brainlike associative recall. Using this system, the machine could work faster than any other at the time on a problem that could be parceled out among the many processors.

# 1988: Posix Threads

---

---

- POSIX = Portable Operating System Interface [for Unix]
- Standardized multithreading in C


# 1987: BBN Butterfly

(BBN = “Bolt, Beranek, and Newman”)

---

---

- Butterfly interconnect
- NUMA (non-uniform memory access) architecture



# 1987: Parallel Inference Machine

---

---

## Japanese 5th Generation Project



# 1989: PVM

---

---

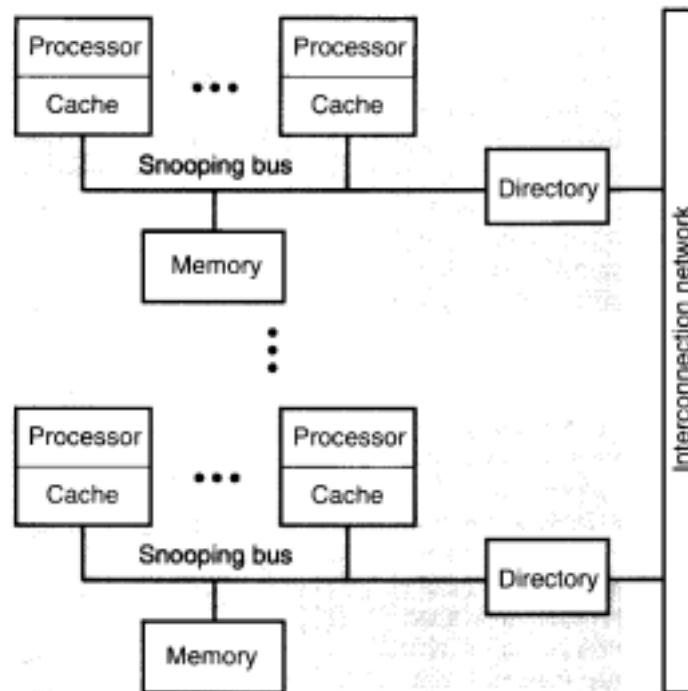
- PVM = Parallel Virtual Machine
- Developed at Oak Ridge National Laboratory (ORNL)

# 1991: Directory-Based Cache Coherence

---

---

## Stanford DASH



# 1991: CM-5

---


---

With the CM-5, announced in 1991, Thinking Machines switched from the CM-2's hypercubic architecture of simple processors to an entirely new MIMD architecture based on a fat tree network of SPARC RISC processors.

# 1993: NOW (Network of Workstations) “Beowulf” Clusters

---

---



John Koza's 1000 node cluster, 1999  
Used for Genetic Programming

# 1994: MPI

---

---

- MPI = “Message Passing Interface”
- Developed at Argonne National Laboratory
- Became the de facto standard for distributed multiprocessing

# 1995: Cray T3E-1200

---

---

- Distributed memory
- 32 to 2048 processing elements (PEs)
- 8 Million words (64 MB) per PE
- PE: DEC Alpha, 1200 Mflops
- 2.4 teraflops maximum
- 3-D torus interconnect
- interprocessor comm. rate: 500 MB/s
- starting at the paltry sum of \$630,000



# 1995: Java

---

---

- Integral multithreading
- Developed at SUN Microsystems

# 1996: Deep Blue: Chess Playing

---

---

Massively parallel, RS/6000 SP Thin P2SC-based system with 30-nodes, with each node containing a 120 MHz P2SC microprocessor for a total of 30, enhanced with 480 special purpose VLSI chess chips.



[http://en.wikipedia.org/wiki/IBM\\_Deep\\_Blue](http://en.wikipedia.org/wiki/IBM_Deep_Blue)

# 1997: ASCI Red



- Intel i4640 nodes
- Housed at Sandia National Laboratory
- 666 MOPs peak per node
- 606 GB memory

# 1997: Distributed Apps

---

---

- distributed.net (Jeff Lawson '99, et al.)


# 2003: IBM SP3

---

---

- 1-8 processors, shared memory
- up to 16 GB memory
- networkable to 512 processors
- PE: IBM RS/6000
- 2-level coherent cache
- 14.2 GB/s crosspoint switch
- Cost for 128 PE version:

only \$8,614,441



# 2006: SGI Origin 3000 series

---

---

- 16-512 processors
- up to 1 TB shared memory
- R12000 processors, 300 MHz
- NUMA architecture
- 38.4 Gflops for 64 processors
- 716 GB/s system bandwidth
- Cost starting around \$300k



# 2005-2008:

## Red Storm (Cray-Designed) at Sandia

Operational Time Frame	2005	2006	2008
Theoretical Peak (TF)	41.47	124.42	284.16
HPL Performance (GF)	36190 on 10,880 processors	101400 on 26,544 processors	204200 on 38,208 processors
Architecture	distr memory MIMD	distr memory MIMD	distr memory MIMD
Number of compute processors	10,368	25,920 (12,960 nodes)	38,400 (12,960 nodes)
Number of service/IO processors	256 + 256	320+320	320+320
Processor	AMD Opteron™ @ 2.0 GHz	AMD dual core Opteron™ @ 2.4 GHz	6720 AMD dual-core Opteron™ @ 2.4 GHz 6240 AMD quad-core Opteron™ @ 2.2 GHz
Total Memory	33.38TB	39.19 TB	78.75 TB
System Memory B/W	57.97 TB/s	78.12 TB/s	126.29 TB/s
Local Disk Storage	240 TB	240 TB	1750 TB

# 2004-2011: IBM Blue Gene

---

---

The first computer in the Blue Gene series, Blue Gene/L, developed through a partnership with Lawrence Livermore National Laboratory (LLNL), originally had a theoretical peak performance of 360 TFLOPS.

...

The archetypal Blue Gene/Q system called Sequoia will be installed at Lawrence Livermore National Laboratory in 2011 as a part of the Advanced Simulation and Computing Program running nuclear simulations and advanced scientific research.

It will consist of 98,304 compute nodes comprising 1.6 million processor cores and 1.6 PB memory in 96 racks covering an area of about 3000 square feet, drawing 6 megawatts of power.

# 2007: Cray XT5 (Jaguar, Kraken)

---



## CRAY XT5

### Cray XT5 – Scalable Computing at Work

As the single most scalable Linux supercomputer, the Cray XT5™ system combines unprecedented sustained application performance with exceptional manageability and reliability, and lower cost of ownership for customers.

Using powerful AMD Opteron™ processor cores, the Cray XT5 system supports both Cray XT4™ blades for optimized compute/interconnect balance and the new Cray XT5 blades optimized for memory-intensive and/or compute-biased workloads.

[Oak Ridge National Laboratory's \(ORNL\) "Jaguar" supercomputer](#), a Cray XT5 system, is the fastest in the world for open science.

[Request a Quote](#)

---

# Moore's "Law"

Moore's original statement that transistor counts had doubled every year can be found in his publication "Cramming more components onto integrated circuits", Electronics Magazine 19 April 1965:


The complexity for minimum component costs has increased at a rate of roughly a factor of two per year...

Certainly over the short term this rate can be expected to continue, if not to increase. Over the longer term, the rate of increase is a bit more uncertain, although there is no reason to believe it will not remain nearly constant for at least 10 years. That means by 1975, the number of components per integrated circuit for minimum cost will be 65,000. I believe that such a large circuit can be built on a single wafer.

Moore slightly altered the formulation of the law over time, in retrospect bolstering the perceived accuracy of his law . Most notably, in 1975, Moore altered his projection to a doubling every two years. Despite popular misconception, he is adamant that he did not predict a doubling "every 18 months". However, David House, an Intel colleague, had factored in the increasing performance of transistors to conclude that integrated circuits would double in performance every 18 months.

[http://en.wikipedia.org/wiki/Moore's\\_Law](http://en.wikipedia.org/wiki/Moore's_Law)

# Moore's Law and Processor Speeds



# 2007: Nvidia GPUs and CUDA

---

---

On high-end NVIDIA video cards, upwards of 300 million keys/second (on a code-breaking problem) has been reported. Considering a very high end single CPU may achieve 50 million keys/second, the CUDA advancement represents a performance increase of roughly 500%.

# 2010 (Nov.): Current Top 500 SC Sites

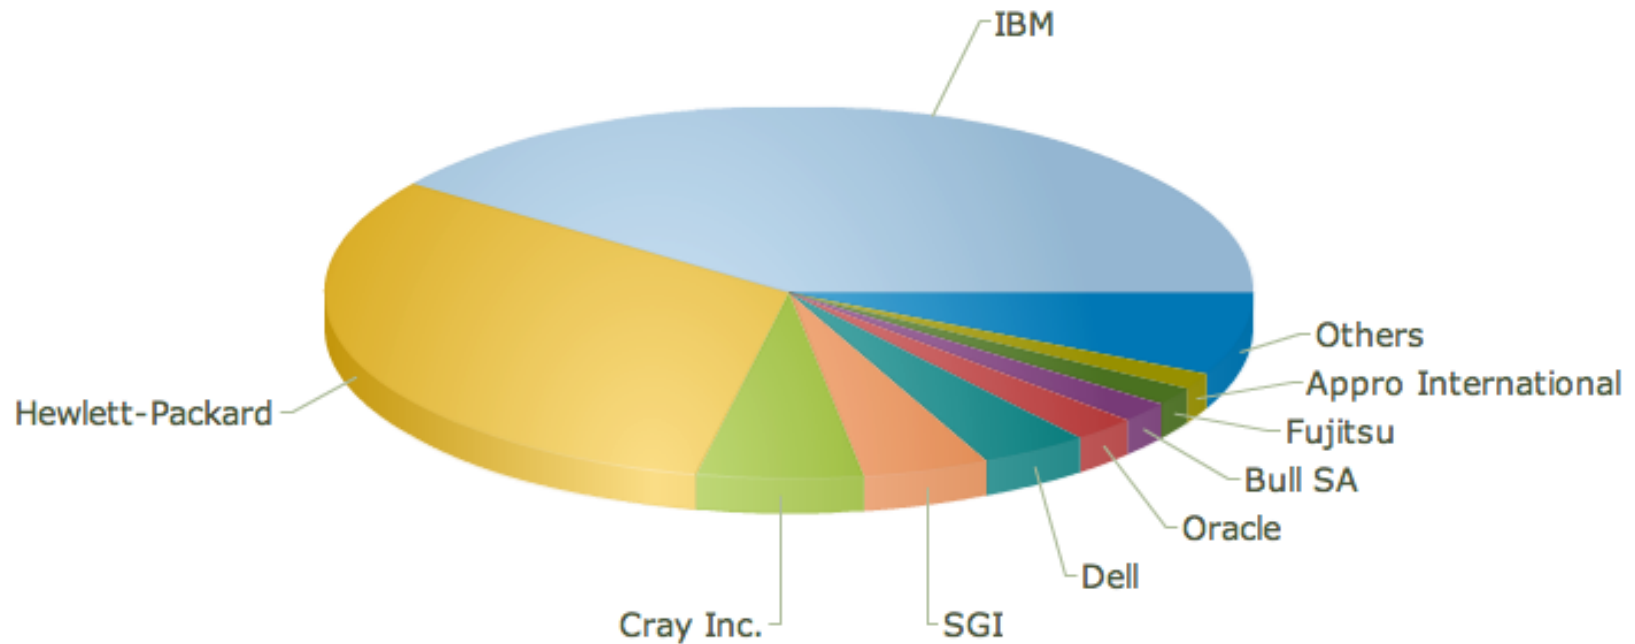
<http://www.top500.org/>

Rank	Site	Computer
1	National Supercomputing Center in Tianjin China	Tianhe-1A - NUDT TH MPP, X5670 2.93Ghz 6C, NVIDIA GPU, FT-1000 8C NUDT
2	DOE/SC/Oak Ridge National Laboratory United States	Jaguar - Cray XT5-HE Opteron 6-core 2.6 GHz Cray Inc.
3	National Supercomputing Centre in Shenzhen (NSCS) China	Nebulae - Dawning TC3600 Blade, Intel X5650, NVidia Tesla C2050 GPU Dawning
4	GSIC Center, Tokyo Institute of Technology Japan	TSUBAME 2.0 - HP ProLiant SL390s G7 Xeon 6C X5670, Nvidia GPU, Linux/Windows NEC/HP
5	DOE/SC/LBNL/NERSC United States	Hopper - Cray XE6 12-core 2.1 GHz Cray Inc.
6	Commissariat a l'Energie Atomique (CEA) France	Tera-100 - Bull bullx super-node S6010/S6030 Bull SA
7	DOE/NNSA/LANL United States	Roadrunner - BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Voltaire Infiniband IBM
8	National Institute for Computational Sciences/University of Tennessee United States	Kraken XT5 - Cray XT5-HE Opteron 6-core 2.6 GHz Cray Inc.
9	Forschungszentrum Juelich (FZJ) Germany	JUGENE - Blue Gene/P Solution IBM
10	DOE/NNSA/LANL/SNL United States	Cielo - Cray XE6 8-core 2.4 GHz Cray Inc.

# Slicing the Top 500 Pie by vendor

---


---



# Slicing the Top 500 Pie by country

---


---



# Slicing the Top 500 Pie by processor family

---


---



# Slicing the Top 500 Pie by interconnect family

---

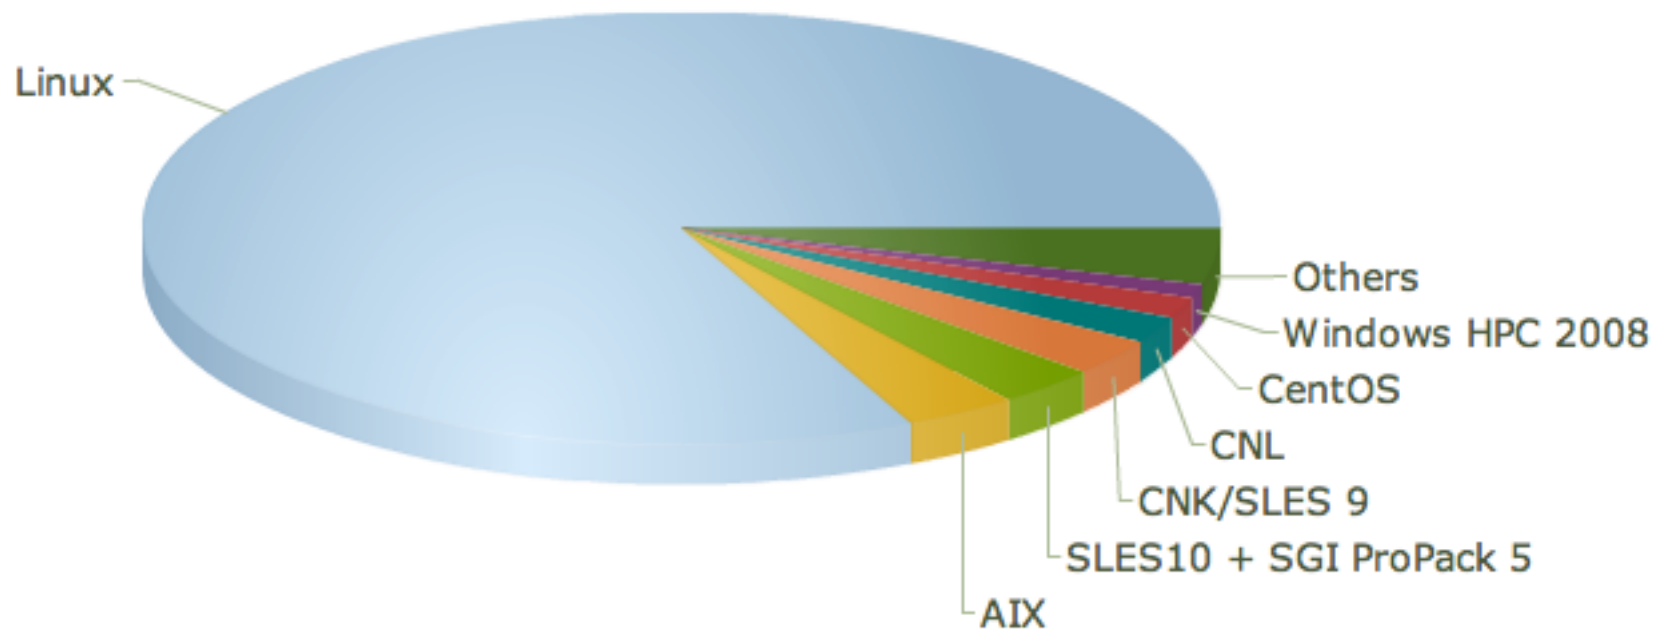
---




# Slicing the Top 500 Pie by operating system

---


---



# Vendor Share vs. Time




# Country vs. Time



**Figure 4: Supercomputer installations worldwide**

[http://www.top500.org/files/TOP500\\_Looking\\_back\\_HWM.pdf](http://www.top500.org/files/TOP500_Looking_back_HWM.pdf)

# Architecture vs. Time



**Figure 17: Architectures / Systems**

---

---


# Example Applications

# Rayleigh-Taylor (Turbulence) Simulation

<http://www.llnl.gov/casc/asciturb/simulations.shtml>

---

A parallelized compressible PPM (Piecewise Perturbation Method) code used to simulate the Rayleigh-Taylor instability and turbulent mixing. The domain is a unit cube spanned by a grid containing 512 points in each of the three directions. This case was run on the ASCI Blue-Pacific ID System at LLNL using 128 nodes.

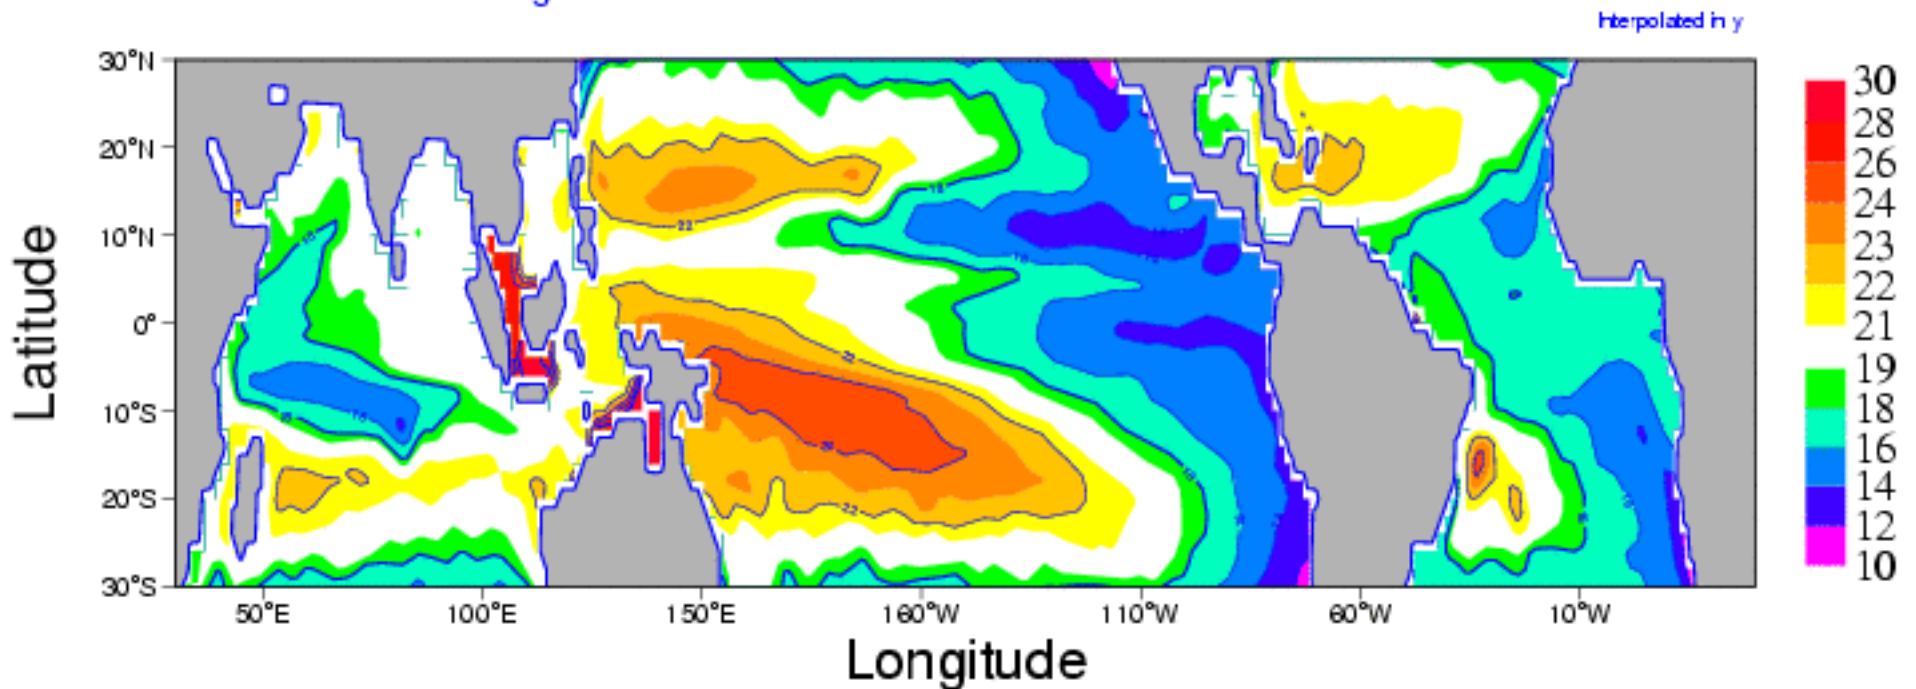


# Global Weather Forecasts

<http://www.ecmwf.int>

ECMWF S3 ocean analysis  
Averaged Temperature upper 300m  
Contour interval = 2 deg C


20110116 (1 days mean)



# Reactor Modeling

<http://www.cs.sandia.gov/CRF/MPSalsa/>

Steady-state solution for the deposition of Gallium Arsenide from GaMe3 and AsH3 in a horizontal CVD reactor with tilted susceptor and rotating substrate. Streamlines and contours of GaMe3 on the reacting surface illustrate the 3D nature of the flow. This MPSalsa calculation of 430,000 unknowns required approximately 15 minutes on 256 processors of an Intel Paragon.



# Virtual Supercomputers (grid, cloud)

---

---

## **Screensaver aids cancer fight**



PC spare capacity will be used to examine molecules  
A computer screensaver is launched on  
Tuesday that will help scientists find a cure for  
cancer.