



KUNGL
TEKNISKA
HÖGSKOLAN

ParallelDatorCentrum - PDC
Enabling Science



Computational Science and Engineering in the 21st Century

Lennart Johnsson
Numerical analysis and Computing Sciences,
KTH
Department of Computer Science,
University of Houston



Summer School in High Performance Computing

- 46 participants
 - 16 International Masters
 - 15 PhD Students
 - 2 NGSSC
 - 2 PSCI
 - 11 KTH Students

- 34 from KTH
- 4 from CTH
- 2 from Lund University
- 2 from Stockholm University
- 2 from Uppsala University
- 1 from Linköping University
- 1 from Institute of Geonics (CZ)



Summer School in High Performance Computing

- Started 1996
- Since 1997 the Summer School is held in collaboration with the National Graduate School in Scientific Computing (NGSSC)
- Focus on parallel and high performance computing

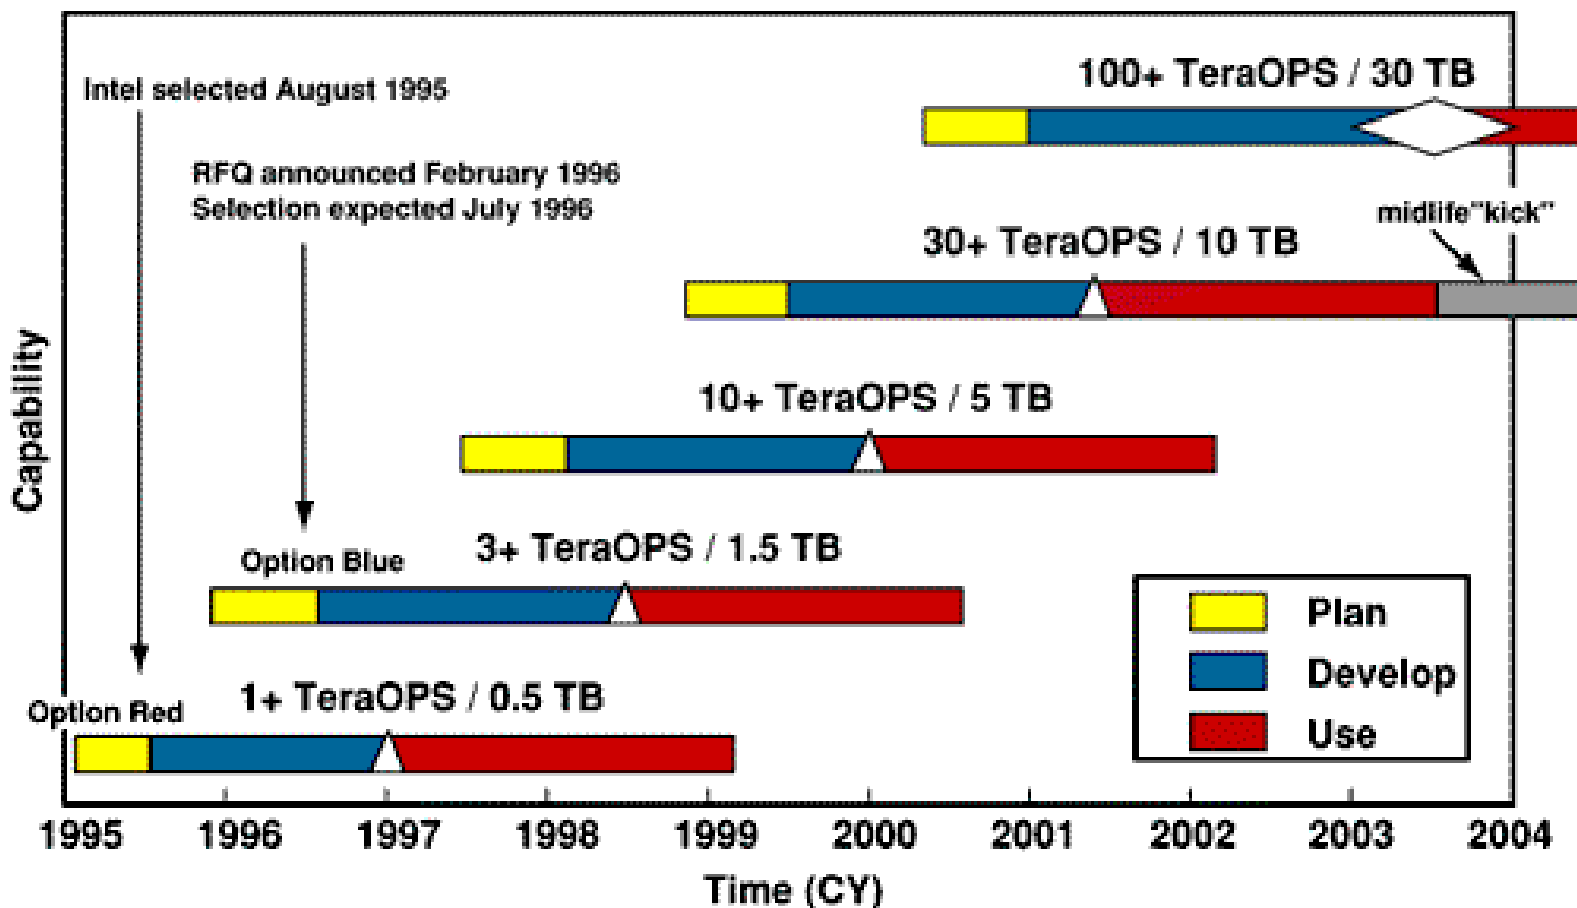


PDC Facilities

- IBM SP: 300 processors, 204 Gflops, 115 Gbytes main memory, 2TBytes of disk
- Fujitsu: 3 processors, 6.6 Gflops, 2GBytes
- SGI Onyx2: 12 processors, 3 Infinite Reality Engines, 4 GBytes
- SGI Octane, 2 processors, EMXI graphics, 1 Gbyte
- PC Cluster: PIII-864, 15 processors
- ImmersaDesk
- VR-CUBE



ASCI Platforms





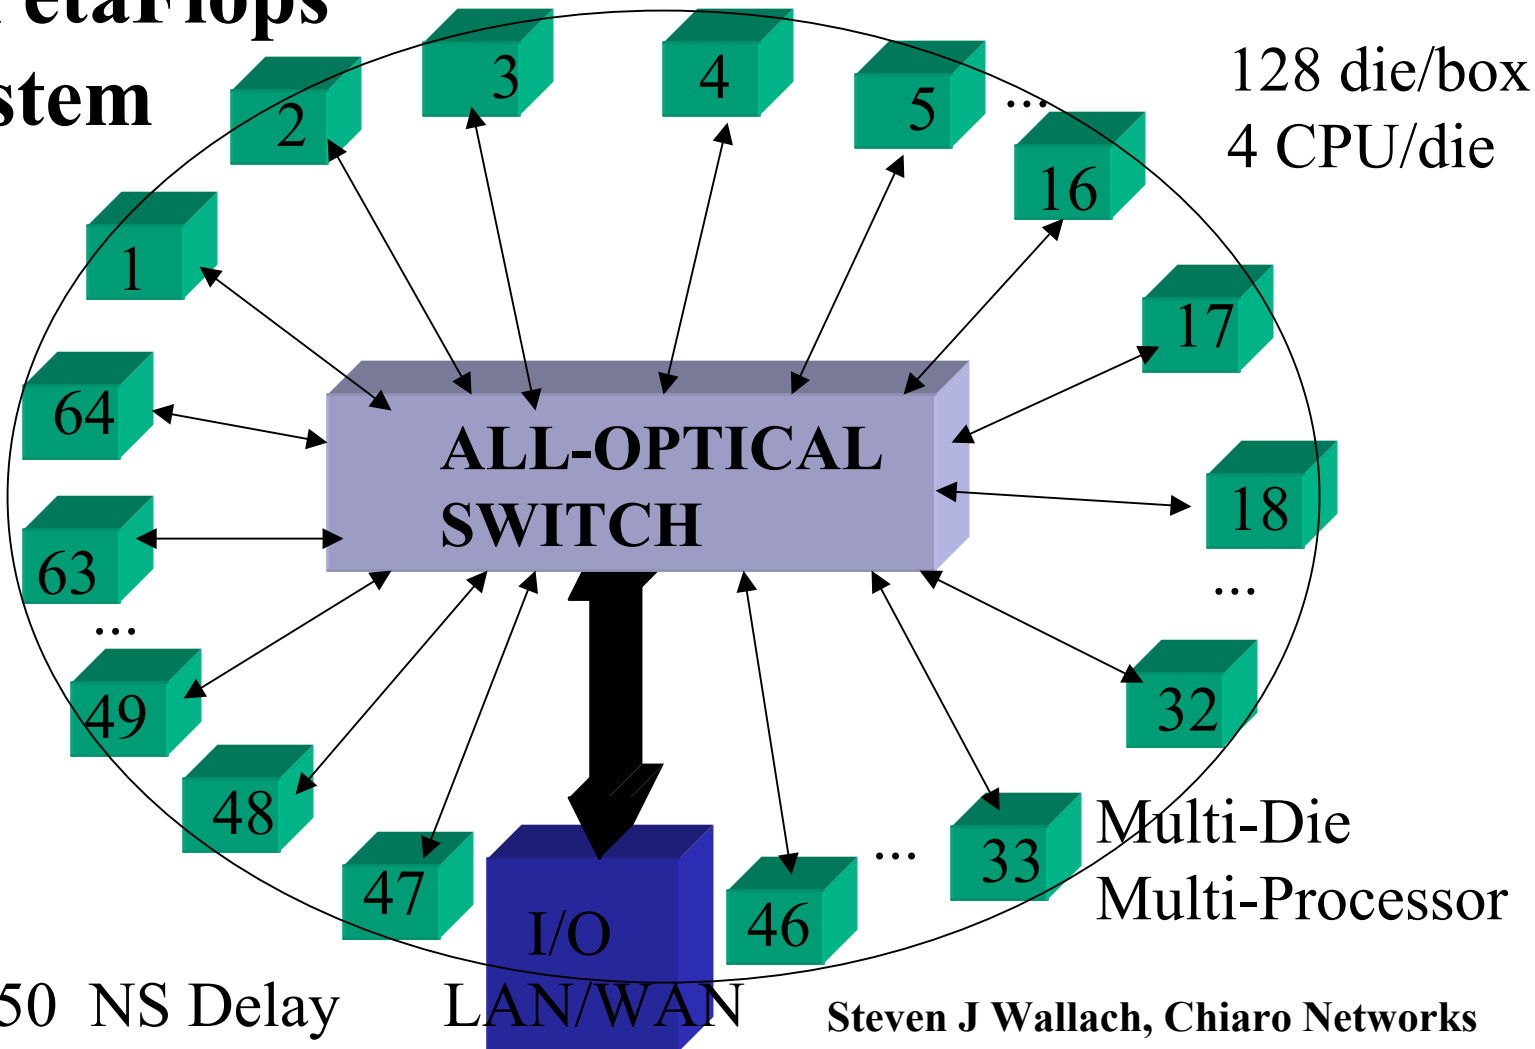
ASCI White





COTS PetaFlops

System



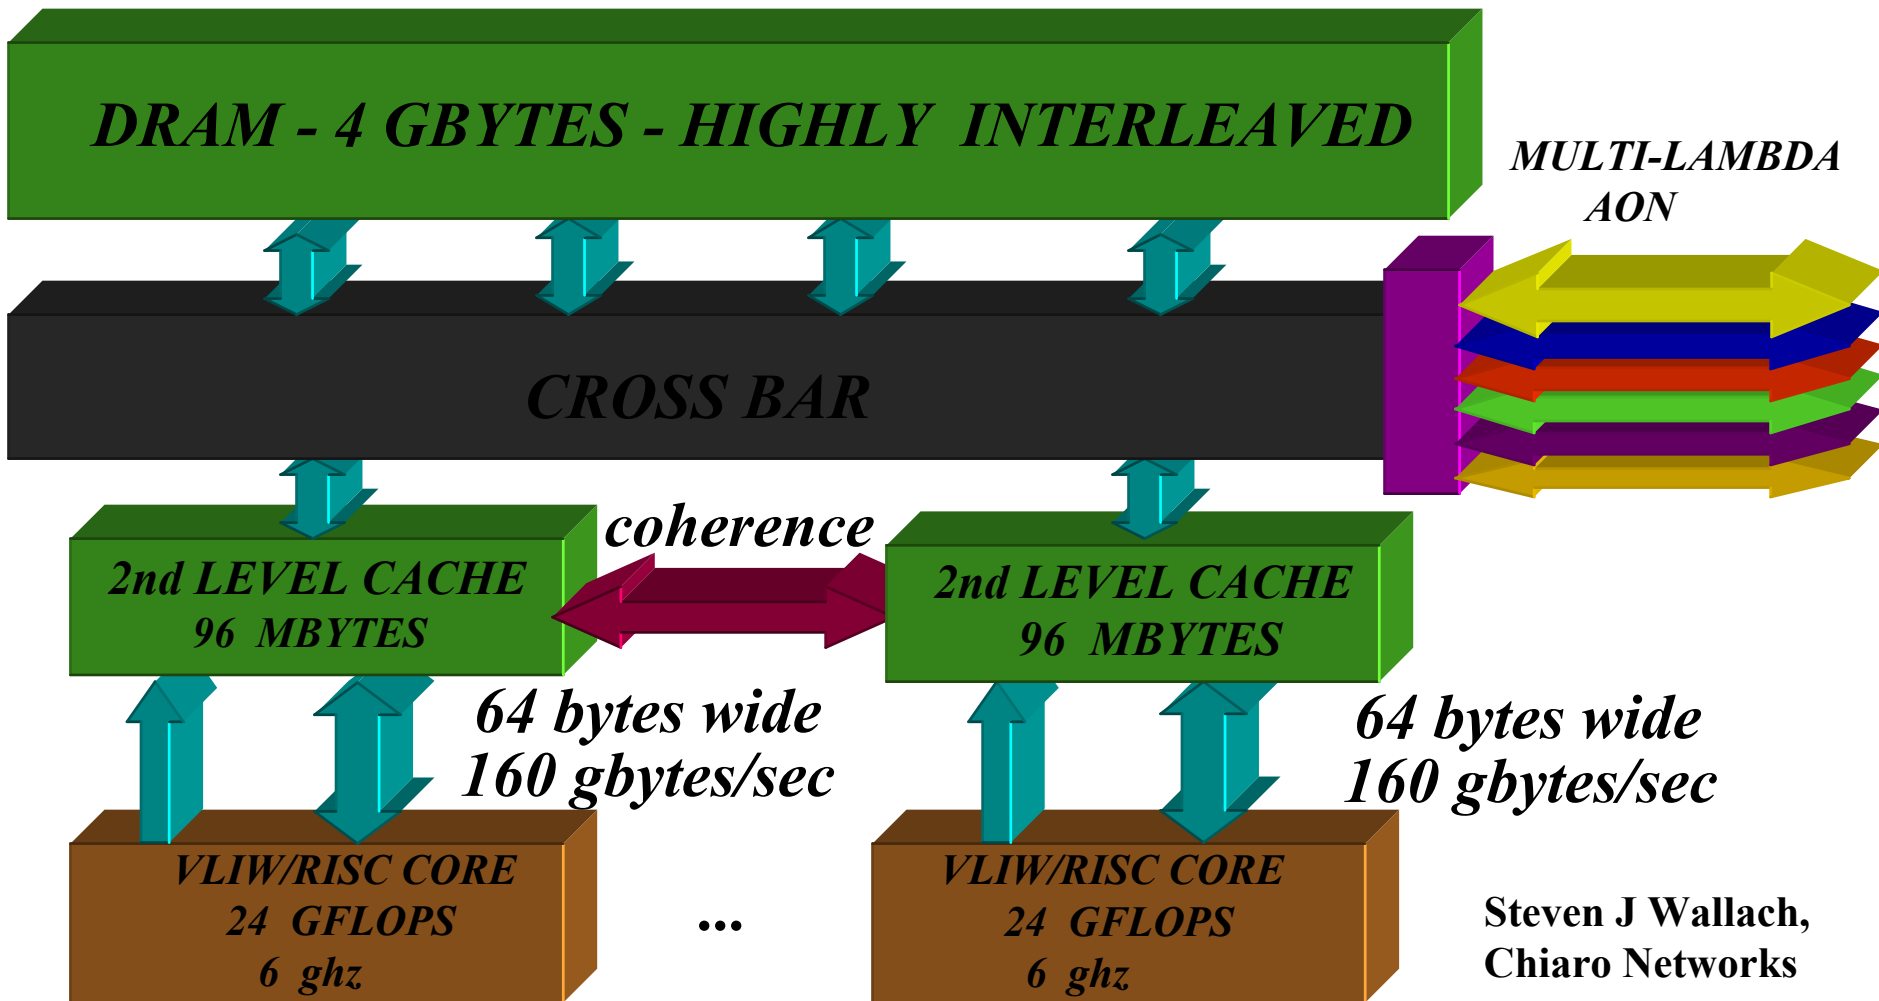
10 meters = 50 NS Delay

LAN/WAN

Steven J Wallach, Chiaro Networks



INTEGRATED SMP - WDM





ARCH. - LONG TERM- 2009

- **THE SIA STUDY TEACHES US:**
 - 64 gbits of dram - (8 gbytes)
 - 8 gbits of sram
 - 520 million MPU transistors
 - 70 nm lithography, 2.54 cm on-a-side
 - 6 ghz clock within vliw/risc core
 - 2.5 ghz across die
 - 2500 external signal pins

Steven J Wallach, Chiaro Networks



ARCH. - LONG TERM - 2009

DESIGN ASSUMPTIONS

- 9 million transistors - vliw/risc core with first level cache.
- 2nd. Level cache - rule of thumb. 1/4 to 1/2 mbyte per 100 mflops peak.
- 96 mbyte 2nd. Level (6 Inst, 90 data)
- 170 watts
- .6 to .9 volts power supply

Steven J Wallach, Chiaro Networks



COTS PetaFlop System

- 15.36 TFlops/Rack (128 die)
- 30 KWatts/Rack - thus 64 racks - 30 inch
- Common System I/O
- 2 Level Main Memory
 - local on chip
 - off-chip within same rack
 - same bandwidth/longer latency
 - meets Byte/Flop metrics
 - reduces external rack bandwidth by a factor of 10

Steven J Wallach, Chiaro Networks



COTS PetaFlop System

- Optical Interconnect
 - OC768 Channels (40 GHz)
 - 128 Channels per Die (DWDM)-5.12 THz
 - 2004/2005 Commercial Transmission
 - ALL Optical Switching
- Bisection Bandwidth of 50 TBytes/sec
 - 15 TFlops/rack*.1bytes/flop/sec*32 racks
- Rack Bandwidth - 15 TFlops*.1= 12 THz
- 2-4 OPTICAL MT-RJ Connectors or equivalent per rack

Steven J Wallach, Chiaro Networks



COTS PetaFlop System

- 8192 Dies (4 CPU/die-minimum)
- Each Die is 120 GFlops
- 1 PetaFlop Peak
- Power 8192 x200 Watts = 1.6 MegaWatts
- Extra Main Memory yields in excess of
3 MegaWatts (512 TBytes)

Steven J Wallach, Chiaro Networks



The HTMT



Thomas Sterling, JPL/Caltech



KUNGL
TEKNISKA
HÖGSKOLAN

ParallelDatorCentrum - PDC

Enabling Science




Strategic Computing Complex





Strategic Computing Complex





The Strategic Computing Complex

The Strategic Computing Complex will be a modern facility and the cornerstone of the revitalized TA-3 at Los Alamos National Laboratory.

303,000 gross sq. ft. , 43,500 sq. ft. unobstructed computer room

Scalable building utilities (30 to 150 TeraOPS)

Power - 7.1 MW expandable to 30 MW

Visualization theaters (2)

Collaboratories (5)

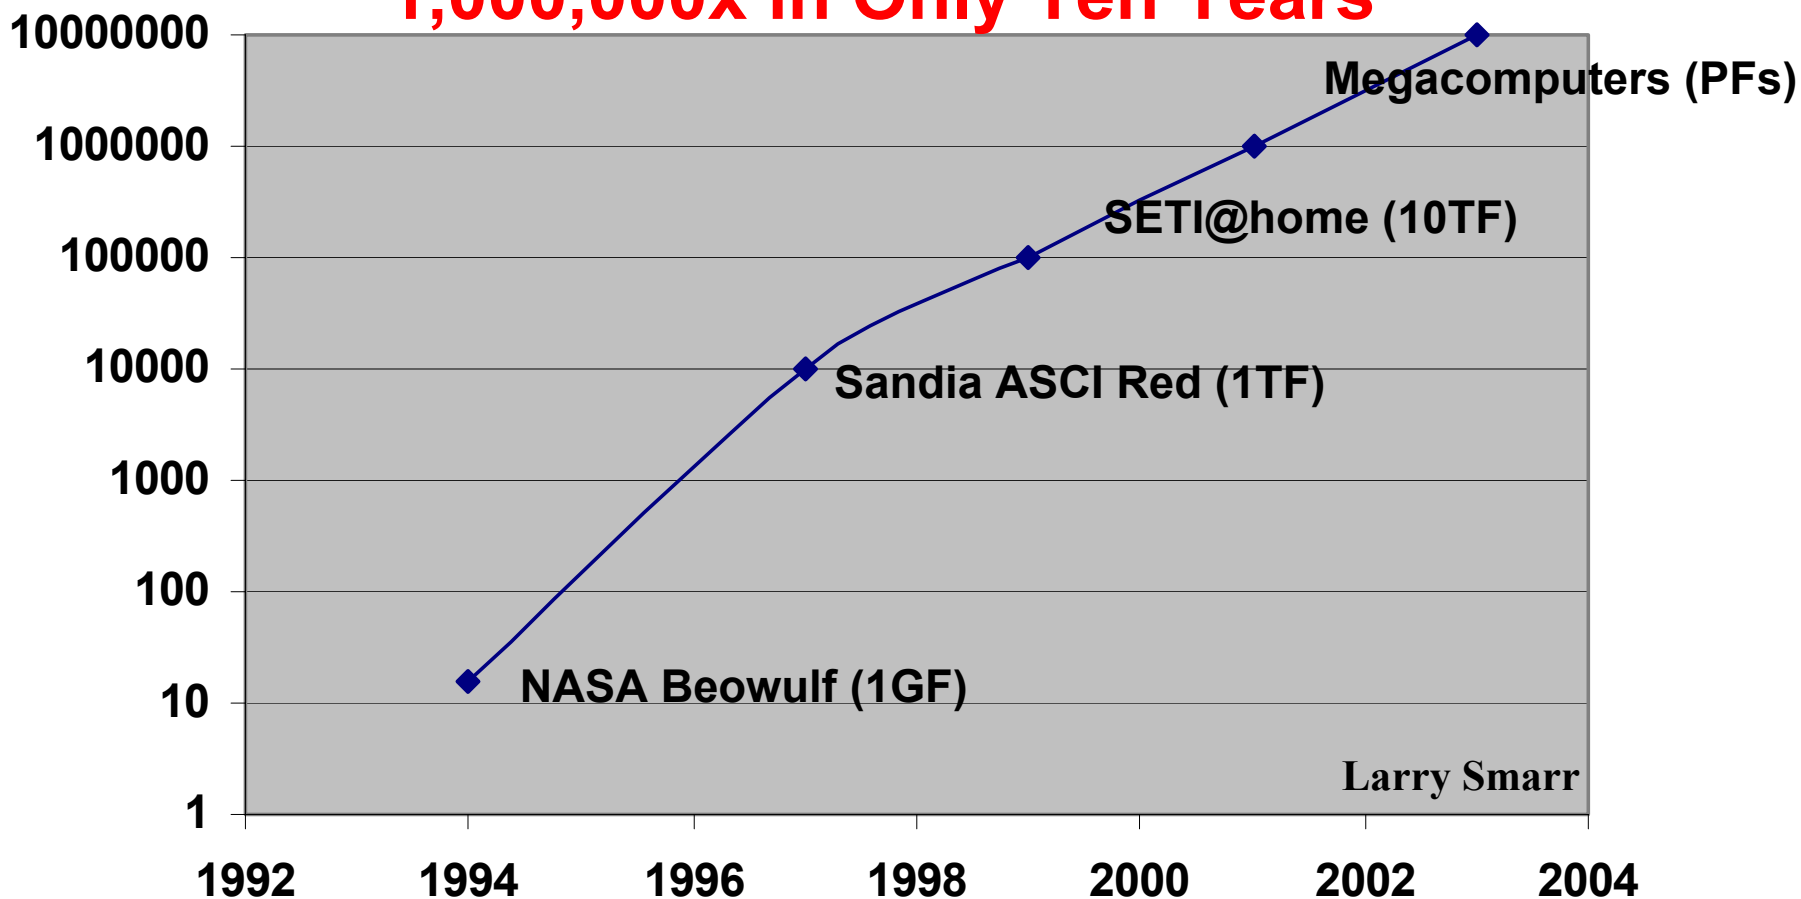
Design Simulation Laboratories (200 classified, 100 unclassified)

200 seat auditorium



Exponential Growth in Scale of PC Parallel Computers

1,000,000x in Only Ten Years





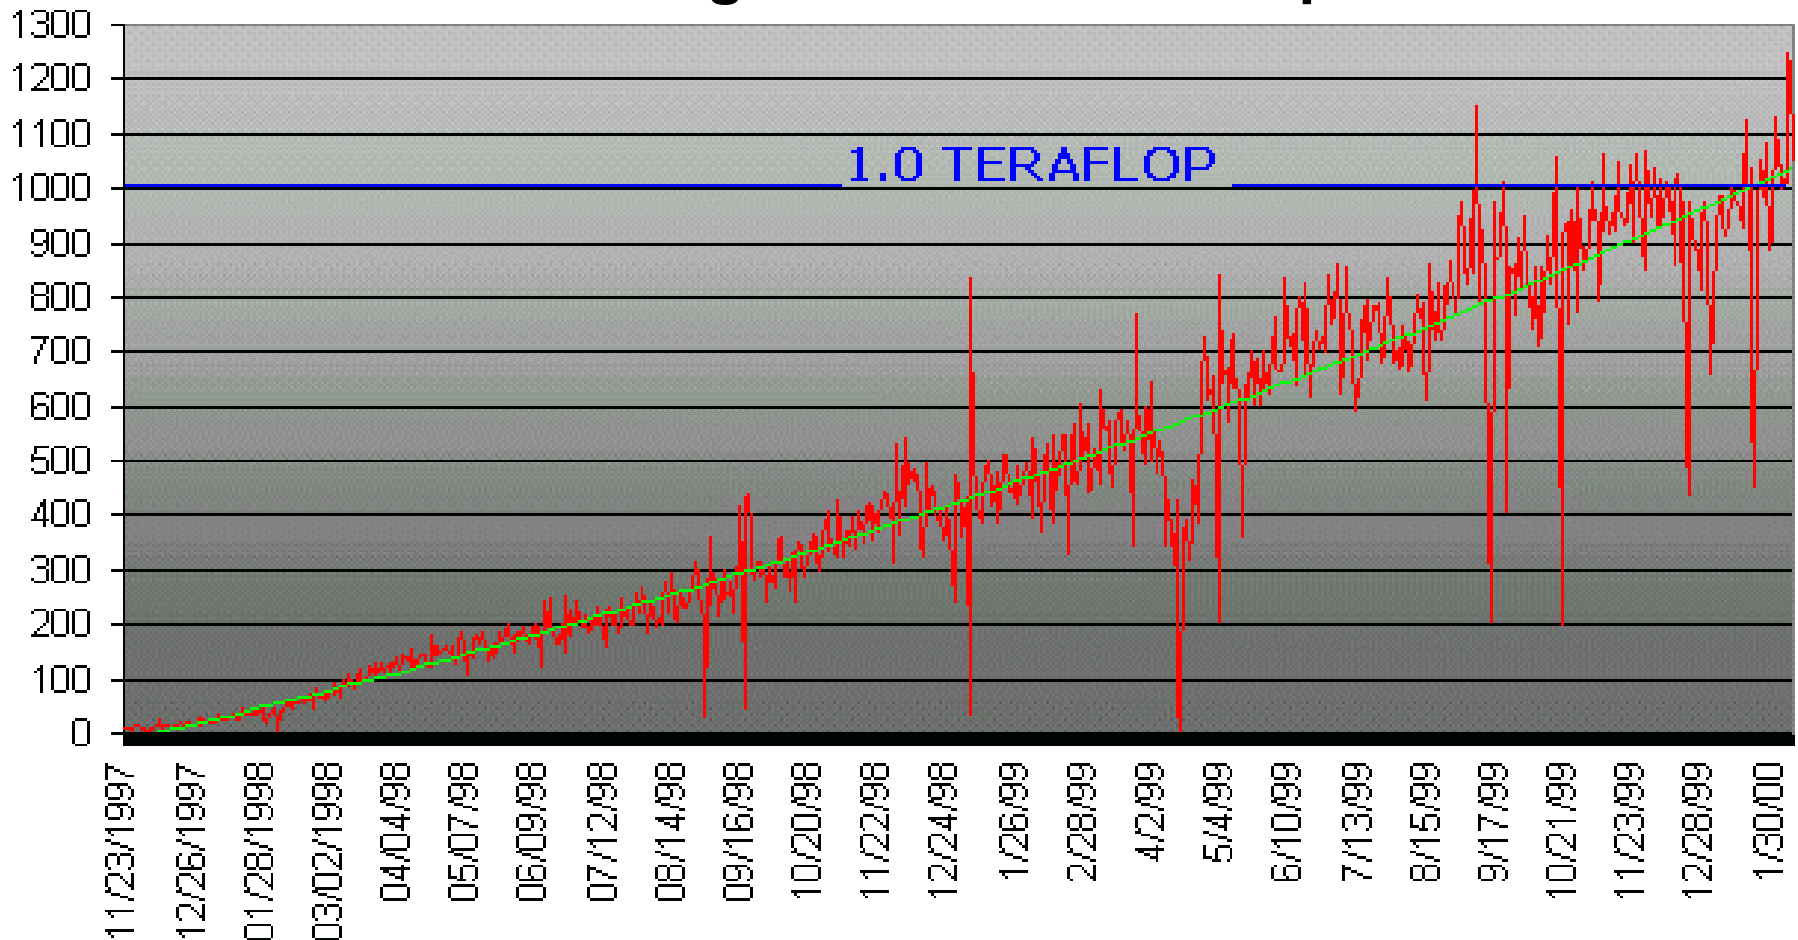
The Network the Becoming the Computer

- Internet Provides Connectivity
- Web Provides Hyperlinked File System
- Distributed Storage Moving from SAN to NAS
- Peer-to-Peer Computing Provides Vast CPU Power
- Result--The Distributed Global Computer
 - Storage everywhere
 - Scalable computing
 - Wireless Interfaces Greatly Outnumber PC Interfaces

Steven J Wallach, Chiaro Networks



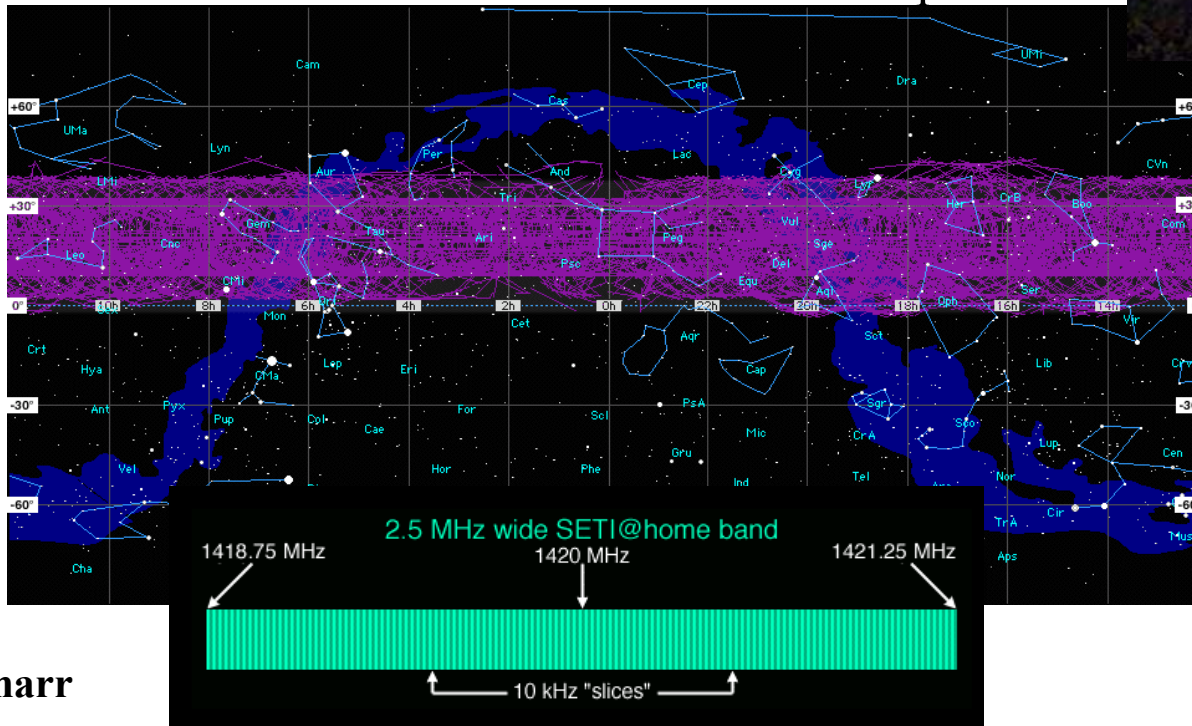
The Great Mersenne Prime (2^P-1) Search (GIMPS) Found the First Million Digit Prime www.entropy.com





SETI @ home was the Breakthrough to PC Internet Computing

- Running on 500,000 PCs, ~1000 CPU Years per Day
 - 428,438 CPU Years so far
- Sophisticated Data & Signal Processing Analysis
- Distributes Datasets from Arecibo Radio Telescope




Next Step-
Allen Telescope Array





Running AutoDock Application Software on 2100 PCs




Art Olson,
TSRI

fight **AIDS** @home the Olson laboratory at
The Scripps Research Institute
computing toward a cure



Larry Smarr



*Theory and Simulation of Biological and Biomimetic
Self-Assembly and Function*

10,000 PCs
70 CPU-Years
as of 10/13/00

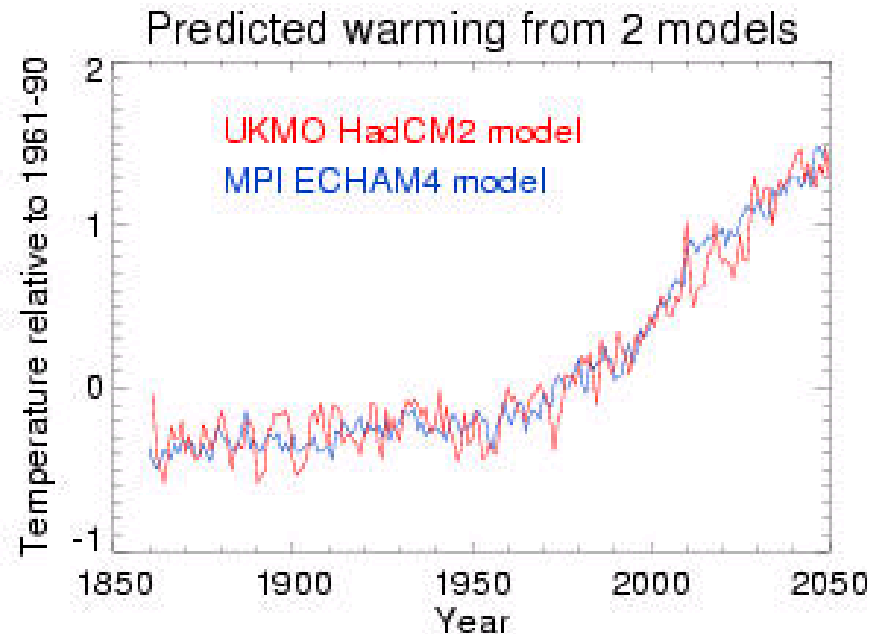


Larry Smarr



Climate Simulations

- Ensemble Computing Varying Model Parameters
- Evaluate Model Against Current Climate
- Home in on Most Realistic Models by Natural Selection
- Then Model 21st Century Climate Evolution
- One Climate Model per PC
- Currently About 20,000 PCs



Larry Smarr, UCSD/CALIT2



Applications

Distributed Terascale Facility

Application	Attributes	Memory Required	Time (5 TF)	Storage Required	Aggregate I/O BW	Network Bandwidth
Cactus	Neutron star collision on a 1024^3 grid	1.4 TB	60 hours	20 TB	1.9 GB/s	100 MB/s
ARPS	U.S. weather model	250 GB	5 hours	16 TB	6.7 GB/s	
TPM	Strong lensing with 2048^3 particles	1.0 TB	200 hours	5 TB	3.8 GB/s	N/A
MILC	$42^3 \times 84$ lattice simulation	0.2 TB	2000 hours	1 TB	1.1 GB/s	N/A
ASPCG	LES of 3-D fins	0.5 TB	100 hours	50 TB	1.1 GB/s	N/A
PPM	Turbulent fluid flow On a 2048^3 grid	0.5 TB	100 hours	54 TB	2.3 GB/s	NA
NAMD	10 ns, 1 million atom bio-MD simulation	0.9 TB	250 hours	250 GB	1.0 GB/s	N/A
LIGO	Gravity wave detection			100 TB/year		9 MB/s
LHC	Particle detection	1.2 TB	150 hours	5-7 PB/year	5.6 GB/s	200 MB/s
NEES	Earthquake engineering	0.5 TB	10 hours	50 TB/year	1.0 GB/s	200 MB/s
Genomics	Structure determination			20 TB/year		N/A



Internet Computing

- **Individual Processors Running at Gigaflops**
 - **One Million Means a Collective Petaflops**
 - **One Petaflops is Roughly a Human Brain-Second**

Larry Smarr

ParallelDatorCentrum - PDC

Enabling Science

