Towards the Teraflops Capability for CFD

Horst D. Simon

Computer Sciences Corporation

NAS Applied Research Branch

Mail Stop T045-1

NASA Ames Research Center

Moffett Field, CA 94035

simon@nas.nasa.gov

NASA Computational Aerosciences Conference

August 20, 1992

OUTLINE

I. A Look at the New Machines

II. Architectural Trends

III. Consequences for CAS Researchers

Comparison of the Machines

Type	Company	Machine	Year	Peak Gflops
MIMD global	Cray Research	C90	1991	16
	Cray Computer	Cray-3	1992	16
	Fujitsu	VP2600	1990	4
	Hitachi	S820-80	1990	3
	NEC	SX-3	1990	22
	SSI		1993	?
MIMD distributed	Kendall Square	KSR1	1991	40
	Tera	T1	1993	300
	Cray Research	MPP	1993	150
MIMD local	Intel	Paragon	1992	300
	NCUBE	NCUBE3	1994	6500
SIMD/MIMD	Thinking Machines	CM-5	1991	1000
SIMD	Maspar	MP2	1992	?

MIMD global - Cray Y-MP16/C90

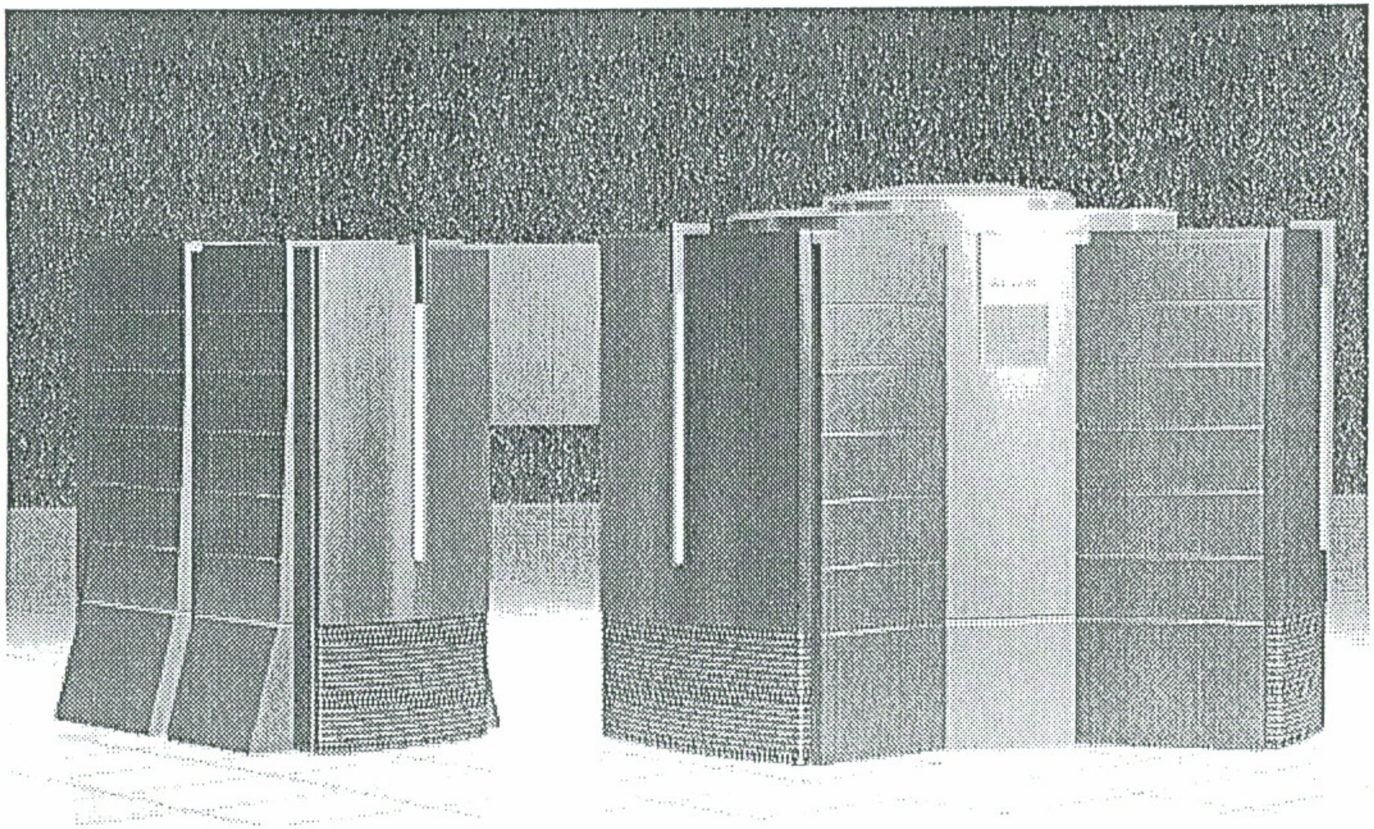

MIMD global - Cray Y-MP16/C90

Table 1. Peak Performance Rates

Computer	Cycle Time (ns)	Single CPU Peak MFLOP Rating	Multiple CPU Peak MFLOP Rating
Cray-1	12.5	160	
Cray X-MP	8.5	233	932, 4 CPUs
Cray-2	4.1	488	1,952, 4 CPUs
Cray Y-MP	6.0	333	2,666, 8 CPUs
Cray C 90	4.0	1,000	16,000 16 CPUs
Cray-3	2	1,000	16,000, 16 CPUs
CYBER 205	20	200	
ETA-10/G (1988)	7	625	5,000, 8 CPUs
Fujitsu VP-400E	7	1,700	
Fujitsu VP-2600	4	4,000	
Hitachi S-810/20	14	630	
Hitachi S-820/80	4	2,000/3,000	
NEC SX-2	6	1,300	
NEC SX-3	2.9	5,500	22,000, 4 CPUs

MIMD distributed - KSR 1

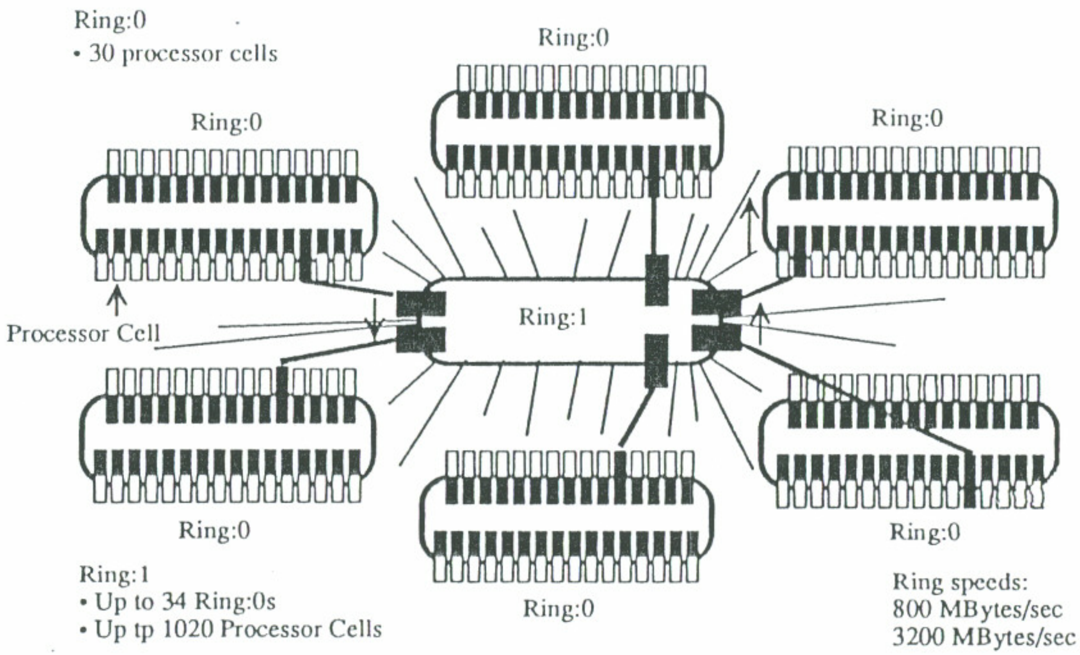

MIMD distributed - KSR 1

- Custom designed 40 Mflops RISC processor
- Up to 1088 processors in hierarchy of rings
- Aggregate performance over 40 Gflops
- 32 Mbyte memory per processor
- Allcache virtual memory system (2^{40} address space)
- Automatic parallelizing Fortran

MIMD distributed - KSR 1

Technical Strengths

- ease of use of the system

Technical Challenges

- scalability unproven
- current processor technology too slow (for custom processor)

MIMD distributed - Cray MPP

- Three year development plan with new generations of machines projected to be available in 1993 (¿150 GigaFLOPS Peak), 1995 (1 TeraFLOPS peak), and 1997 (1 TeraFLOPS sustained).
- The MPP component will be closely coupled with a vector/scalar GP
- The first generation MPP PE will be based on the DEC Alpha microprocessor providing a peak of 150 MFLOPS per PE.
- Future generations may use a different micro-architecture.
- First MPP will scale up to 100s of GFLOPS peak
- Physically distributed memory with 8 - 32 Mbytes of memory per processor

MIMD distributed - Cray MPP

Technical Strengths

- experience in scientific computing (hardware, software, applications, customer support)
- existing user base to build on

Technical Challenges

- integration of outside technology (DEC Alpha)
- software model

MIMD distributed - Tera

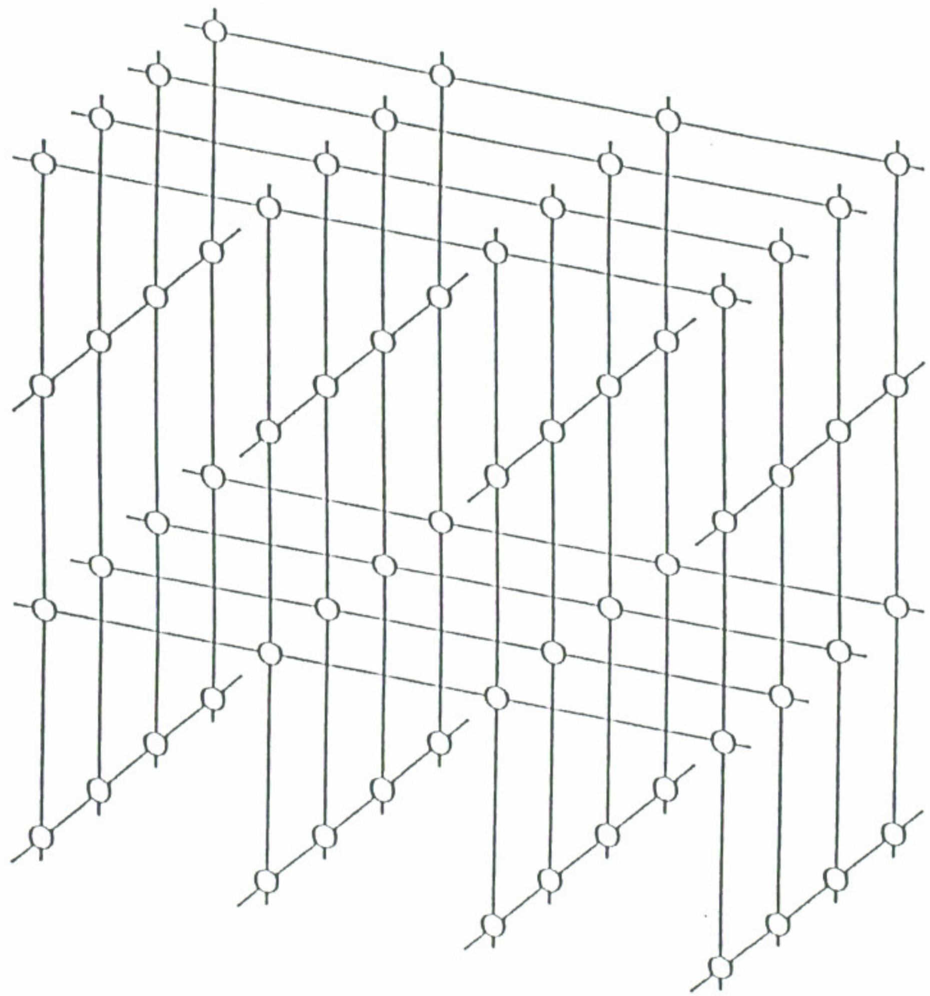

TERA

Interconnection Network

MIMD distributed - Tera

- Shared memory parallel processor - data can be anywhere
- Scalar instructions
- Fast floating point - 300 Gflops peak performance
- Efficient synchronization - up to once per cycle per proc
- General purpose - numeric/nonnumeric
- Automatic parallelizing compilers - both Fortran and C++
- Unix, TCP/IP, NFS

MIMD Multicomputer - Intel Paragon

MIMD distributed - Tera

Technical Strengths

- very powerful processor
- potential ease of use of the machine

Technical Challenges

- small company with large number of technological risks
- will latency hiding really work

MIMD Multicomputer - Intel Paragon

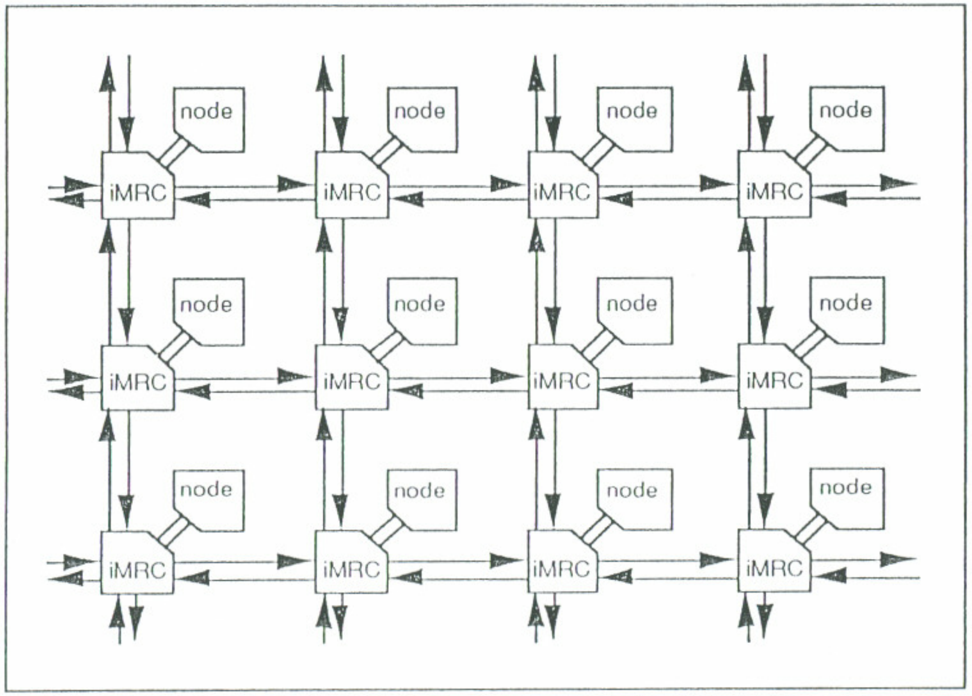

MIMD Multicomputer - Intel Paragon

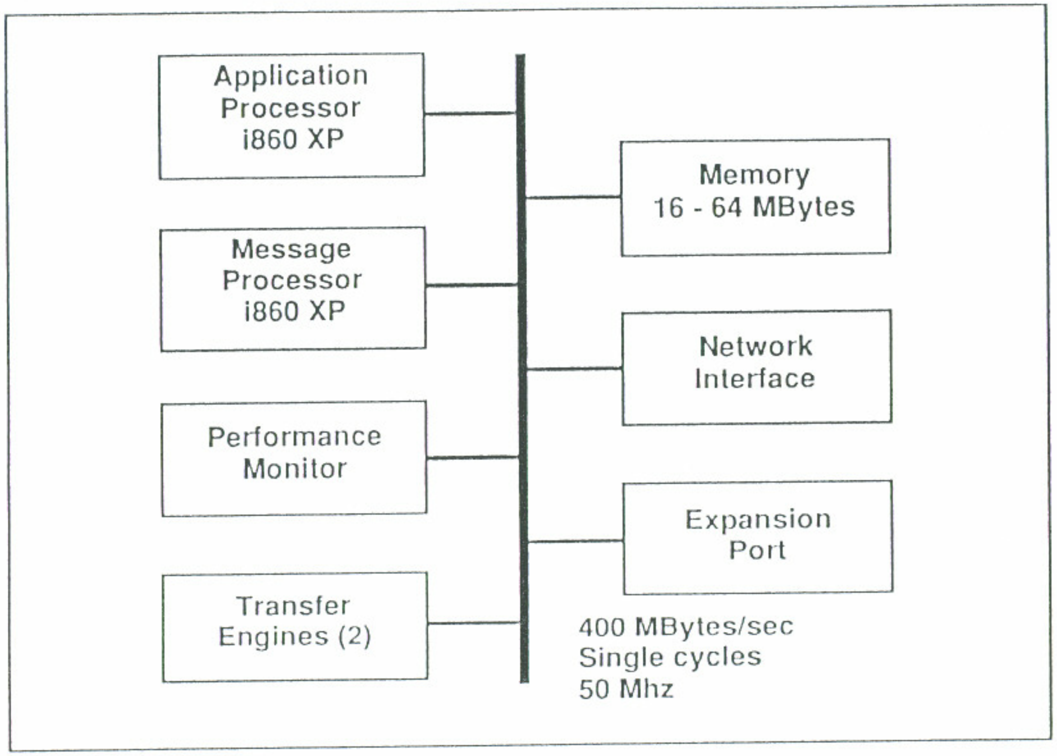

MIMD Multicomputer - Intel Paragon

- Uses i860XP microprocessor at 75 Mflops
- Up to 4000 processors 2D mesh
- Aggregate performance over 300 Gflops
- 16 - 64 Mbyte memory per processor
- New message routing chip at 200 Mbytes/sec

MIMD Multicomputer - Intel Paragon

Technical Strengths

- long experience and user base in parallel processing
- “free” hardware development in house

Technical Challenges

- software still immature
- weak I/O system
- transition to real supercomputer company

MIMD Multicomputer - NCUBE3

- complete system redesign with new 50 Mhz custom processor
- up to 65,536 processors for a total 6.5 Tflops peak
- 1 Gbyte memory per processor for up to 65 Tbytes total
- hypercube interconnect
- scalable I/O and network capability

SIMD/MIMD - CM 5

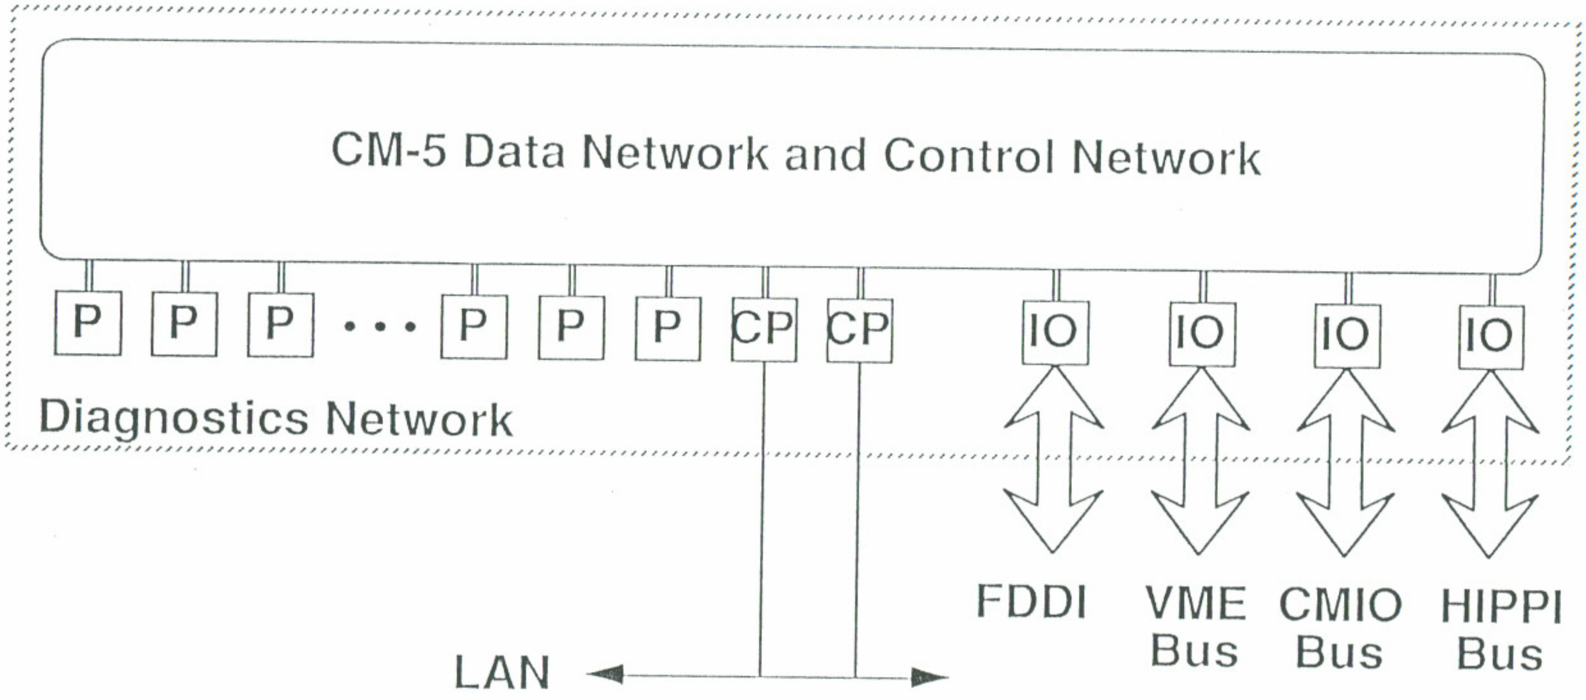

CM-5 Processing Node with Vector Units

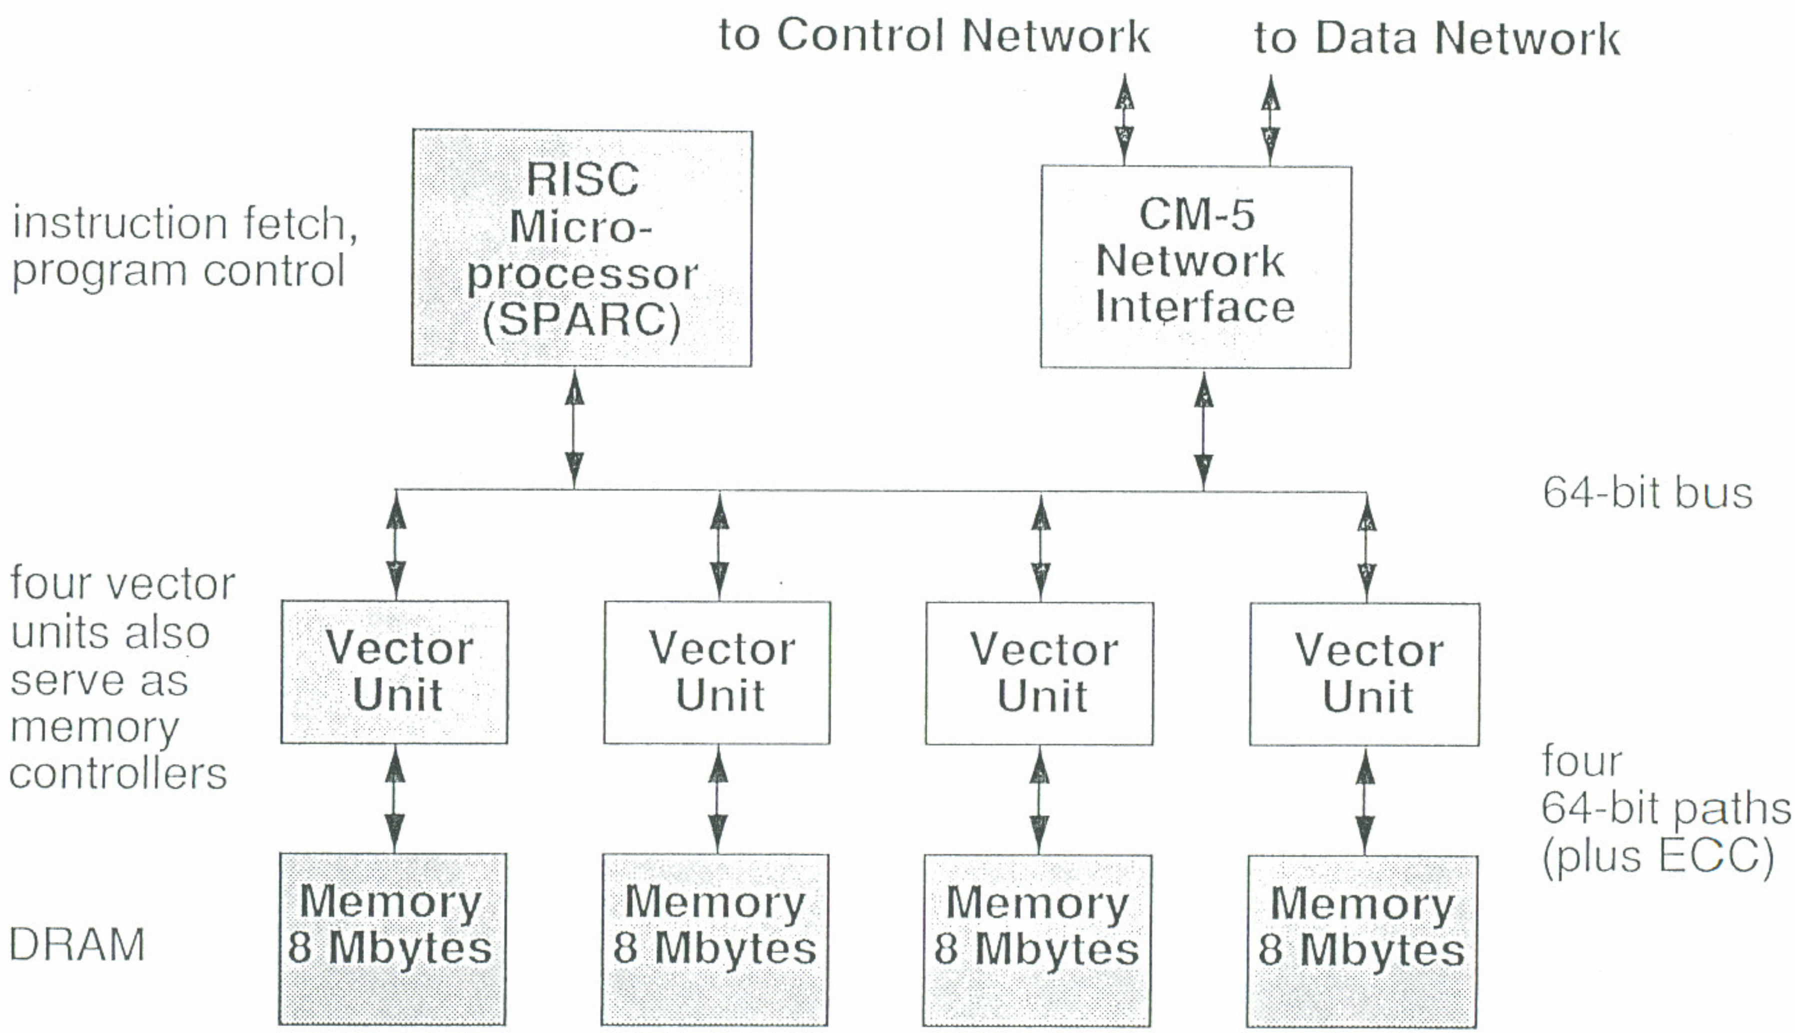

Data Network with 16 Nodes

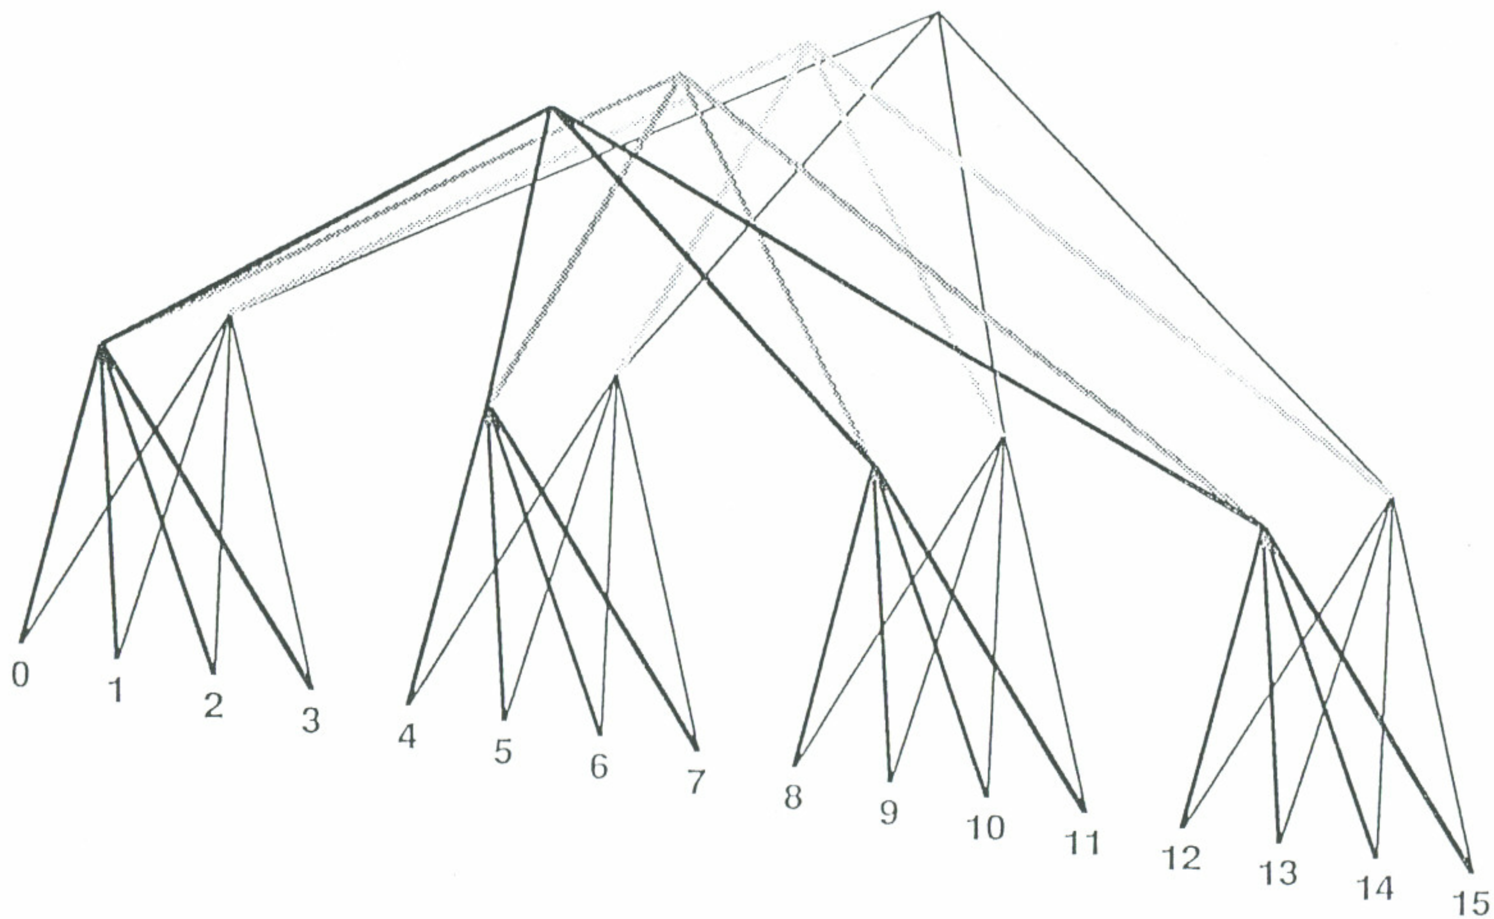

SIMD/MIMD - CM 5

Technical Strengths

- long experience and user base in parallel processing
- good system software and support

Technical Challenges

- transition to new node architecture
- transition to MIMD computational model

Trends and Issues in the Quest for the Teraflop

- SIMD approach has been dropped
- MIMD distributed memory (KSR, Cray, Tera) has re-emerged as a valid alternative
- custom processors versus commodity micro processors is still an open issue
- the field is still full of vaporware and premature announcements
- a Tflops peak machine at a reasonable price ($25 M) by 1995 is feasible
- sustained performance of CFD codes on such a Tflops machine is probably in the 10's Gflops range
- networked workstations will provide the most useful and cost-effective computing environment

- system software, algorithms, and applications will remain a major challenge for years to come

Fundamental Forces

Floating Point Performance

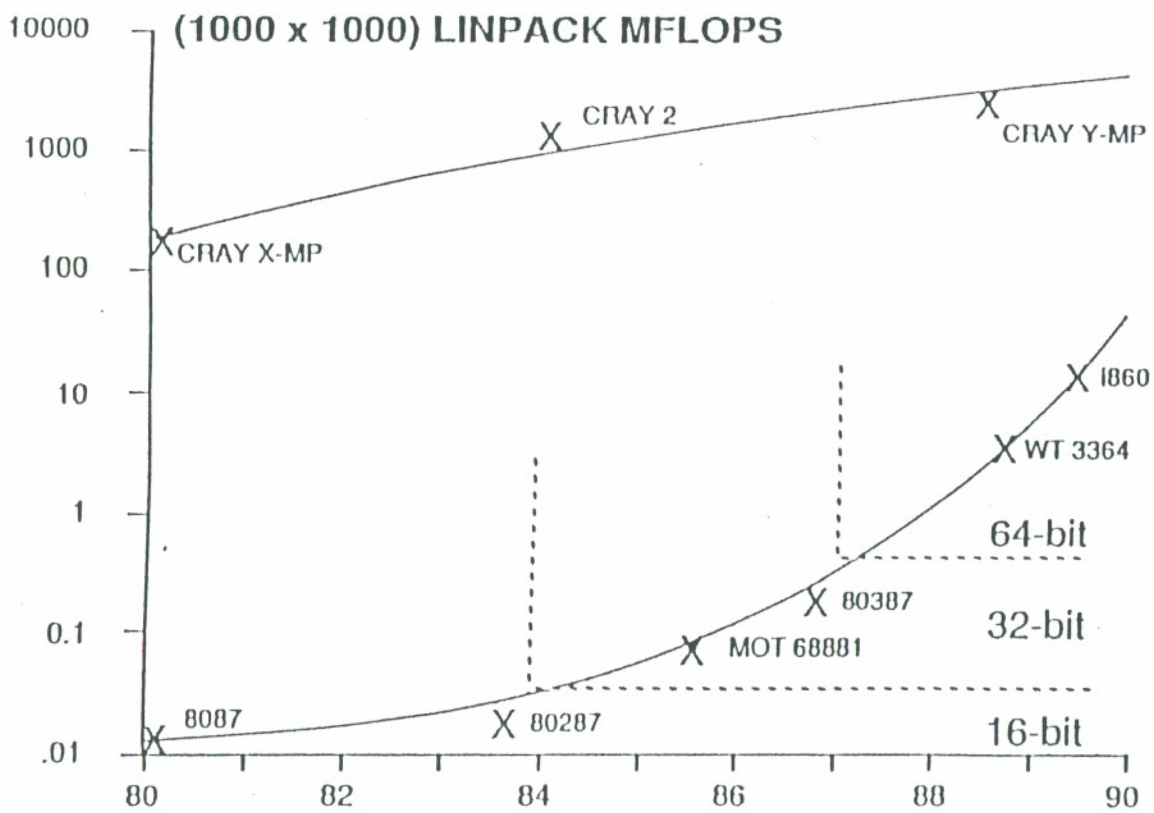

Floating Point Performance (1000 by 1000 LINPACK)

(single processor performance for supers)

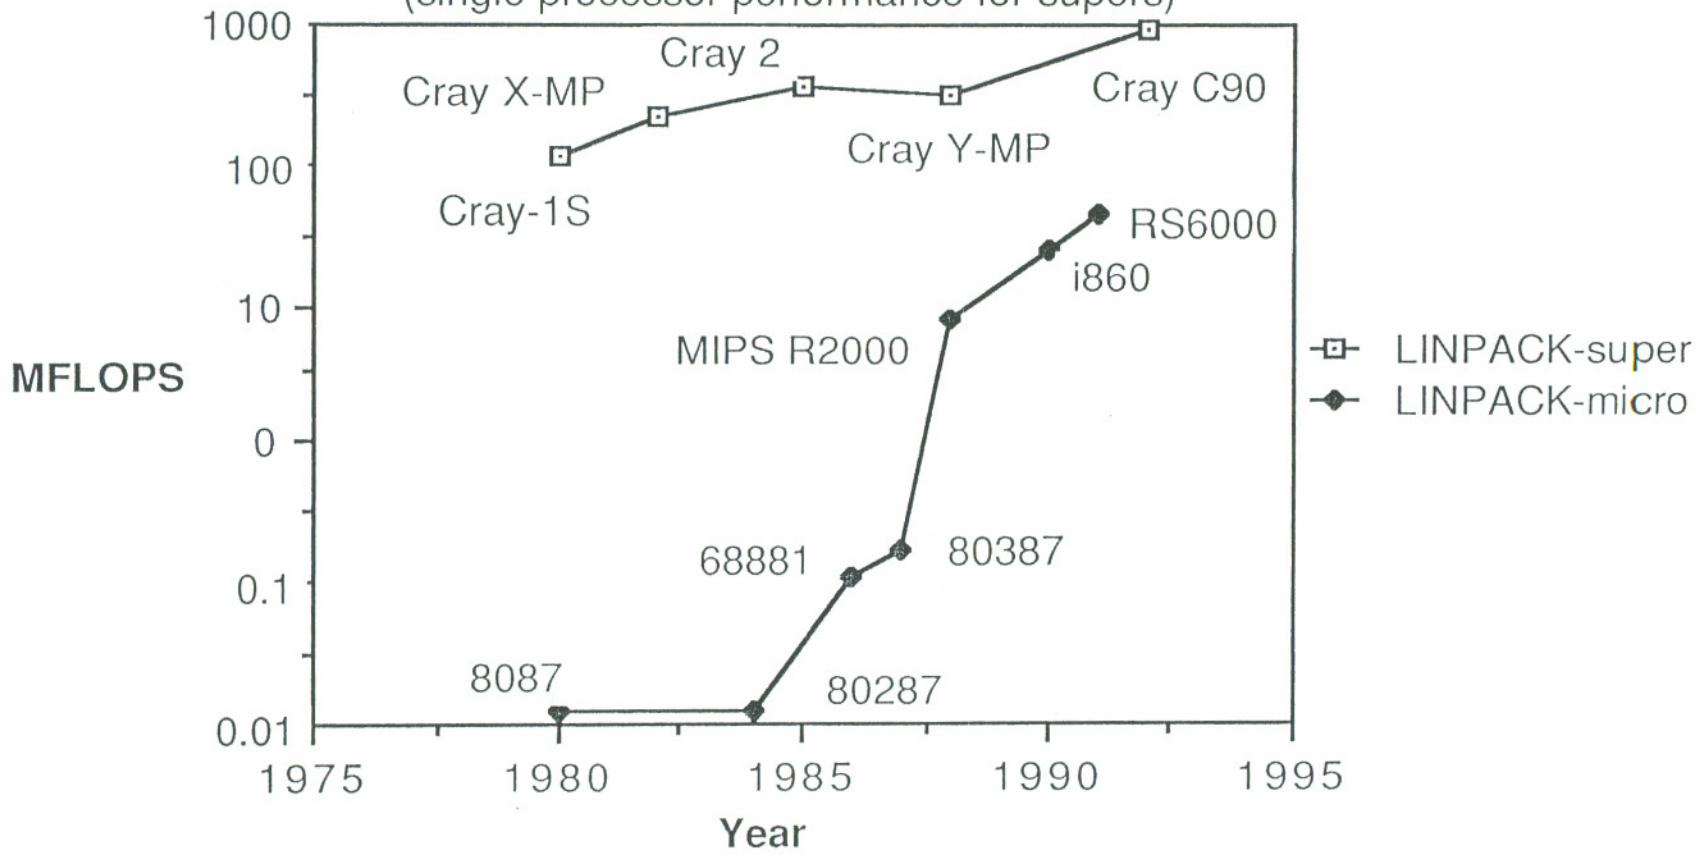

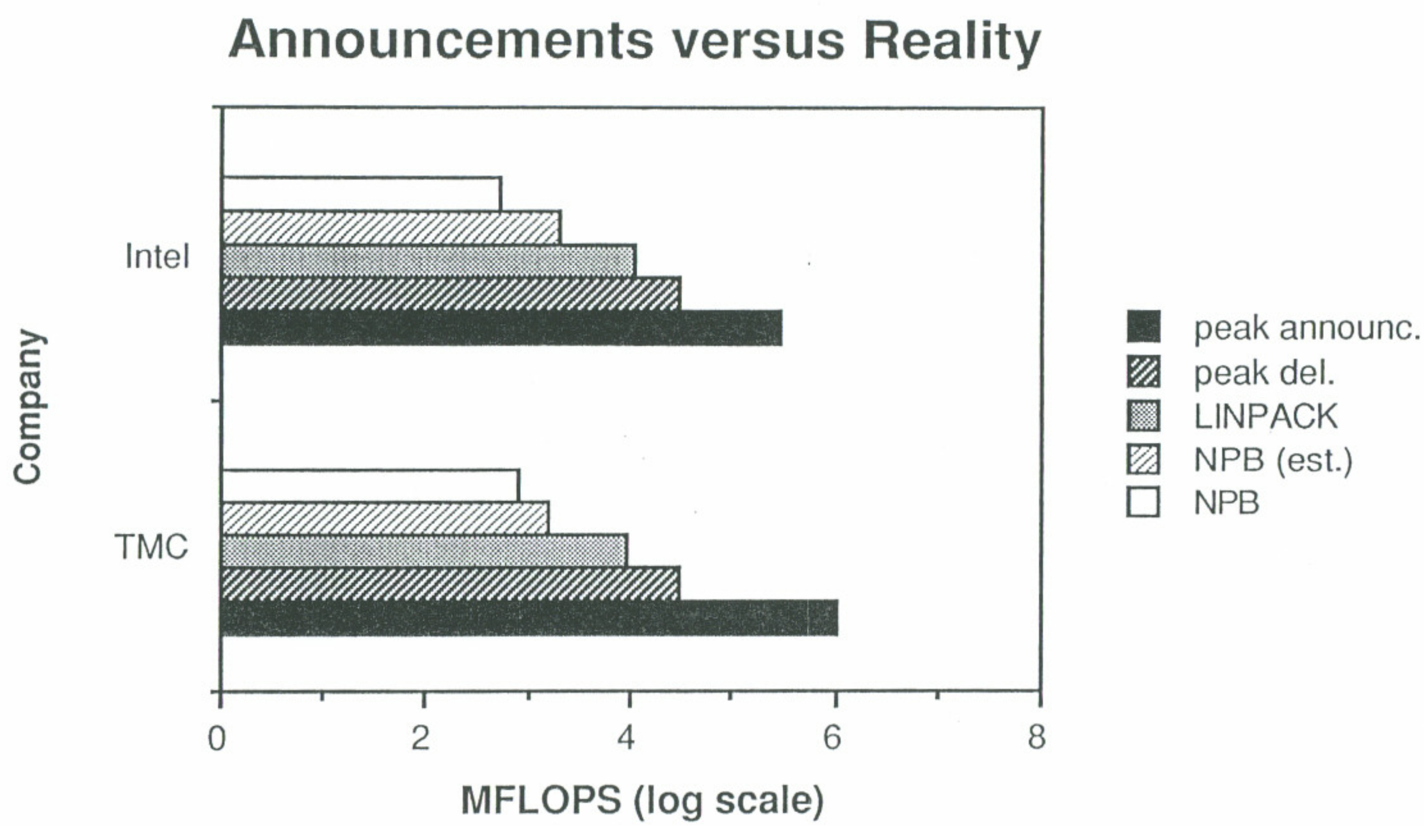
Announcements versus Reality
Company
Intel
TMC
0
2
4
6
8
MFLOPS (log scale)
peak announc.
peak del.
LINPACK
NPB (est.)
NPB

NPB Results in Megaflops

Computer System	No. Proc.	Kernels EP	MG	CG	FT	IS	CFD Applications LU	SP	BT
C-90	1	542	451	331	544	133	401	535	508
	16	8364	4067	2600	4469	+919	3665	7810	6386
Y-MP	1	211	176	127	196	68	194	216	229
	8	1681	1319	634	1344	422	1304	1579	1590
iPSC/860	128	435	403	70	521	24	146	206	430
CM-2	32K	792	150	105	309	14	113	94	95

Notes:
Megaflops rates are computed from standard NPB flop counts.
+ Only 8 CPUs were used.

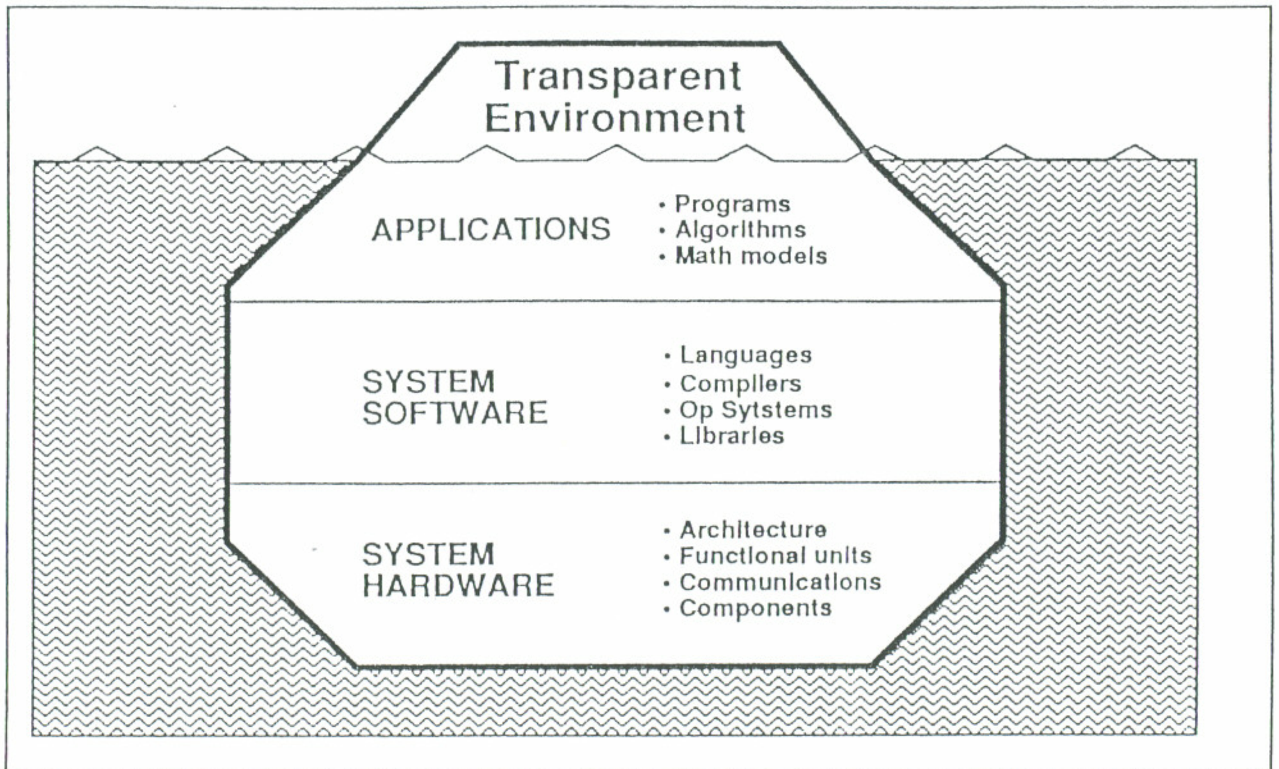

Figure 12. The Iceberg Metaphor for Mature High-Performance Computing.

source: Jack Worlton

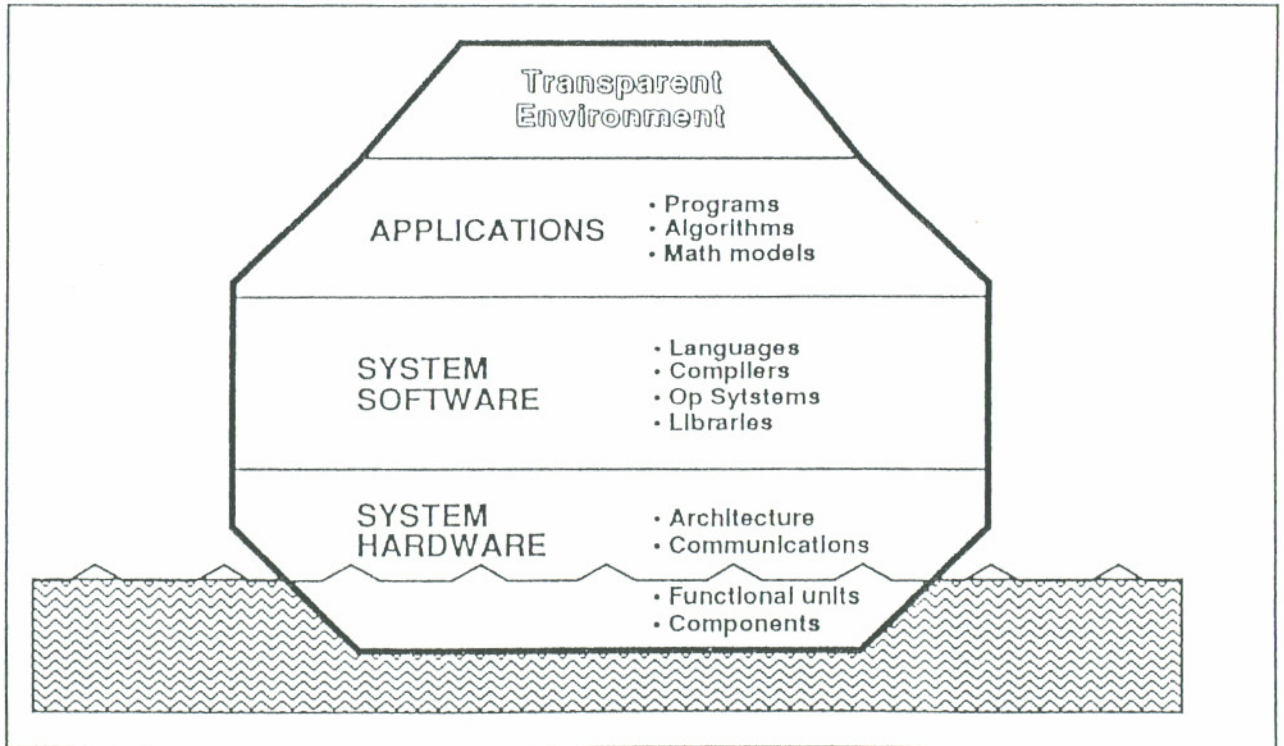

Figure 12a. The Iceberg Metaphor for Massively Parallel Computing.

Source: Jack Worlton