Silicon Graphics, Inc.

Linux64 im CAE Umfeld

3. LS-DYNA Forum

14.-15. Oktober 2004 in Bamberg

Josef Hellauer
josefh@sgi.com

Agenda

- **HPC im CAE-Bereich**
- **Systemanforderung CAE**
- **Linux64-Server SGI Altix**

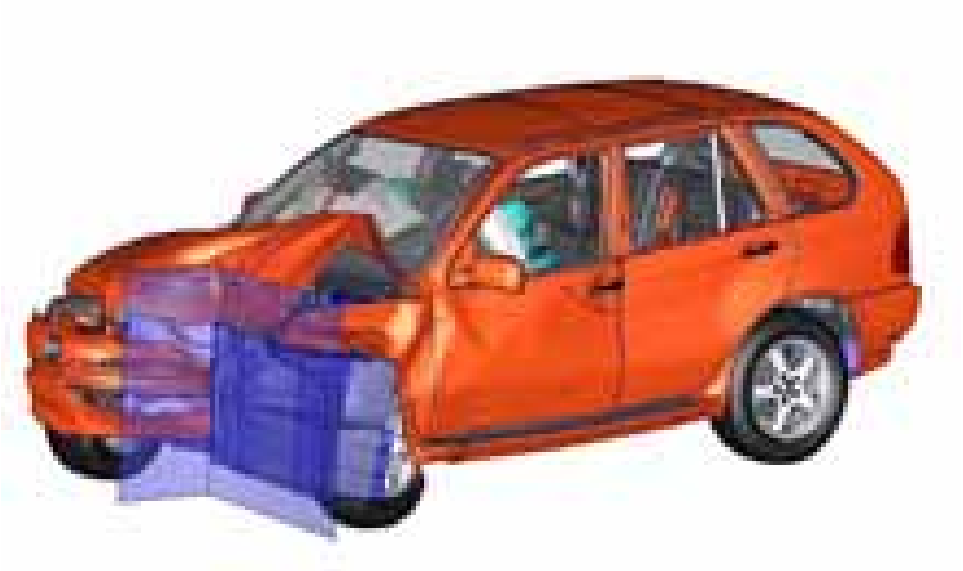

CAE Application Segments in Automotive 2003

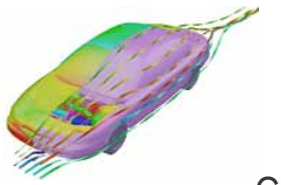

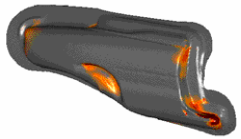

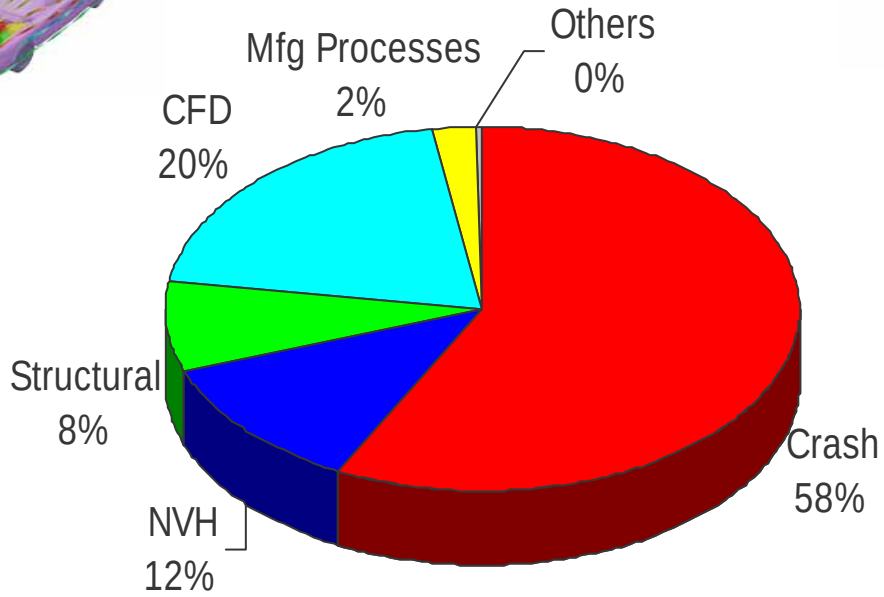

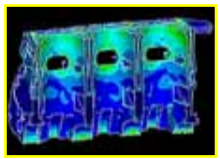

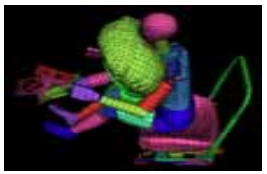

- Crash is the application segment #1
- The fastest growing application is CFD, won 3% shares from 2002
- NVH is the third application segment. However, it is the most demanding in terms of memory and IO bandwidth.
- 10 CAE Applications drive 89% of the installed compute power

CFD - Computational Fluid Dynamics
NHV - Noise, Vibration, Harshness
MFG Processes - Stamping, Casting, Forging

Source: Top20Auto, Ch.Tanasescu, SC'2003, www.top500.org
http://www.top500.org/lists/2003/11/Top20Auto_Top500V2.pdf

sgi

Key CAE Applications in Automotive

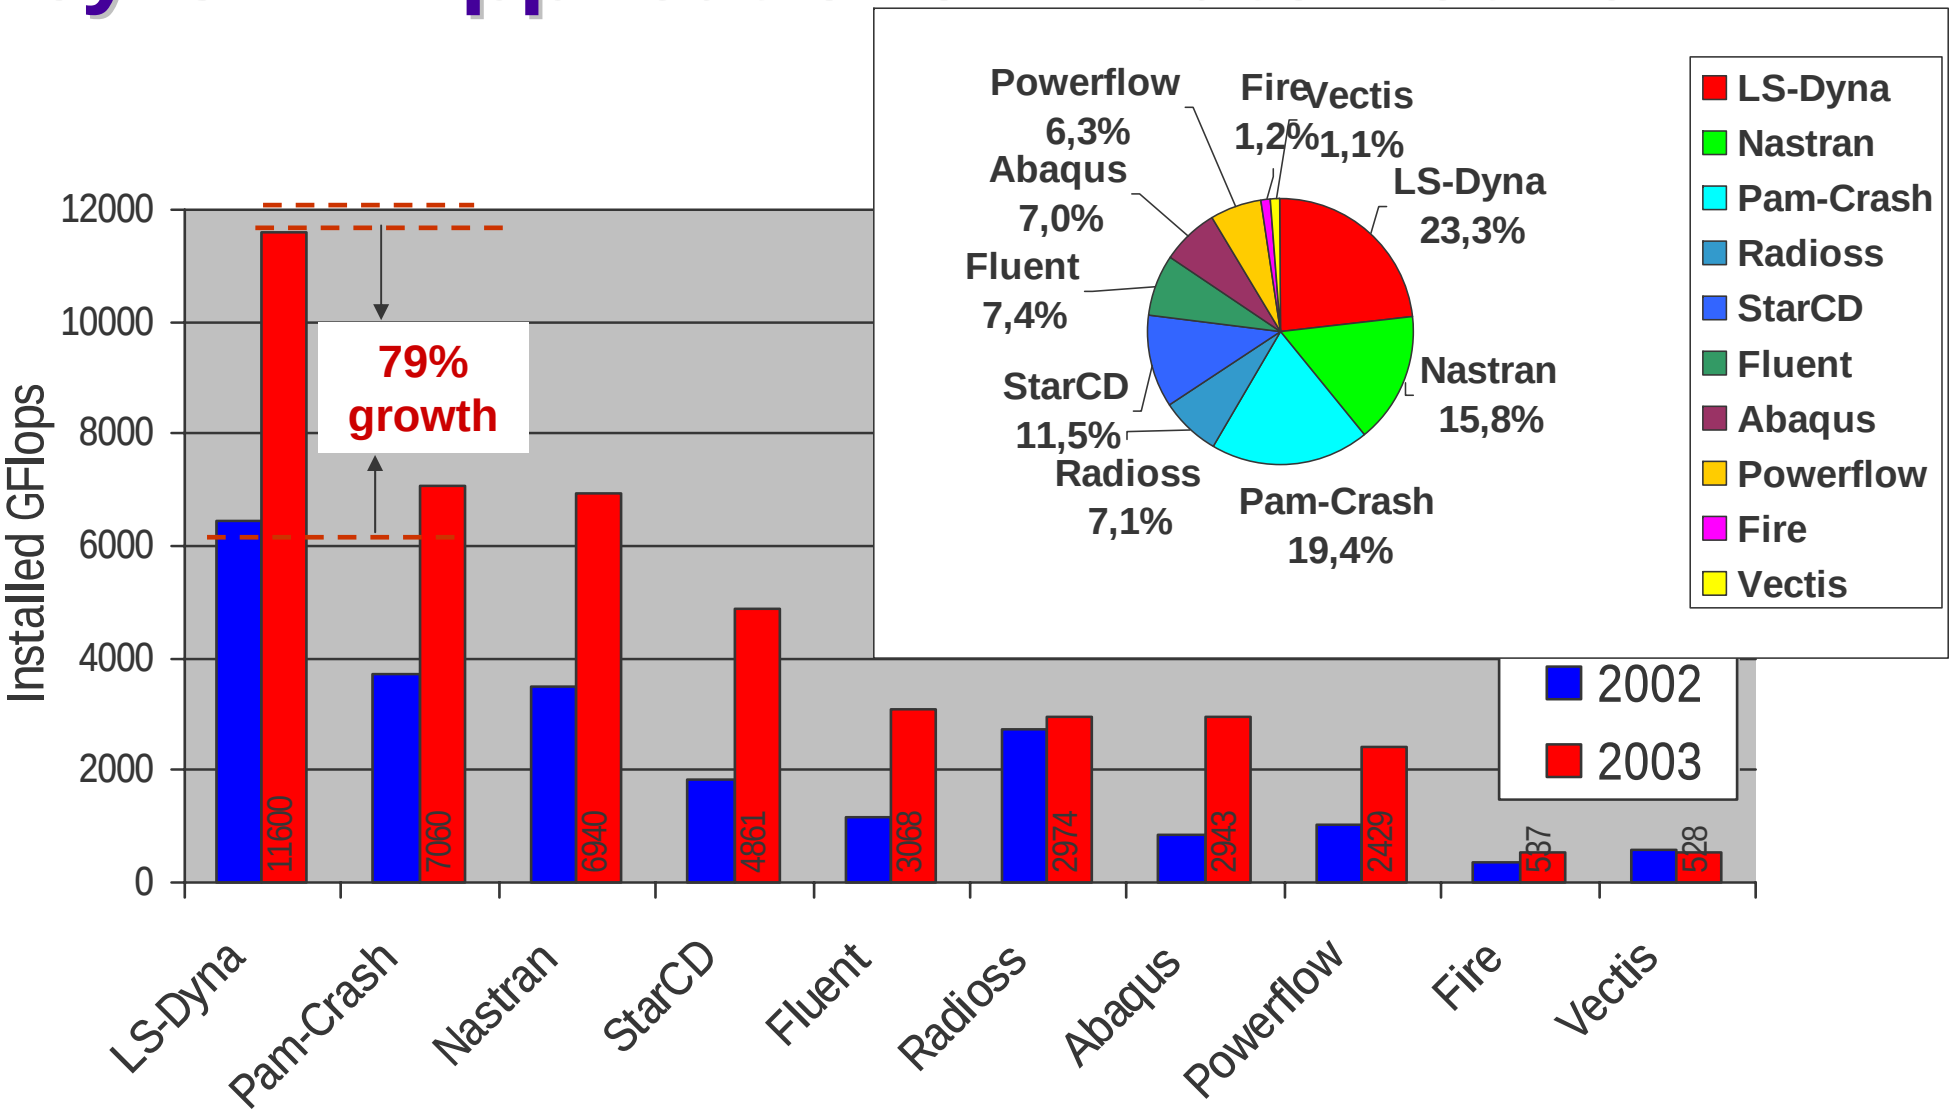

Other applications: Marc, Ansys, Madymo, AMLS, Pam-Stamp, Magma, Autoform, Permas

Source: Top20Auto, Ch.Tanasescu, SC'2003, www.top500.org
http://www.top500.org/lists/2003/11/Top20Auto_Top500V2.pdf

System Architectures in Automotive World Wide - 2003

Ex: Cluster of IA32 (1-2p per node)

62% are SMP based systems

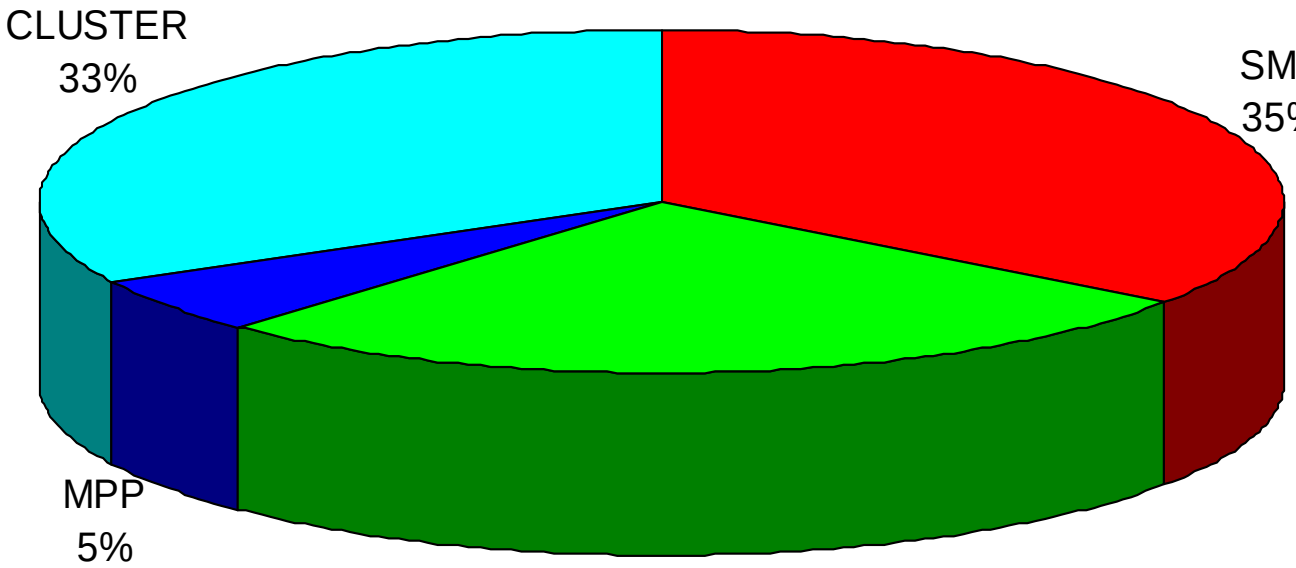

Ex: SGI Altix 3000,
IBM p690,
HP Superdome, NEC SX-6,
SGI Origin3000

Ex: IBM SP3,
Fujitsu VPP5000

Ex: Cluster of SMP nodes like
SGI Altix 3000, IBM p690, CPQ E45,
HP Superdome, SGI Origin3000

Notes

CONSTELLATION - cluster of SMP nodes

CLUSTER - cluster of IA32 (1 or 2p)

Source: Top20Auto, Ch.Tanasescu, SC'2003, www.top500.org
http://www.top500.org/lists/2003/11/Top20Auto_Top500V2.pdf

A History of Innovation of SGI

Jim Clark founded SGI on the vision of Computer Visualization

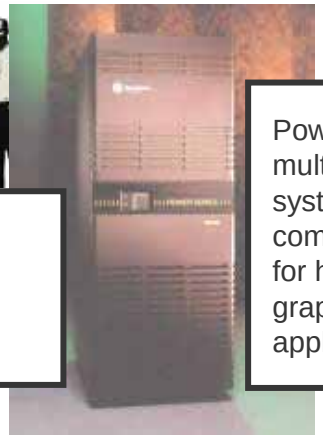

Power Series™, multi-processing systems provide compute power for high-end graphics applications

Challenge® XL media server fuels Steven Spielberg's Shoah project to document Holocaust survivor stories

SGI introduces its first 64-bit operating system

First systems deployed in Stephen Hawking's COSMOS system

NASA Ames and Altix® set world record for STREAMS benchmark

Altix®, first scalable 64-bit Linux® Server

1982 **1984** **1988** **1994** **1995** **1996** **1997** **1998** **2001** **2003** **2004**

IRIS® Workstations become first integrated 3D graphics systems

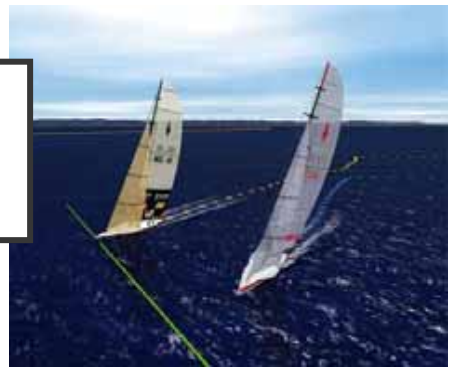

Dockside engineering analysis on Origin® 2000 and Indigo2™ helps Team New Zealand win America's Cup

First generation NUMA System: Origin® 2000

DOE deploys 6144p Origin 2000 to monitor and simulate nuclear stockpile

Introduced modular NUMAflex™ architecture with Origin® 3000

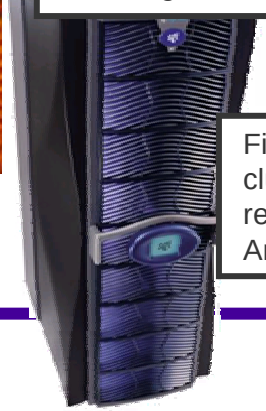

First 512p Altix cluster drives ocean research at NASA Ames

Images courtesy of Team New Zealand and the University of Cambridge

SGI Family of Scalable Linux® Solutions

SGI® Altix® 350
Servers and Clusters

SGI® Altix® 3000
Servers and Superclusters

Other vendors' Linux® Solutions

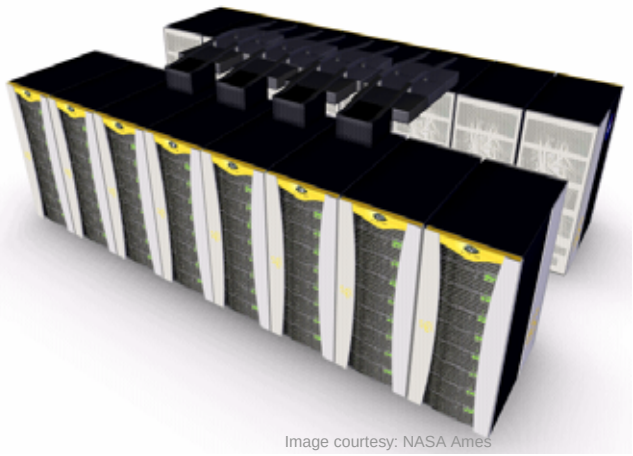

Image courtesy: NASA Ames

Desktops **Edge servers** **Mid-range** **Divisional** **Capability**

AA

Modular SGI® NUMAflex™ Architecture

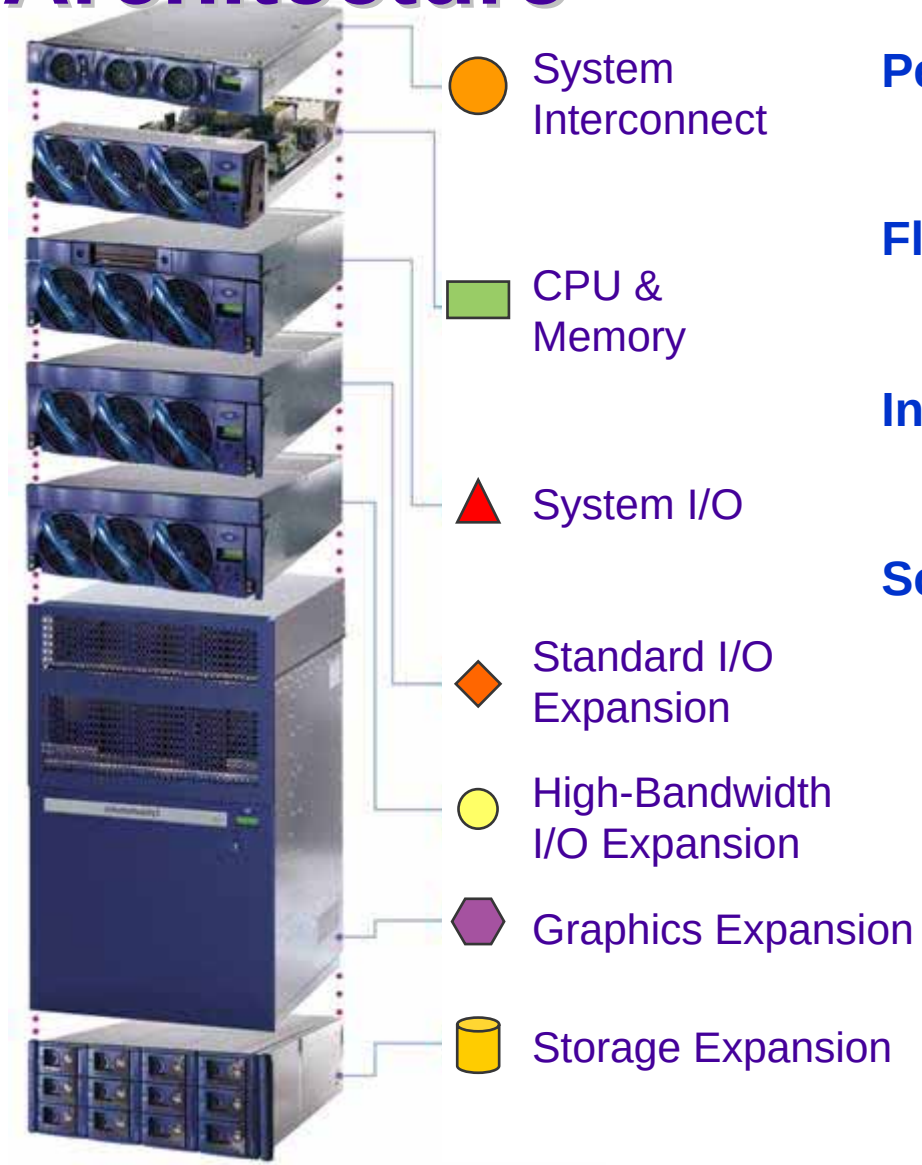

Performance: High-bandwidth interconnect with very low latency

Flexibility: Tailored configurations for different dimensions of scalability

Investment protection: Add new technologies as they evolve

Scalability: No central bus or switch; just modules and NUMAlink™ cables

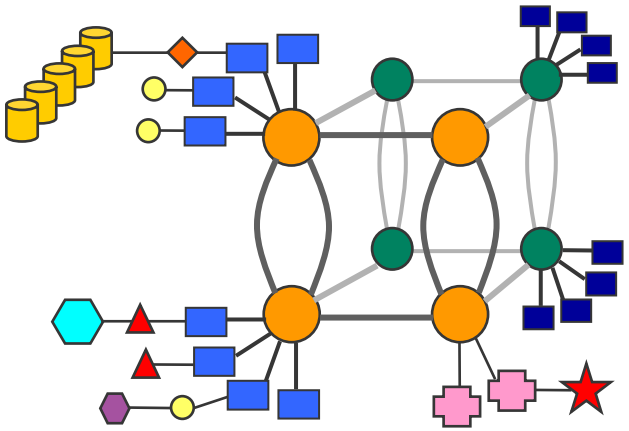

Silicon Graphics Prism™ Visualization System *Innovation that SCALES*

Power User Visualization	Small-Group Visualization	Ultimate Visualization Capability
2 - 4 GPUs, 2 - 8 CPUs	**4-8 GPUs, 8-16 CPUs**	**8-16 GPUs, 16-32 CPUs**
Delivers productivity	***Delivers breakthroughs***	***Delivers ultimate advantage***

Image courtesy of Laboratory for Atomospheres, NASA Goddard Space Flight Center

Silicon Graphics Prism™ Visualization System

SGI® NUMAflex™ independently scales all resources

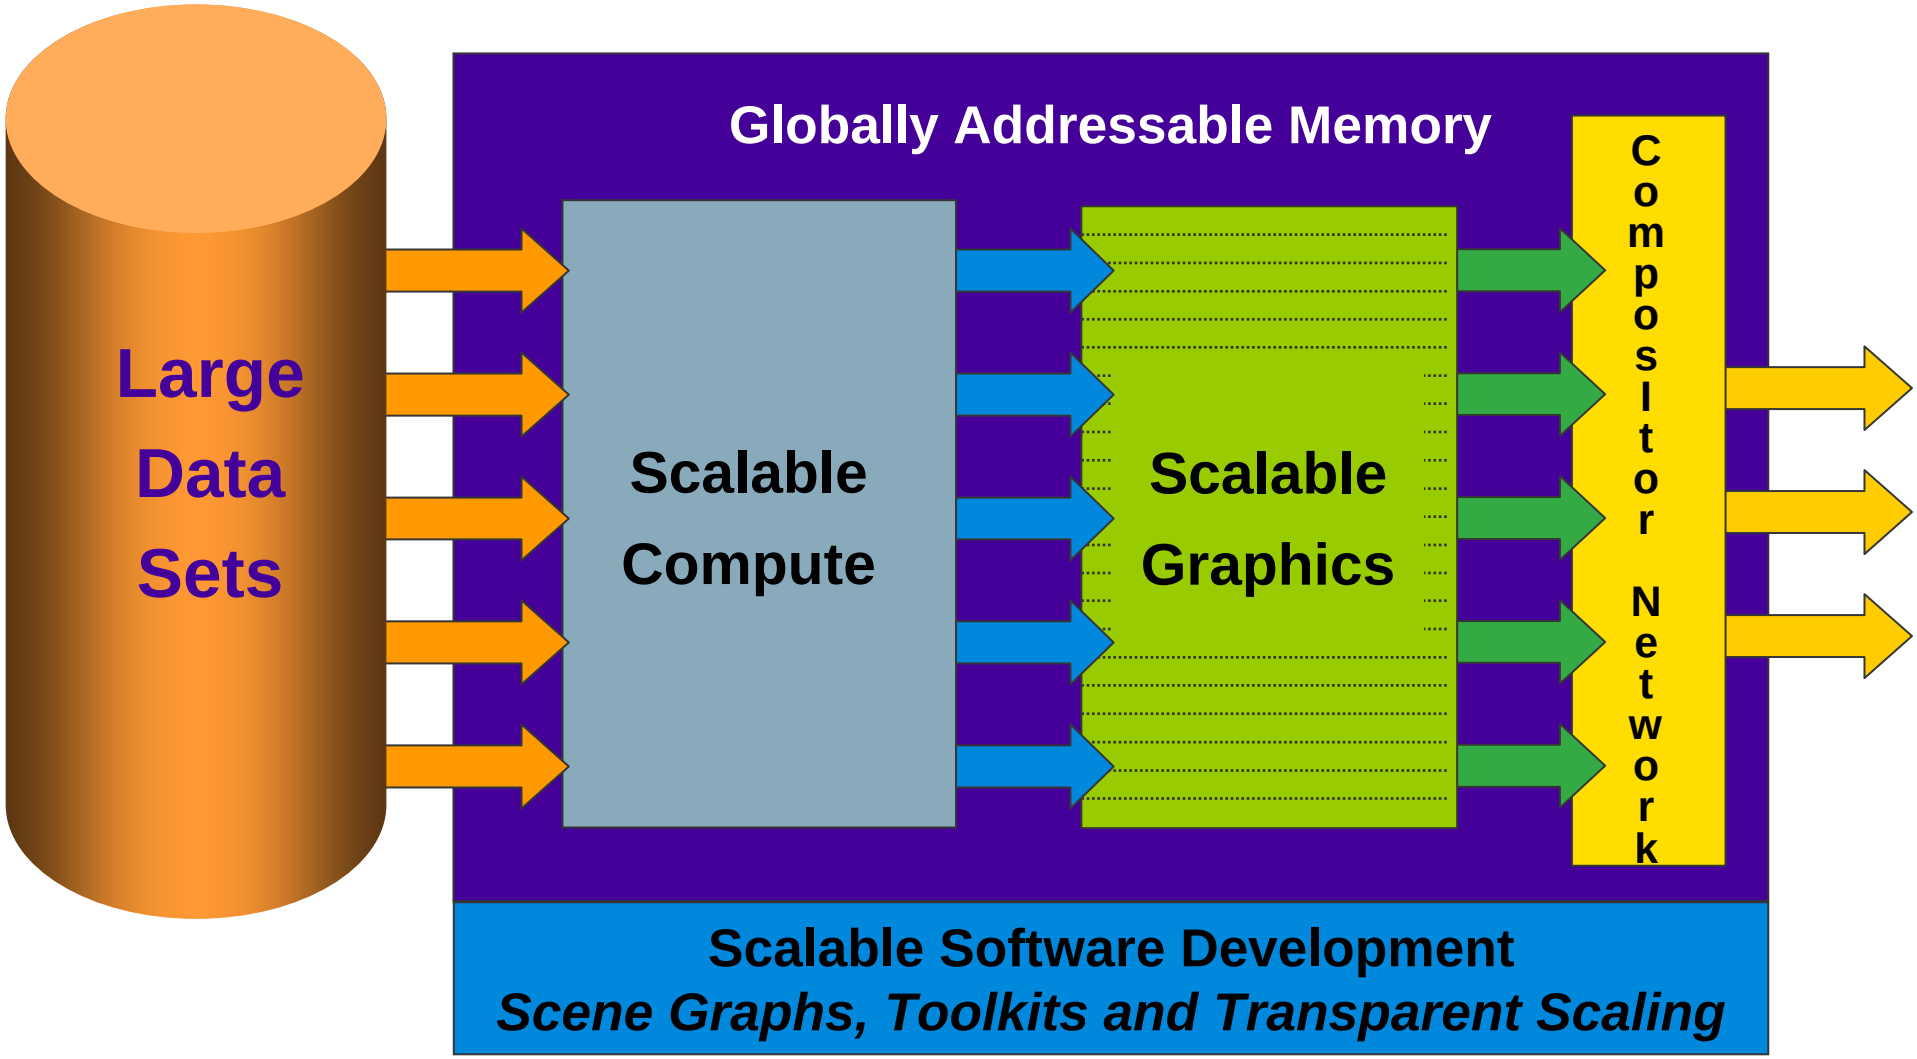

SGI Altix 3700 and 350 Servers With Global Addressable Memory

Altix™ 3700

- Supercomputer, supercluster configurations
- SSI up tp 512 processors,
- Supercluster up to 2048p via NUMAlink
- Up to 8TB of shared memory
- 6.4GB/sec dual plane fat tree → 12.8GB/sec with NUMAlink 4 router
- 1.3GHz / 3.0MB to 1.6GHz / 9.0MB Intel Itanium2 processors
- Infinitely scalable using commercial interconnects

Image courtesy: NASA Ames

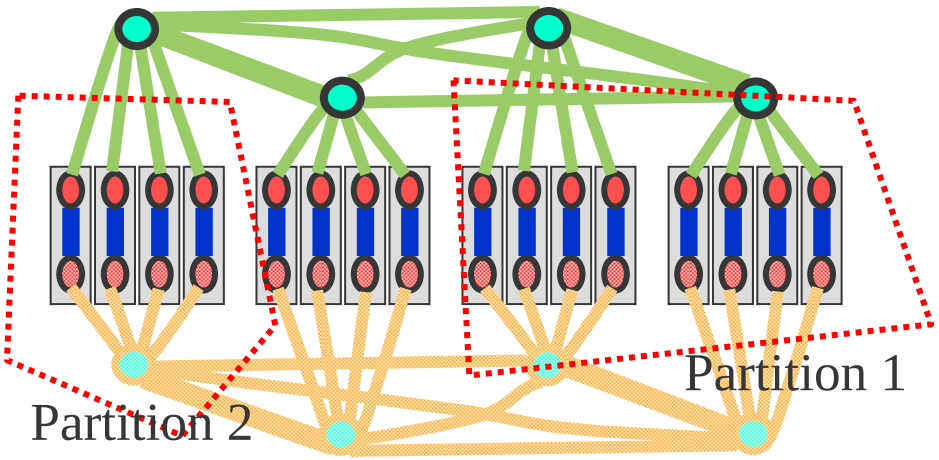

Altix™ 350

Infiniband®/Altix 350 Clusters

- **Competitive advantages:**
  - **large host size (16 P)**
  - **SGI ProPack software features and support**
  - **MPT MPI**

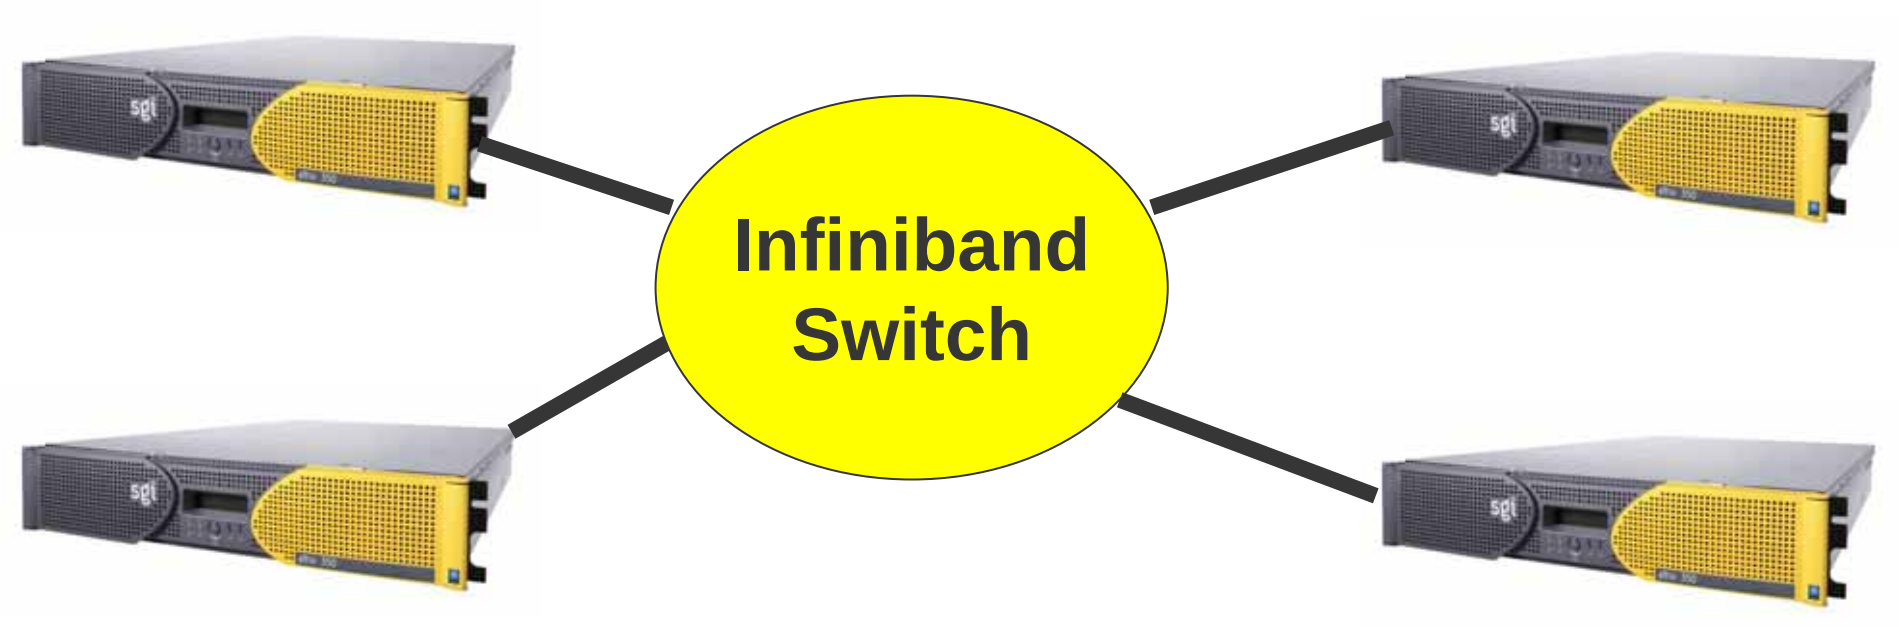

SGI Production Linux Solution Stack

SGI ProPack™ HPC Value-Add Enhancements **Optimized high-productivity computing**	System management: Partitioning (Altix 3000), Performance Co-Pilot Resource management: CPU sets/memory placement, MPT, array services, SCSL math libraries Data management: XVM
SGI® Open-Source Enhancements **Enabling features and functionality**	Contributing SGI expertise in scalability, NUMA XFS® high-performance filesystem
Standard Linux® Distribution **Runs standard 64-bit Linux applications**	Red Hat® Enterprise Linux® AS 3.0* compatible for highest performance Certified SUSE Linux 9 option for highest certified support levels Easy to develop and administer Intel® compilers and tools

HPC Resource Demands for MCAE

Generally Five Types of MCAE Computational Behavior to Consider:

MCAE Segment	Software	System Resource Benefits/Requirements					
		CPU	Memory BW	I/O	Bandwidth	Latency	Scalability
IFEA Statics	ABAQUS® ANSYS® MSC.NASTRAN™	H	H	M	L	L	< 10p
IFEA Dynamics	ABAQUS® ANSYS® MSC.NASTRAN™	M	H	H	H	L	< 10p
EFEA	LS-DYNA®	H	M	L	M	H	~ 64p
CFD Unstructured	FLUENT® STAR-CD™ PowerFLOW®	M	H	M	H	H	~ 100p
CFD Structured	FLOWER OVERFLOW	H	H	L	M	M	> 100p

LD-Dyna mpp970 Communication Matrix

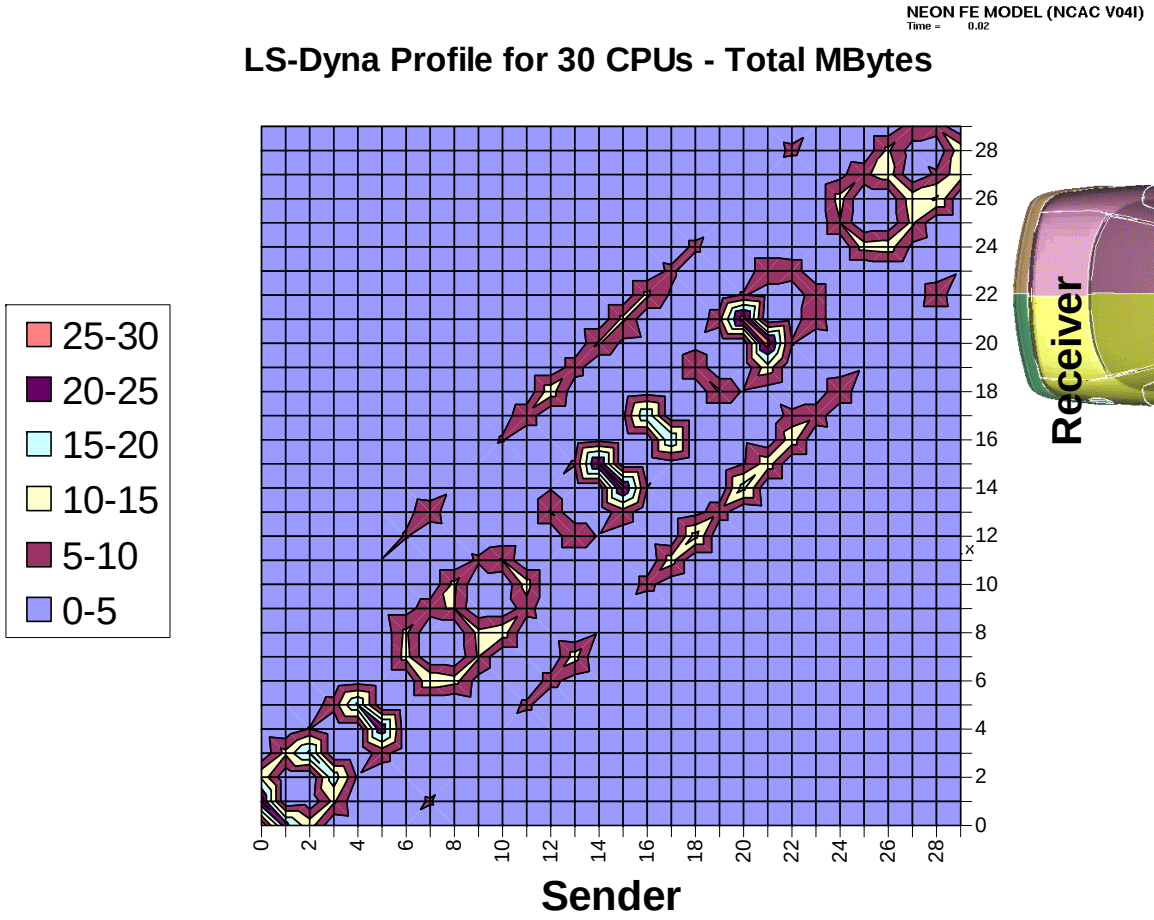

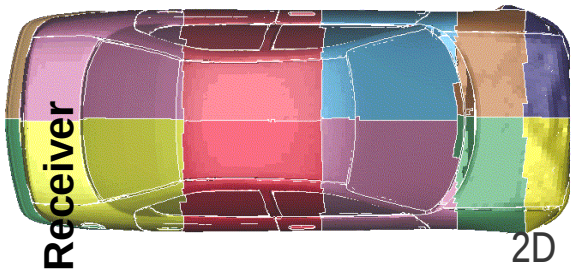

1D

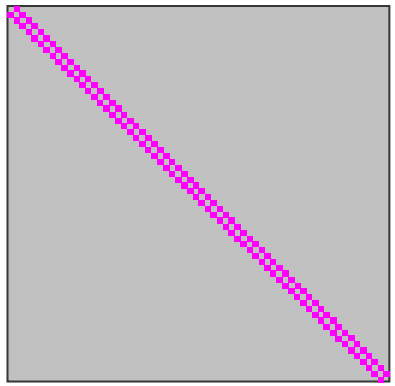

2D

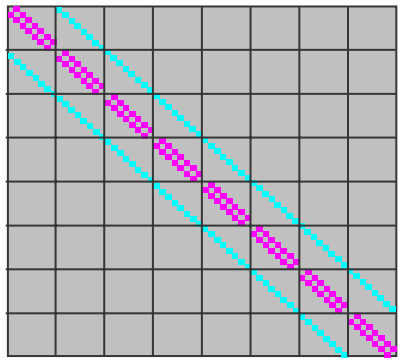

3D

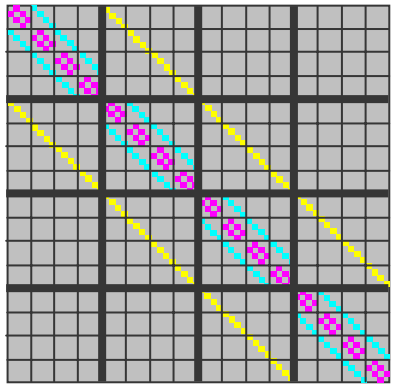

90% or messages shorter than 3kbyte , Latency is important for scalability

Interconnect Contention depends on Topology

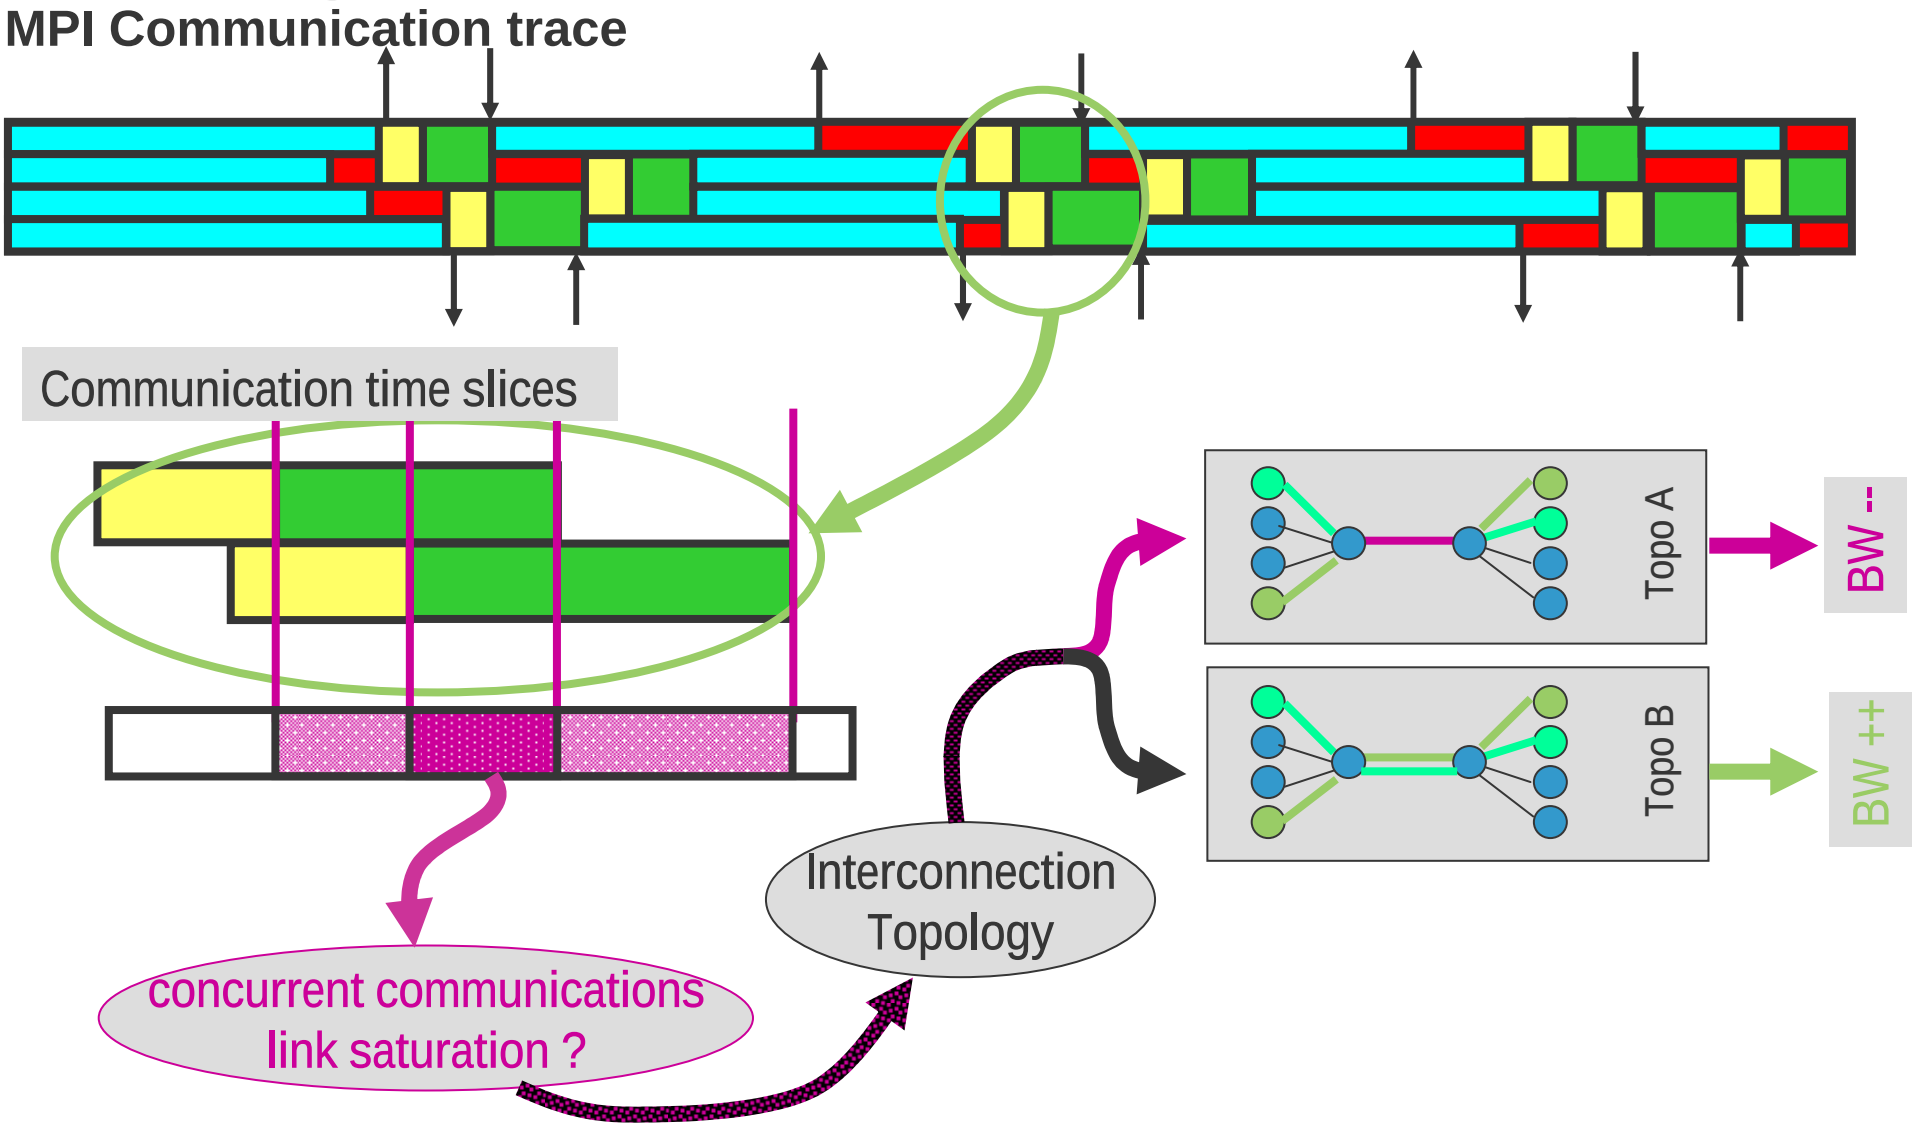

Scalable CPU/Memory Interconnect Fabric

Scalable Compute and Large Memory

Scalable Compute

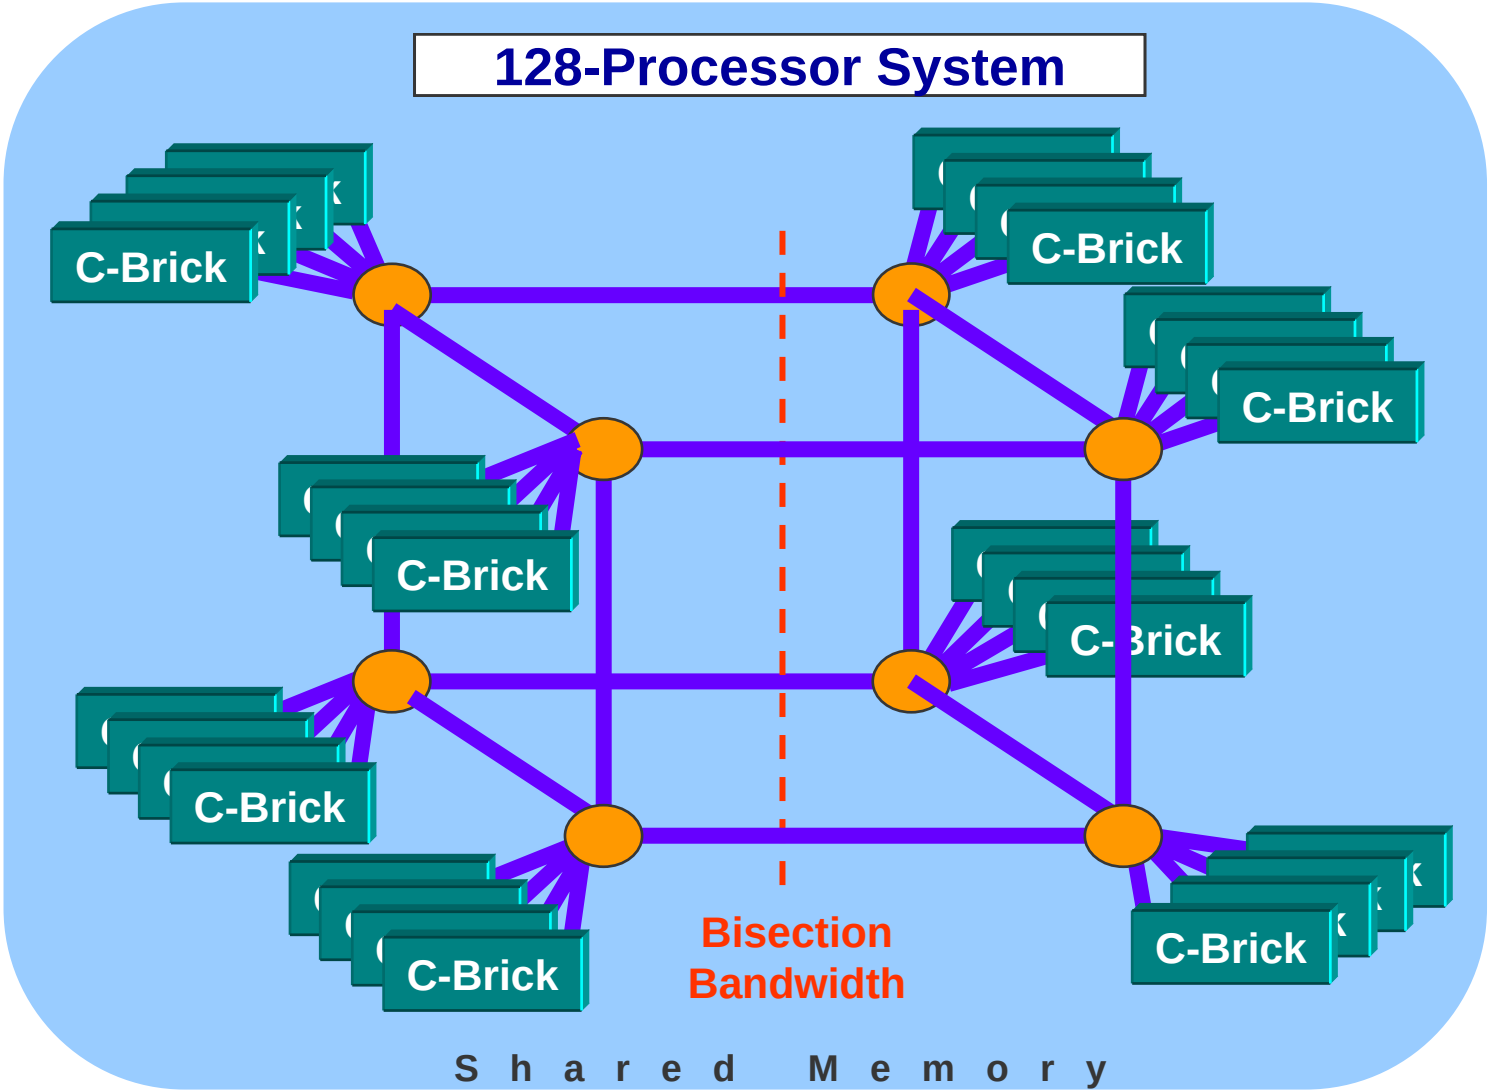

Total system bisection bandwidth actually scales as system grows!

Skalierbare Leistung im Vergleich

CPUs	PC Cluster 3.06GHz			Altix 1.3GHz			PCcluster/ Altix
	Elapsed Time		Scalability	Elapsed Time		Scalability	
	sec	hours		sec	hours		
4	56479	15:41:19	1.00	48611	13:30:11	1.00	1.16
8	30461	8:27:41	1.85	25463	7:04:23	1.91	1.20
16	18157	5:02:37	3.11	13012	3:36:52	3.74	1.40
32	14610	4:03:30	3.87	7220	2:00:20	6.73	2.02
64	13162	3:39:22	4.29	4456	1:14:16	10.9	2.95

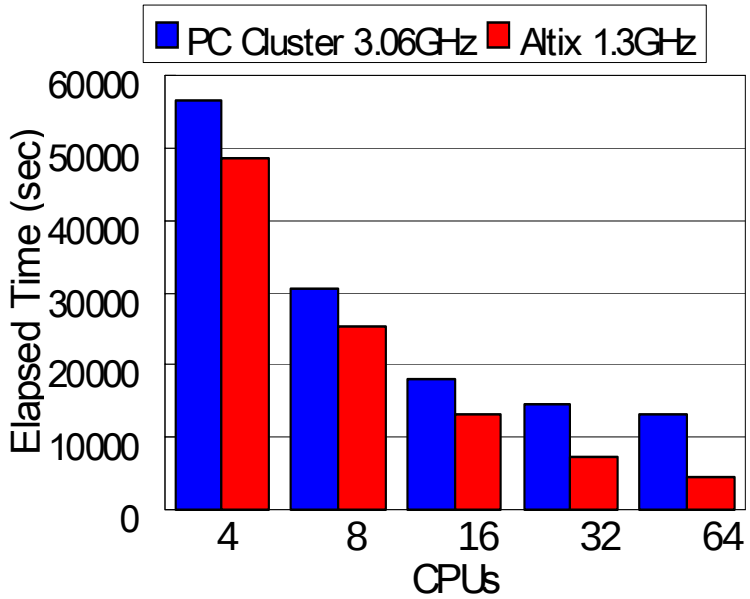

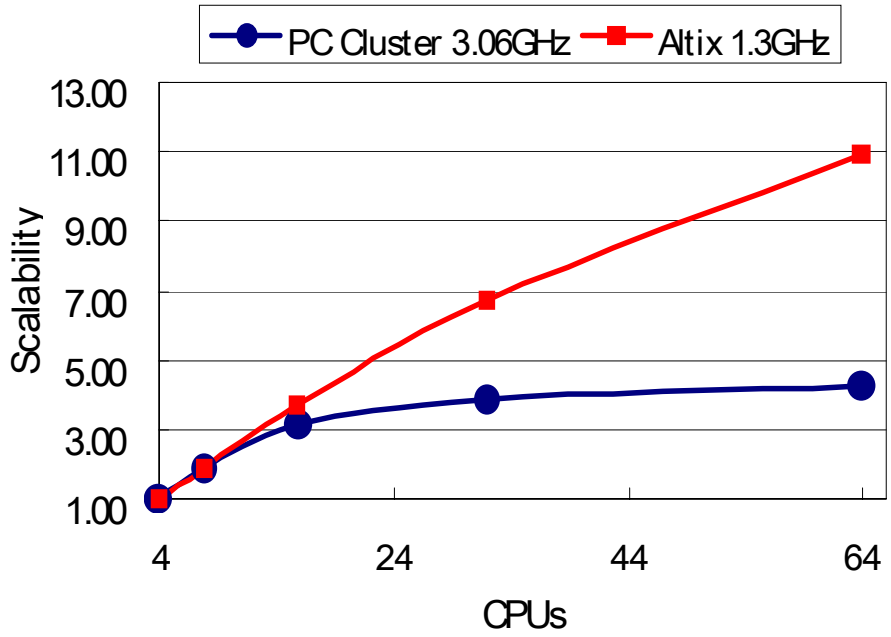

Beispiel: LS-DYNA Stamping

> 800.000 ElementeSimulation Time = 21 msec

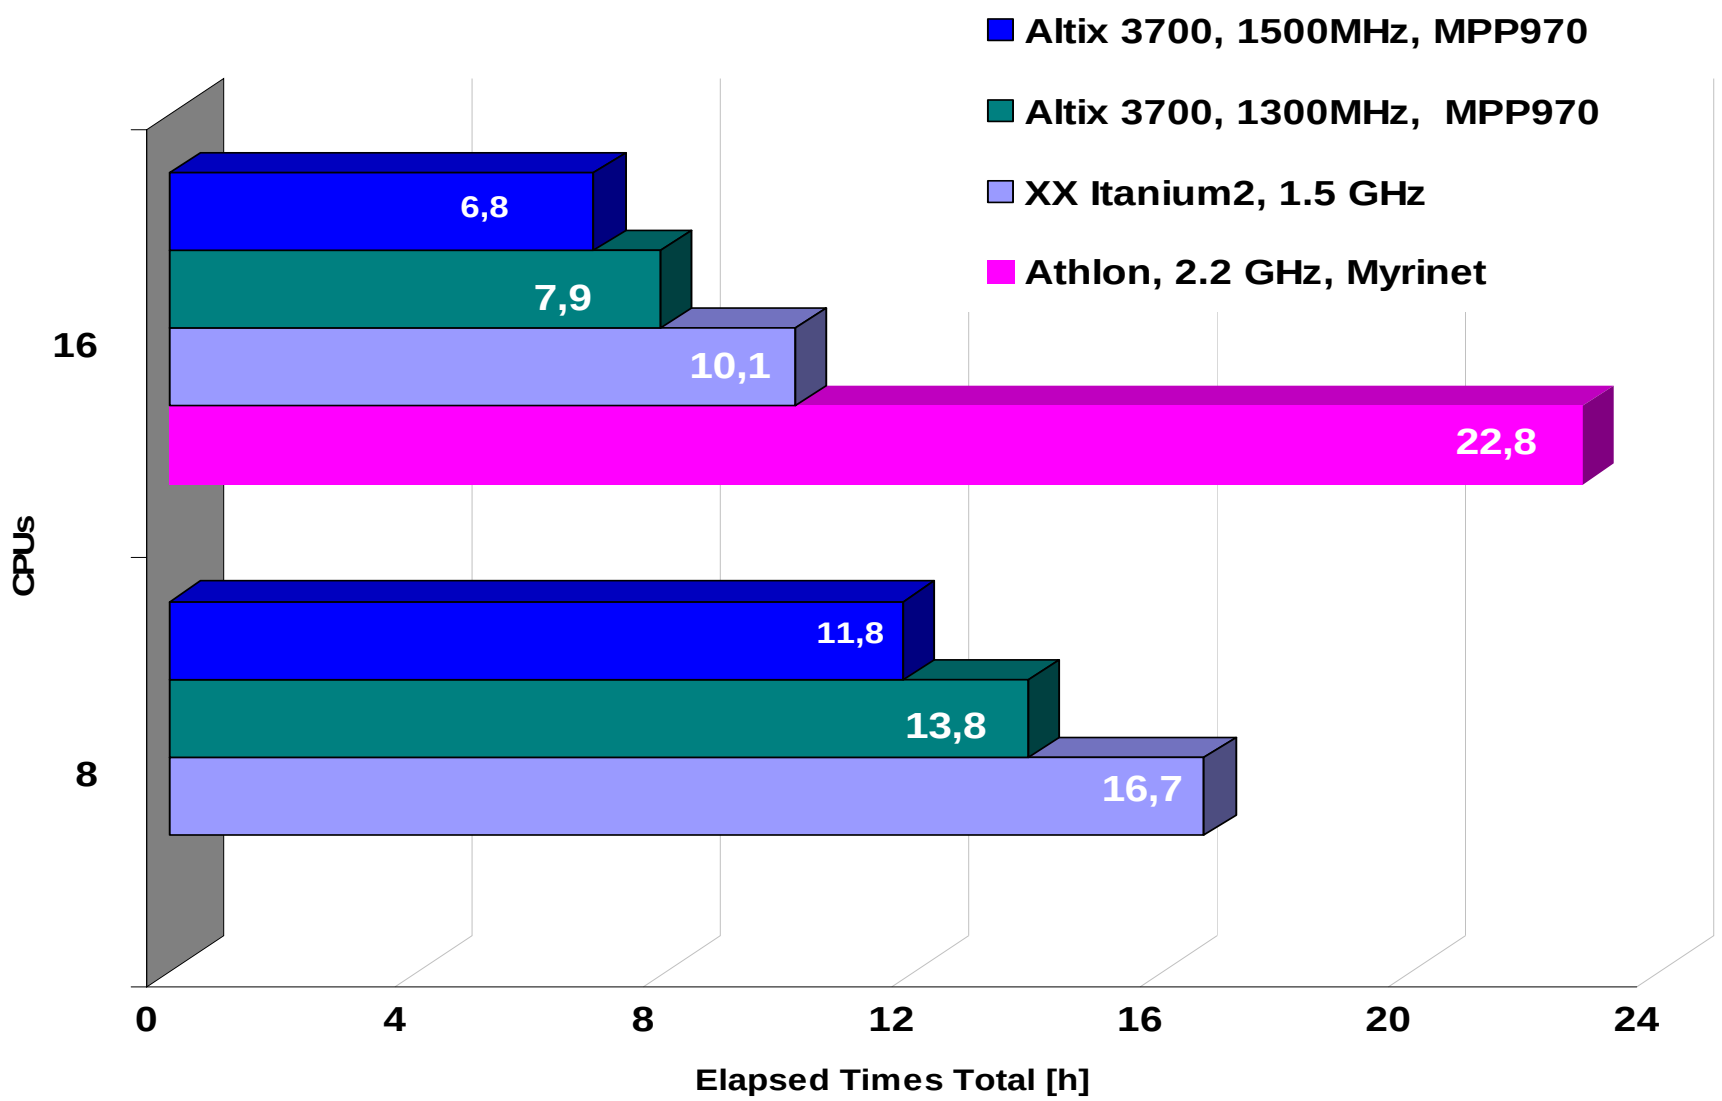

Beispiel: LS-DYNA Stamping

> 800.000 ElementeSimulation Time = 21 msec

Performance Relation vs. 16@2.2GHz, Athlon

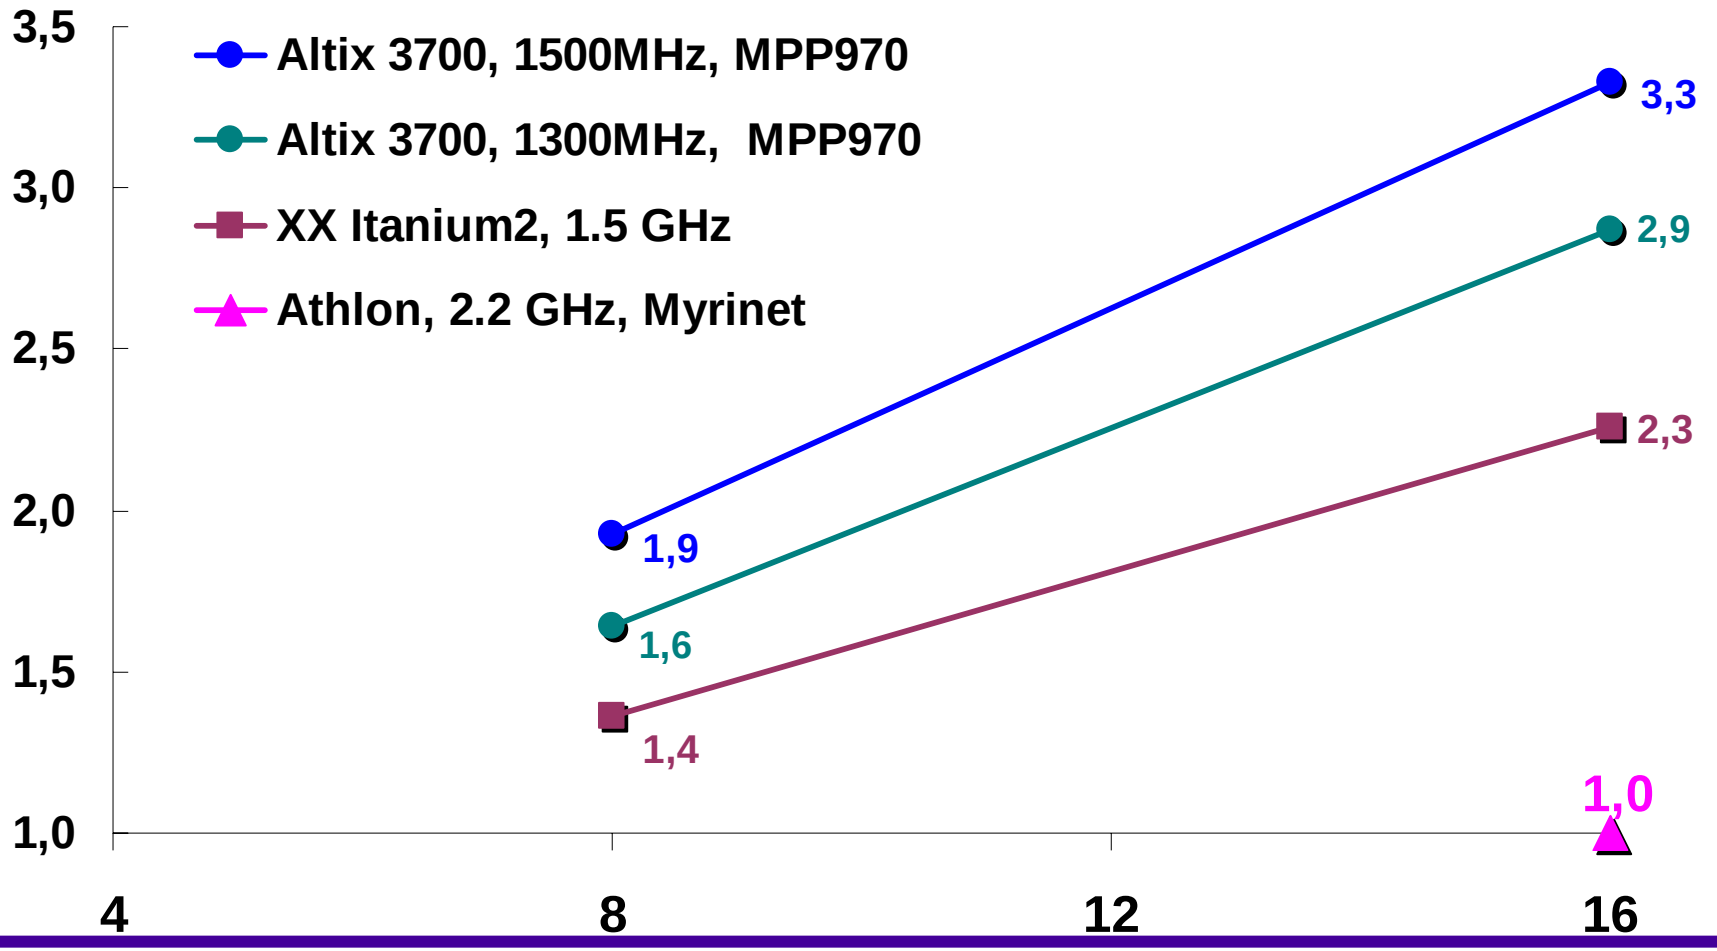

SGI® Altix™ und ccNUMA im CAE

- **Der 64-bit Adressraum der Itanium®2 CPUs erlaubt CAE-Anwendungen mit Speicherbedarf >> 4 GB**
  - **kurzfristige Erweiterung (höhere Diskretisierungen, größere Modelle) des virtuellen und physikalischen Speichers ohne Einschränkungen und Systemwechsel (HW+SW),**
  - **die performante Nutzung von In-Core-Solver >> 4 GB**
  - **das Testen und Debuggen von „Problem-Modellen“ auf 1 CPU im SMP und MPP-Modus:**
- **Hohe I/O-Bandbreiten + SGI-Filesystem XFS beschleunigen Out-of-Core Solver im NVH-Bereich.**
- **Numaflex + optimierte OS-System Funktionen ergeben beste Durchsatzleistungen (s. SPECfp_rate2000)**

New Discoveries with NASA ECCO
Estimating the Circulation and Climate of the Ocean

The Challenge

- Assimilate real-time satellite data directly into into global ocean circulation models
- Drive resolution from today's .25 degree (~25km) to .1 degree

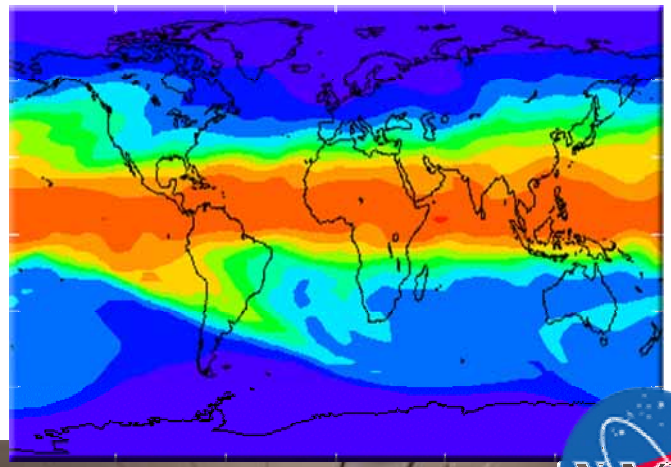

The Solution

- 512p Altix®, 28TB storage
- Reducing typical ocean simulation times from months to 2-3 days
- World Record STREAM Triad benchmark result, first system to break 1,000 GB/sec

Optimizing Decisions at Škoda Auto

Škoda Auto
Mladá Boleslav, Czech Republic

The Challenge

Reduce processing times for crash analysis and fluid dynamics analysis

The Solution

- Deployed 96p Altix® 3700 cluster, 16TB InfiniteStorage disk array
- Reduced processing times allow Skoda to optimize decisions around car development and design efficiency, quality and safety

Images courtesy of Škoda Auto

Altix®: Breakthrough Deployments

Breakthrough applications

Chemistry and Biology
Gaussian, Amber, Blast, Charm, Accelrys and more

Oak Ridge National Lab, University of Minnesota Supercomputing Institute, National Cancer Institute, many universities

Manufacturing
LS-Dyna, Abaqus, Star-CD, NASTRAN, Ansys and more

DaimlerChrysler,VW, Skoda, many more

Physical Sciences
Modeling of oceans, weather, earth, astronomy, energy, and more

COSMOS, NASA, University of Queensland, Total, Japan Institute of Statistical Mathematics, many more

Breakthrough technologies

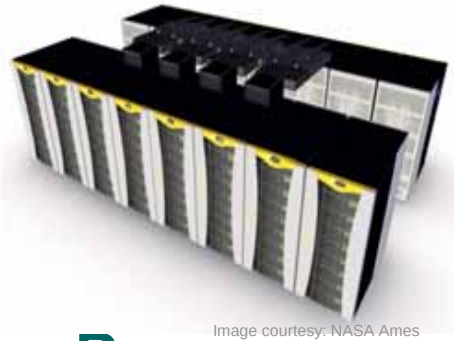

Image courtesy: NASA Ames

Large shared memory

Oak Ridge National Labs (2TB), Japan Institute of Statistics and Mathematics (1.96TB), NRL (1TB), and others

Large processor scalability

NASA Ames (512p), NRL (384p), SARA (416p), Total (552p), and others

Grid computing

COSMOS, Netherlands National Supercomputing Facility at SARA, and others

Hybrid solutions

Wichita State University, BP, and others

Proven performance

Early adopters expanding their deployments

Total, NASA Ames, NRL, and others

Images courtesy of the University of Cambridge, Molecular Simulations Inc., Skoda Auto, and NASA Ames

Zusammenfassung

- Leistungsfähigstes Linux® Computersystem der Welt
- Intel® Itanium®2 CPUs mit SGI® NUMAflex™ Architektur
- Single System Image mit bis zu 512 CPUs
- Shared-Memory Skalierbarkeit bis zu 8TB
- MPI mit Shared Memory Features
- Linux Erweiterungen für HPC Workflow optimiert
- Höchste Bandbreite und kürzeste Latenzzeiten
- Standard Linux Distribution und Entwicklungsplatform
- Differenzierte “Middleware” und Treiber
- Hohe Verfügbarkeit von 64bit Anwendungen
- Integration von skalierbarer Graphik möglich