

# PERFORMANCE of VARIOUS COMPUTERS in COMPUTATIONAL CHEMISTRY

14<sup>th</sup> Daresbury Machine Evaluation Workshop  
CCLRC Daresbury Laboratory

**Martyn F. Guest**

**Computational Science and Engineering Department**

**m.f.guest@dl.ac.uk**

## Acknowledgements:

M. Deegan (HP/Compaq),


M. Ehrig (HP/Compaq)

# OUTLINE

- Processor Performance Overview
  - Single-processor performance, Performance Metrics
- SPEC (Standard Performance Evaluation Corporation)
  - SPEC 95 and **SPEC CPU 2000**
- Computational Chemistry Benchmark (serial) - SPECfp ?
  - Matrix and application “kernels”
  - Application packages (GAMESS-UK, DL\_POLY)
  - Comparison involves 200+ computers (vector supercomputers, workstations, PCs and MPP nodes)
- Now extended to include “RATE” benchmarks
  - URLs: Powerpoint presentation and Paper:  
*<http://www.cse.clrc.ac.uk/disco/hw-perf.shtml>*
- *Benchmarking of Parallel Commodity Systems*


# MEW13 (2002) - Summary PI relative to the HP RX2600 Itanium2/1000

The MATRIX-97, Chemistry  
Kernels & GAMESS-UK  
Benchmarks



# Fortran Compilers - Performance

Performance relative to GNU g77 on Pentium 4 Xeon/2666



# MACHINES UNDER EVALUATION

<u>Machine</u>	<u>Processor</u>
AMD Athlon <u>AMD Opteron</u>	MP 1.667 GHz, 2 GHz <u>848 (2.2 GHz), 246 (2.0 GHz), 244 (1.8 GHz), 842 (1.6 GHz), 840 1.4 GHz</u>
<u>Intel IA32</u>	Pentium 4 2.5, 2.66 GHz, <u>3.06 GHz</u>
SUN Blade 2000 / 1056 Cu SUN Fire V880 / 900 Cu	UltraSPARC-3 / 1056 MHz UltraSPARC-3 / 900 MHz
HP PA-9000 / RP7410	PA8700+ / 875MHz
HP RX2600 (2-way) HP RX5670 (4-way) <u>HP RX4640 (4-way)</u> <u>HP RX2600 (2-way)</u> <u>HP RX5670 (4-way)</u> HP ZX6000	Itanium 2 / 1000 MHz (3 MB L3) Itanium 2 / 1000 MHz (3 MB L3) <u>Itanium 2 / 1300 MHz (3 MB L3)</u> <u>Itanium 2 / 1500 MHz (6 MB L3)</u> <u>Itanium 2 / 1500 MHz (6 MB L3)</u> Itanium 2 / 900 MHz (1.5 MB L3)
Intel Tiger (4-way) <u>Intel Tiger (4-way) +</u> <u>Intel Tiger (4-way)</u>	Itanium 2 / 1000 MHz (3 MB L3) <u>Itanium 2 / 1200 MHz (3 MB L3)</u> <u>Itanium 2 / 1500 MHz (6 MB L3)</u>

<u>Machine</u>	<u>Processor</u>
HP/Compaq ES45 HP/Compaq ES45 Compaq Marvel	AXP A21264C / 1000 MHz AXP A21264C / 1250 MHz EV7 1000 MHz
<u>SGI Altix 3700</u> SGI Origin3800/R14k SGI Origin300/R14k SGI O2/R12k-SC	<u>Itanium 2 / 1300 MHz</u> R14000/R14010 600 MHz R14000/R14010 500 MHz R12000/R12010 270 MHz
IBM p-Series 630 IBM p-Series 690 <u>IBM p-Series 690</u> Intellisation Z Pro	RS/6000 /POWER4 1.0GHz RS/6000 /POWER4 1.3GHz <u>RS/6000 /POWER4 1.7GHz</u> Itanium / 800 MHz (4 MB L3)

### Vector Supercomputers

NEC SX-5, SX-4      Cray Y-MP/J90-10,  
FUJITSU VPP/300,      Cray YMP C98/4256

### MPP Nodes

Cray T3E/1200      AXP EV56 600 MHz  
IBM SP / Power3      RS/6000 WH2 - 375 MHz

# SPEC Benchmarks

- Compute intensive categories
  - integer versus floating point
  - conservative versus aggressive compilation
  - speed versus throughput
- Composite Metrics - SPEC CPU 2000

NOT:

- Graphics
- Network
- I/O

<http://www.specbench.org>

	SPEED	THROUGHPUT
Aggressive	SPECint2000 SPECfp2000	SPECint_rate2000 SPECfp_rate2000
Conservative	SPECint_base2000 SPECfp_base2000	SPECint_rate_base2000 SPECfp_rate_base2000

- SPECratio -  $T(\text{measured system}) / T(\text{reference})$
- Reference = 300 MHz Ultra 5/10 (100)
- SPECfp2000 - geometric mean of 14 ratios, one for each benchmark
- SPECint2000 - geometric mean of 12 ratios, one for each benchmark


# SPEC CPU 2000 - Floating point Benchmark Suite (SPECfp2000)

Benchmark	Language	Description
168.wupwise	F77	Physics: Quantum chromodynamics
171.swim	F77	Shallow water modelling
172.mgrid	F77	Multigrid solver: 3D potential field
173.applu	F77	Partial differential equations
177.mesa	C	3D graphics library
178.galgel	F90	Computational fluid dynamics
179.art	C	Image recognition / neural networks
183.quake	C	Seismic wave propagation simulation
187.facerec	F90	Image processing: Face recognition
188.amp	C	Computational chemistry
189.lucas	F90	Number theory / primality testing
191.fma3d	F90	Finite-element crash simulation
200.sixtrack	F77	Nuclear physics accelerator design
301.apsi	F77	Metereology: Pollutant distribution

Reference: 300 MHz Ultra 5/10 = 100


# SPEC CPU 2000 - SPECfp2000

Values relative to IBM p-series 690/pwr4 1.3 GHz




# SPEC CPU 2000 - SPECfp2000\_base

Values relative to IBM p-series 690/pwr4 1.3 GHz




# SPEC CPU 2000 - SPECint2000

Values relative to IBM p-series 690/pwr4 1.3 GHz



# IBM p-series 690Turbo:Multi-chip Module (MCM)


Four POWER4 chips (8 processors) on an MCM, with two associated memory slots



*L3 cache shared 4 GX Bus links for external connections across all processors*


# SPEC CPU 2000 - SPECfp2000

## Units of SPECfp as a function of Peak Performance



# SPEC CPU 2000 - SPECfp vs SPECfp\_rate (32 CPUs)

Values relative to IBM 690 Turbo 1.7 GHz



# Computational Chemistry Benchmark Suite

- Benchmark suite developed to incorporate:
  - Matrix “kernels” (MATRIX-89 and MATRIX-97)
  - Application “kernels”
  - Application packages (e.g. GAMESS-UK, DL\_POLY)
- Implemented on Supercomputers, servers (superminis), workstations, PCs and parallel machines
  - Matrix Operations / Matrix multiplication and matrix diagonalisation
  - Computational Chemistry Kernels - four typical application kernels (direct-SCF, MD, QMC and Jacobi eigen solver)
  - STREAM (memory bandwidth)
  - Quantum Chemistry Calculations - twelve typical applications, including SCF, direct-SCF, CASSCF, MCSCF, direct-CI and MRD-CI, MP2, 2nd derivatives (GAMESS-UK-89 and GAMESS-UK-99)
  - Molecular Dynamics Calculations - six typical simulations

# COMPUTATIONAL CHEMISTRY BENCHMARKS I. Matrix Operations

## MATRIX OPERATIONS (MATRIX-89 and MATRIX-97)


- SPARSE Matrix Multiply BenchMark
  - MMO operation is central to the efficient operation of modern QC codes. In this benchmark a series of MMOs ( $R = A \times B$ ) are performed involving matrices of increasing order:
    - MATRIX-89: 10, 20, 30, ... , 100 (B is sparse)
    - MATRIX-97: 50, 100, 150, ... , 500 (B is sparse)
- Diagonalisation Benchmark
  - Based on diagonalising a series of real symmetric matrices. Measures the performance of 8 routines from mathematical libraries and QC codes:
    - MATRIX-89: 10, 20, 30, ... , 100
    - MATRIX-97: 50, 100, 150, 200, 250, 300
- Q<sup>+</sup>HQ Benchmark
  - Designed to extend MMO benchmark by allowing for the use of library routines e.g. BLAS. Uses both a scalar and vector algorithm:
    - MATRIX-89: 10, 20, 30, ... , 150
    - MATRIX-97: 20, 40, 60, ... , 300

Memory bandwidth

-lblas


# Matrix-97: SPARSE MMO Benchmark.

Performance relative to the IBM p-series 690/pwr4 1.3 GHz




# Matrix-97: Q<sup>†</sup>HQ MMO Benchmark.

Performance relative to the IBM p-series 690/pwr4 1.3 GHz




# Matrix-97: Diagonalisation Benchmark

Performance relative to the IBM p-series 690/pwr4 1.3 GHz



# The Matrix-97 Benchmarks.

Performance relative to the IBM p-series 690/pwr4 1.3 GHz



# COMPUTATIONAL CHEMISTRY KERNELS

Programs that are realistic models of actual chemical applications or algorithms

## ■ Self consistent Field (SCF)

- SCF code using distributed primitive 1s gaussian functions as a basis (thus emulating the use of s, p, functions); performs direct-SCF calculation on Be<sub>4</sub> (60 functions).

## ■ Molecular Dynamics (MD)

- This code bounces a few thousand argon atoms around in a box with periodic boundary conditions. LJ pair-wise interactions are used with integration of the Newtonian equations of motion.

## ■ Quantum Monte Carlo (QMC)


- This code evaluates the energy of the simplest explicitly correlated electronic wavefunction for the He atom using a variational monte-carlo method without importance sampling.

## ■ Jacobi iterative linear equation solver (JACOBI)

- JACOBI uses a naive jacobi iterative algorithm to solve a linear equation. All the time is spent in a large matrix-vector multiplication.


# Computational Chemistry Kernels - Direct-SCF.

Performance relative to the IBM p-series 690/pwr4 1.3 GHz




# Chemistry Kernels - Molecular Dynamics.

Performance relative to the IBM p-series 690/pwr4 1.3 GHz




# Chemistry Kernels - Quantum Monte Carlo.

Performance relative to the IBM p-series 690/pwr4 1.3 GHz




# Chemistry Kernels - Jacobi Solver.


Performance relative to the IBM p-series 690/pwr4 1.3 GHz



# STREAM: Measured Sustainable Memory Bandwidth in HPC (TRIAD)

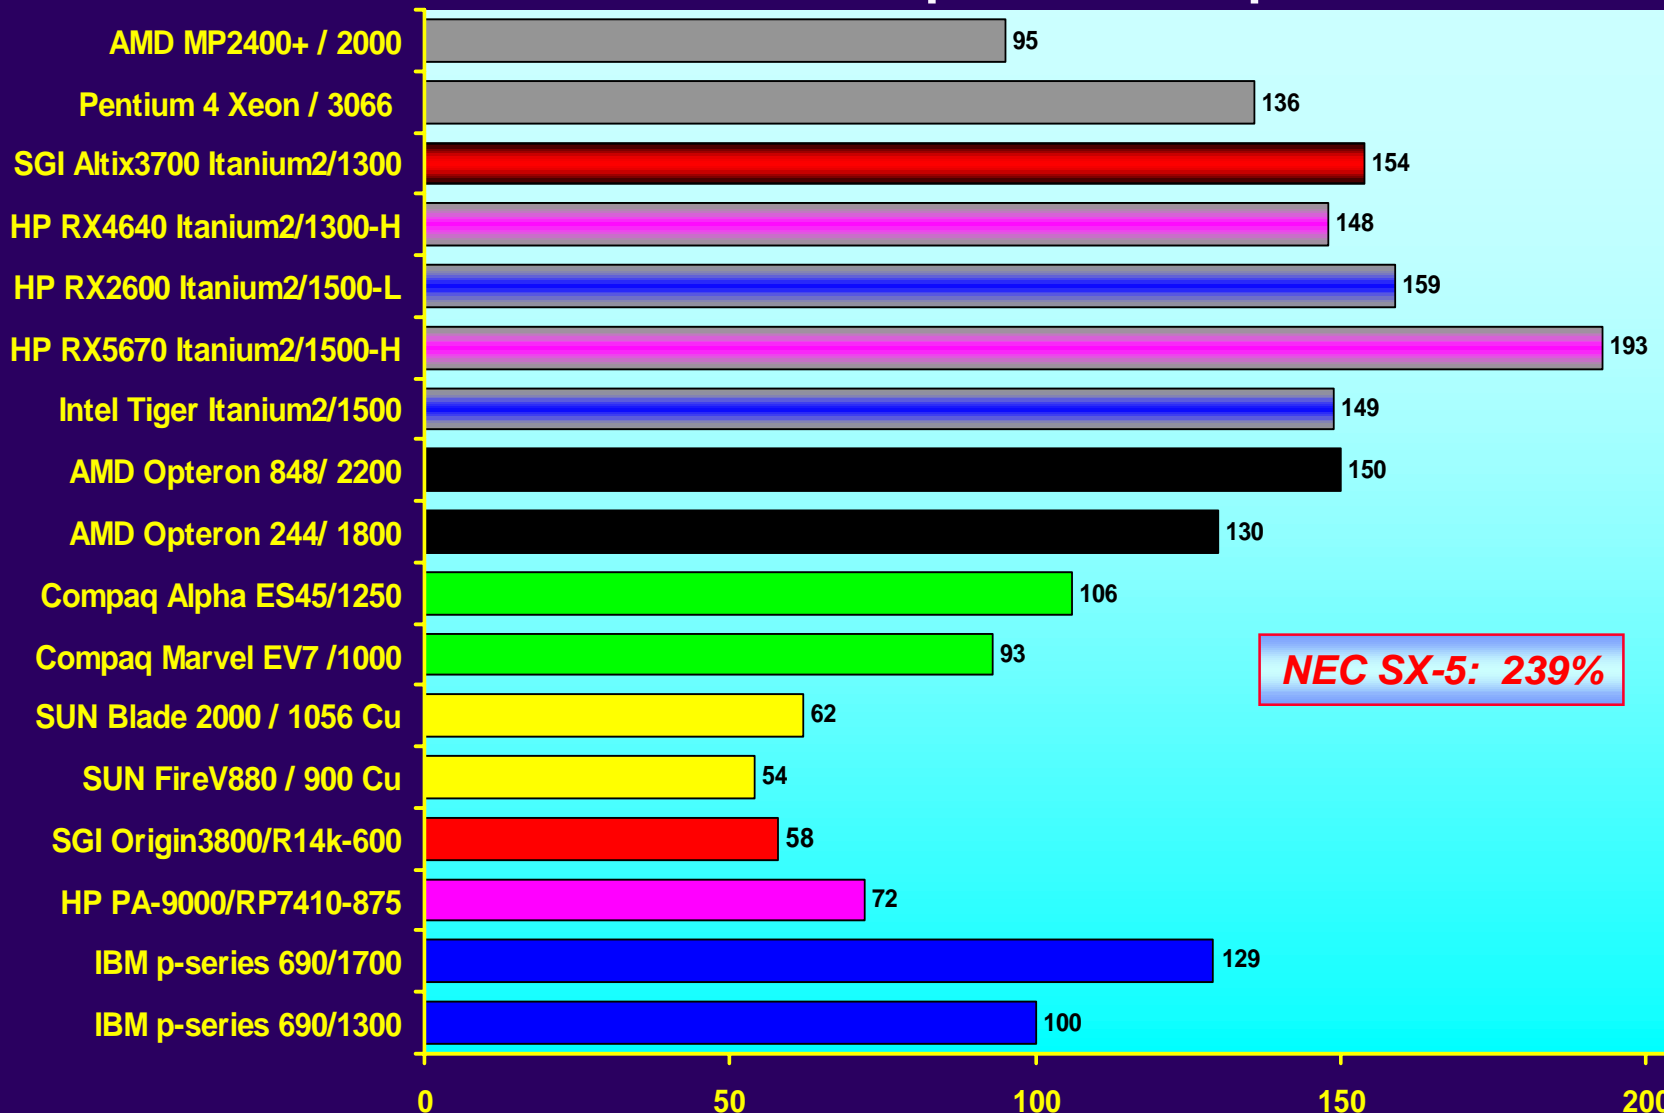


# STREAM: Sustainable Memory Bandwidth (TRIAD) per process



# Computational Chemistry Kernels.

Performance relative to the IBM p-series 690/pwr4 1.3 GHz



# GAMESS-UK and DL\_POLY Benchmarks

## GAMESS-UK

## DL\_POLY

### 12 Typical QC Calculations


<u>Module</u>	<u>Basis (GTOs)</u>	<u>Species</u>
1. SCF	STO-3G (124)	Morphine
2. SCF	6-31G (154)	C <sub>6</sub> H <sub>3</sub> (NO <sub>2</sub> ) <sub>3</sub>
3. ECP Geometry	ECPDZ (70)	Na <sub>7</sub> Mg <sup>+</sup>
4. Direct-SCF	6-31G (82)	Cytosine
5. CAS-geometry	TZVP (52)	H <sub>2</sub> CO
6. MCSCF	EXT1 (74)	H <sub>2</sub> CO
7. Direct-CI	EXT2 (64)	H <sub>2</sub> CO/H <sub>2</sub> +CO
8. MRD-CI (26M)	ECP (59)	TiCl <sub>4</sub>
9. MP2-geometry	6-31G* (70)	H <sub>3</sub> SiNCO
10. SCF 2nd derivs.	6-31G (64)	C <sub>5</sub> H <sub>5</sub> N
11. MP2 2nd derivs.	6-31G* (60)	C <sub>4</sub>
12. Direct-MP2	DZP (76)	C <sub>5</sub> H <sub>5</sub> N

### Six Typical Simulations


<u>Simulation</u>	<u>Atoms</u>	<u>Time steps</u>
1. Na-K disilicate glass	1080	300
2. Metallic Al with Sutton-Chen potential	256	8000
3. Valinomycin in 1223 water molecules	3837	100
4. Dynamic shell model water with 1024 sites	768	1000
5. Dynamic shell model MgCl <sub>2</sub> with 1280 sites	768	1000
6. Model membrane, 2 membrane chains, 202 solute and 2746 solvent molecules	3148	1000

# The GAMESS-UK Benchmark I. CPU

Performance relative to the IBM p-series 690/pwr4 1.3 GHz



# SPECfp2000 and the GAMESS-UK Benchmark



**Performance relative to the IBM p-series 690/1300**

# The GAMESS-UK-99 Benchmark


## 10 Typical QC Calculations

<u>Module</u>	<u>Basis (GTOs)</u>	<u>Species</u>
1. Direct- SCF	6-31G (227)	Morphine
2. SCF	6-31G** (265)	C <sub>6</sub> H <sub>3</sub> (NO <sub>2</sub> ) <sub>3</sub>
3. DFT B3LYP	6-311G* (167)	Cytosine
4. MCSCF	CC-PVTZ (100)	H <sub>2</sub> CO
5. Direct-CI	CC-PVTZ (100)	H <sub>2</sub> CO/H <sub>2</sub> +CO
6. CCSD(T)	TZV+2d+1f (144)	C <sub>4</sub>
7. MP2-geometry	TZVP (105)	H <sub>3</sub> SiNCO
8. SCF 2nd derivs.	6-311G** (144)	C <sub>5</sub> H <sub>5</sub> N
9. MP2 2nd derivs.	TZVP(C2d) (104)	C <sub>4</sub>
10. Direct-MP2	6-31G* (130)	Cytosine

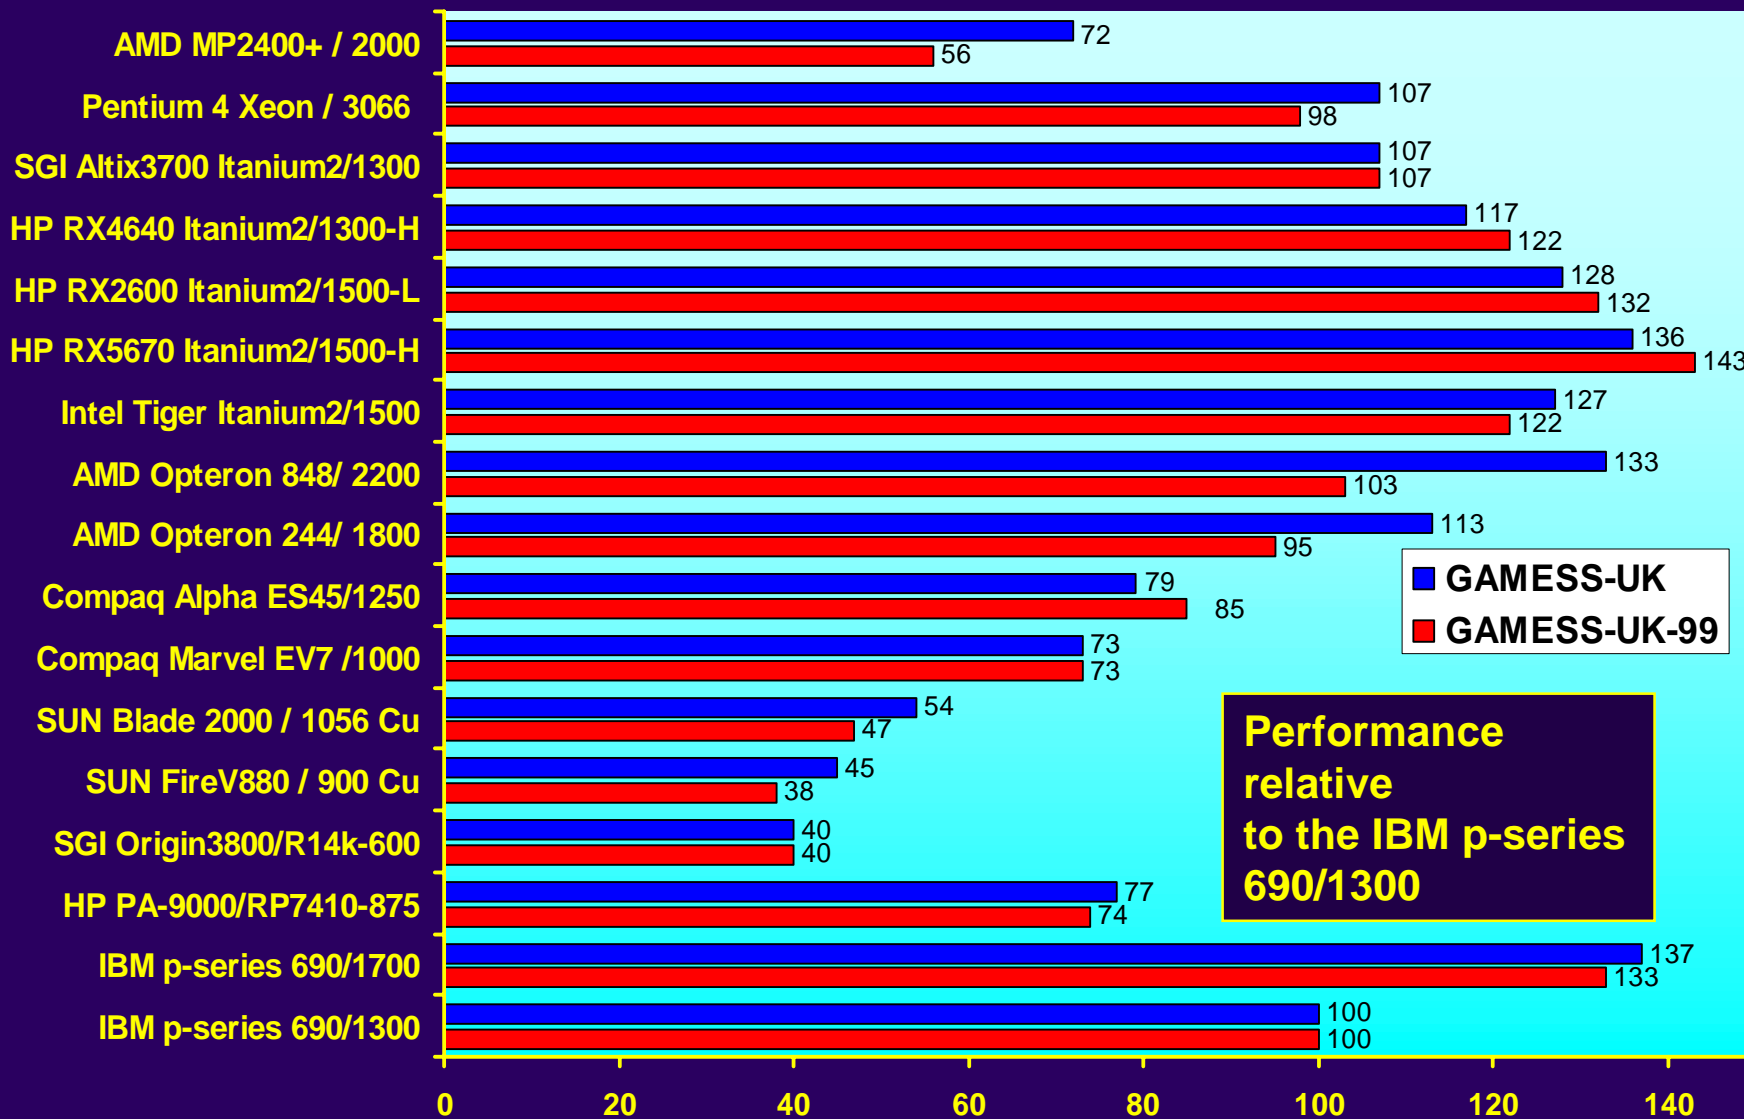
Benchmark Time: 30.8 minutes on IBM p-series 690/pwr4 1.3 GHz

# The GAMESS-UK-99 Benchmark

Performance relative to the IBM p-series 690/pwr4 1.3 GHz




# The GAMESS-UK and GAMESS-UK-99 Benchmark

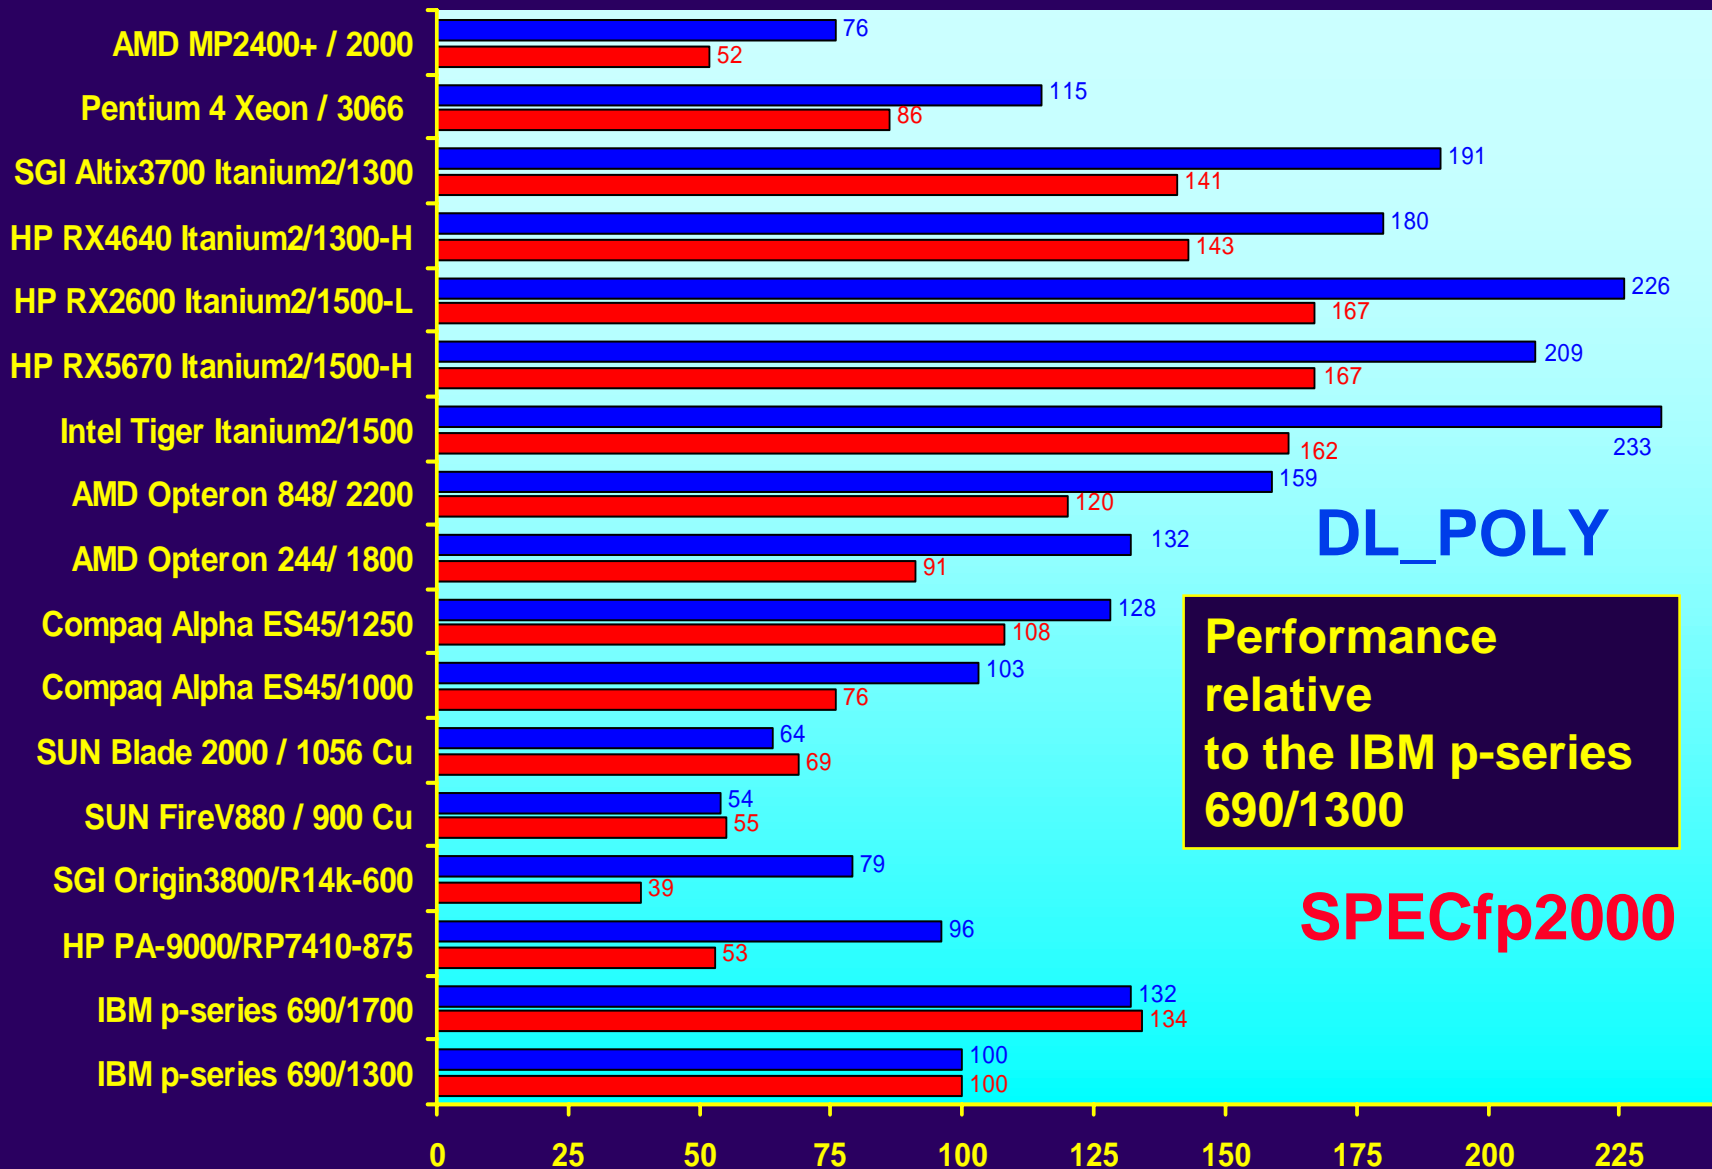


# The DL\_POLY Benchmark.

Performance relative to the IBM p-series 690/pwr4 1.3 GHz




# SPECfp2000 and the DL\_POLY Benchmark



# Summary Index relative to the IBM p-series 690/pwr4 1.3 GHz

## The MATRIX-97, Chemistry Kernels, GAMESS-UK and DLPOLY Benchmarks



# Computational Chemistry Rate Benchmark

- “Rate” Benchmark suite recently developed to incorporate:
  - Matrix “kernels” (MATRIX-97+)
  - Application packages (e.g. GAMESS-UK-99, DL\_POLY+)
- Multi-component benchmark:
  - Matrix Operations / Matrix multiplication and matrix diagonalisation
  - Quantum Chemistry Calculations
  - Molecular Dynamics Calculations
- Rate Procedure:
  - For each benchmark (i) run  $n$  instances at once and take elapsed time (last to finish - first to start).
  - The rate for this benchmark is  $R_i = n \times T_{ref} / T_i$
  - $T_{ref}$  is the elapsed time on a reference system. (IBM p-series 690/pwr4 1.3 GHz scaled to a single processor ( $n=1$ ) elapsed time of 100 units)
  - Take “the geometric mean” of all the benchmarks (with the same  $n$ ).

# Components of Chemistry Rate Benchmark

## 1. MATRIX OPERATIONS (MATRIX-97+)

- SPARSE Matrix Multiply BenchMark: Vector FORTRAN and DGEMM
  - series of MMOs ( $R = A \times B$ ) are performed involving matrices of increasing order:
    - MATRIX-97+: 100, 200, ... , 1200 (B is sparse)
- Diagonalisation Benchmark
  - performance of 7 routines from mathematical libraries and QC codes:
    - MATRIX-97+: 100, 200, ... , 600
- Q<sup>†</sup>HQ Benchmark
  - involves use of library routines e.g. BLAS. Uses both a scalar and vector algorithm (DGEMM):
    - MATRIX-97+: 100, 200, ... , 1000

Memory bandwidth

-lblas


## 2. DLPOLY

- 5 simulations (increased no. of time steps)


## 3. GAMESS-UK

- 8 Calculations (GAMESS\_UK-99 : Direct- SCF, DFT B3LYP, MCSCF, Direct-CI, MP2-geometry, SCF 2<sup>nd</sup> derivs., MP2 2<sup>nd</sup> derivs., Direct-MP2)

# Rate Benchmark: DGEMM Component




# Rate Benchmark: MMO FORTTRAN




# Rate Benchmark: Diagonalisation Component


## R-Diagonalisation



Decline in Rate with 32-bit commodity solutions




# Rate Benchmark: DL\_POLY Component



# Rate Benchmark: GAMESS-UK Component

## R-GAMESS-UK



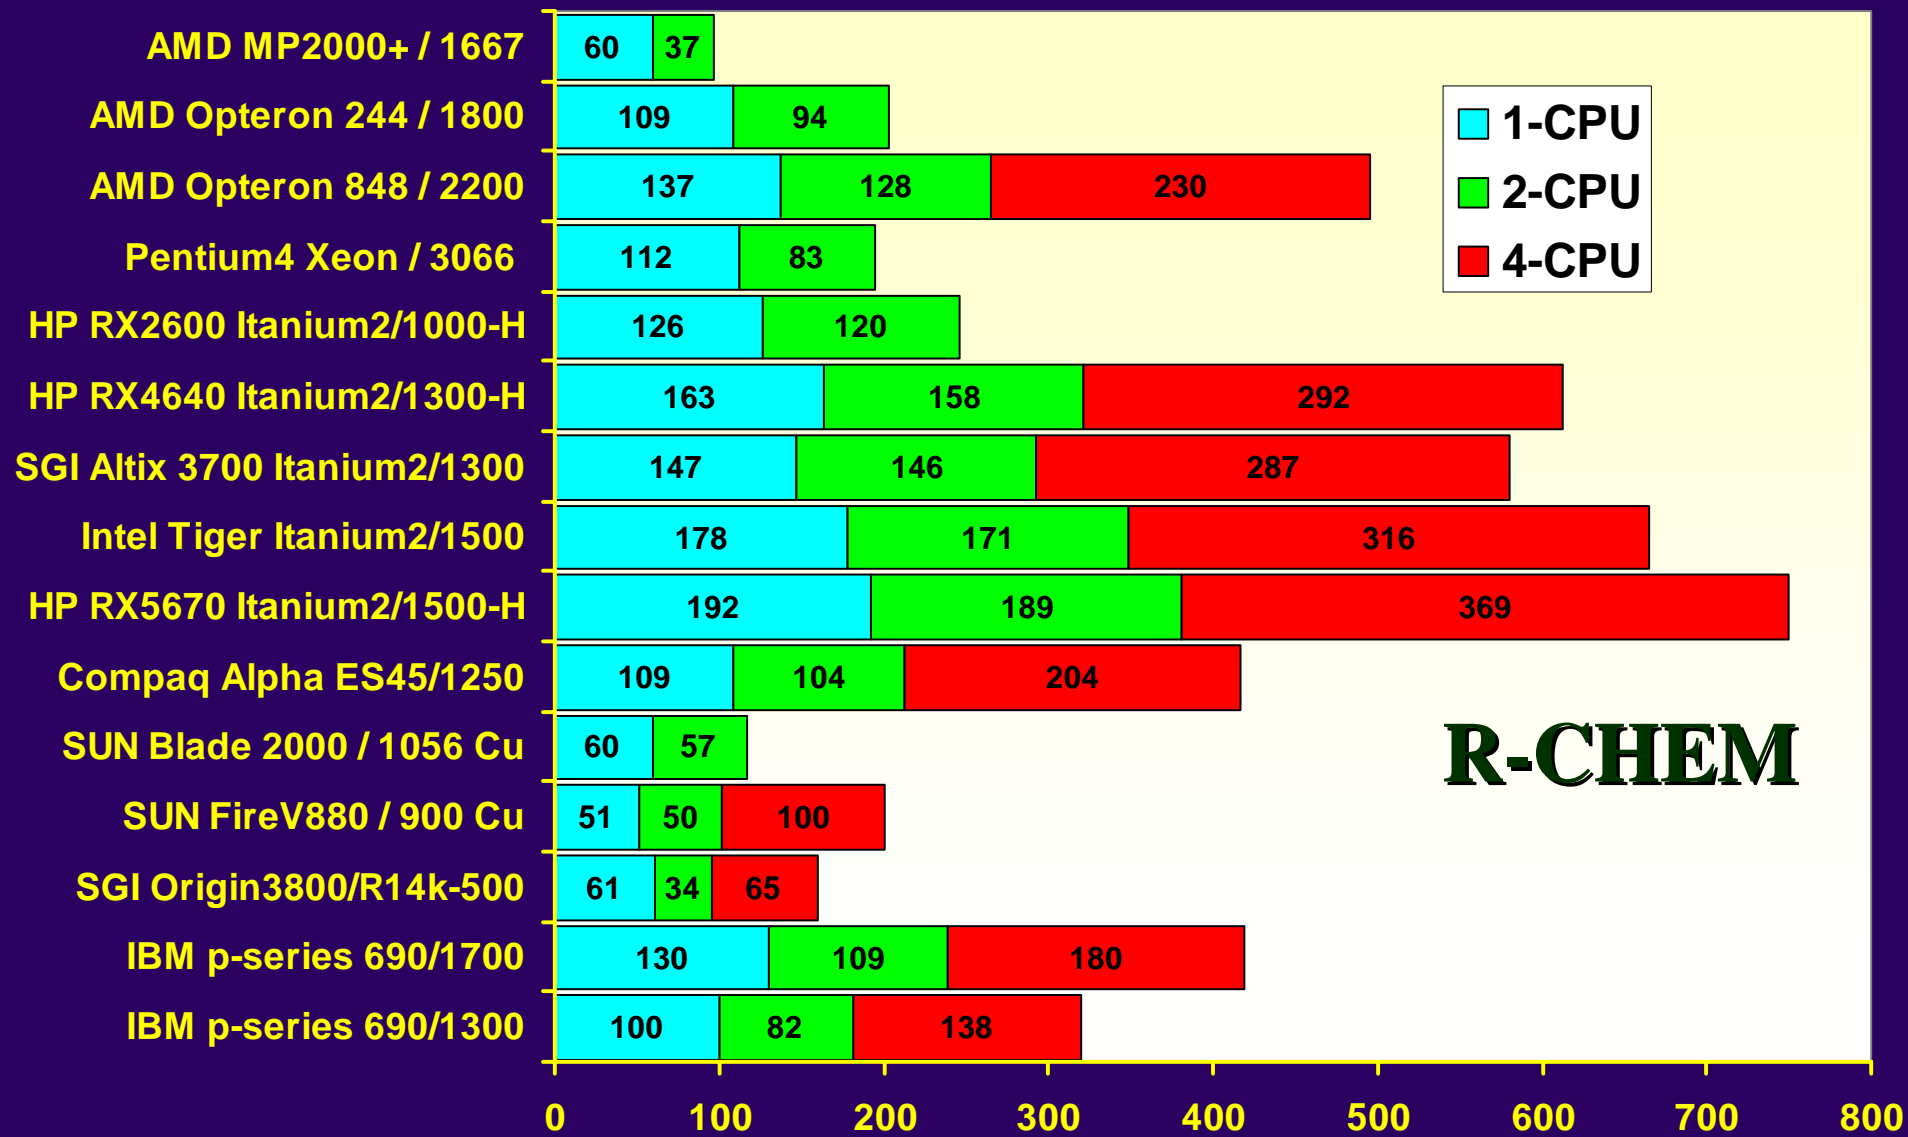
General decline in Rate with number of Runs

# Computational Chemistry Rate Metric

- Chemistry Rate Metric:
  - For each component benchmark (i) run  $n$  instances at once and take elapsed time (last to finish - first to start).
  - Component rate  $R_i = n \times T_{ref} / T_i$
  - $T_{ref}$  is the elapsed time on the IBM p-series 690/pwr4 1.3 GHz scaled to a single CPU time of 100 units.

$$R_{CHEM} = 0.1 \times R_{DGEMM} + \\ 0.1 \times R_{Diagonalisation} + \\ 0.4 \times R_{DLPOLY} + \\ 0.4 \times R_{GAMESS-UK}$$

# Chemistry Rate Benchmark



**R-CHEM**

# Evaluation and Benchmarking of High-end and Commodity-based Systems

## Parallel Systems and Capacity Computing

# High-End Systems Evaluated

- Cray T3E/1200E
  - 816 processor system at Manchester (CSAR), 600 Mz Alpha EV56 CPU, 256 MB
- IBM SP/WH2-375 and SP/p690
  - 32 CPU system at DL, 4-way WH2 SMP “thin nodes” with 2 GB memory, 375 MHz Power3-II processors with 8 MB L2 cache
  - **IBM Regatta-H** (32-way node, 1.3 GHZ power4 CPUs) at Montepelier
  - **IBM SP/p690 (8-way LPAR'd nodes, 1.3 GHZ) at Daresbury (1280 CPUs, HPCx)**
- Compaq AlphaServer SC
  - 4-way ES40/667 A21264A (APAC) and 833 MHz SMP nodes (2 GB RAM);
  - **TCS1 system at PSC** (comprising 750 4-way ES45 nodes - 3,000 EV68 CPUs - with 4 GB memory per node, 8MB L2 cache)
  - Quadrics “fat tree” interconnect (5 usec latency, 250 MB/sec B/W)
- SGI Origin 3800
  - SARA (1000 CPUs) - Numalink with MIPS R14k/500 CPUs
- SGI Altix 3700
  - **Linux Cluster - Numalink with Itanium 2 1.3 GHz CPUs, 3MB L3 cache**
  - **CSAR (“newton” 256 CPUs) and SARA (“aster” - 416 CPUs - 7 nodes)**

# Commodity Systems (CSx)

## Prototype / Evaluation Hardware

<u>Systems</u>	<u>Location</u>	<u>CPUs</u>	<u>Configuration</u>
CS1	Daresbury	32	PentiumIII / 450 MHz + FE (EPSRC)
CS2	Daresbury	64	24 X dual UP2000/EV67-667, QSNNet Alpha/LINUX cluster, 8 X dual CS20/EV67-833 ("loki")
CS3	RAL	16	Athlon K7 850MHz + myrinet
CS4	Sara	32	Athlon K7 1.2 GHz + FE
CS6	CLiC	528	PentiumIII / 800 MHz; fast ethernet (Chemnitzer Cluster)
CS7	Daresbury	64	AMD K7/1000 MP + SCALI/SCI ("ukcp")
CS8	NCSA	320	160 dual IBM Itanium/800 + Myrinet 2k ("titan")
CS9	Bristol	96	Pentium4 Xeon/2000 + Myrinet 2k ("dirac")
<a href="http://www.cse.clrc.ac.uk/Activity/DisCo">www.cse.clrc.ac.uk/Activity/DisCo</a>			
<u>Protoype Systems</u>			
CS0	Daresbury	10	10 CPUS, Pentium II/266
CS5	Daresbury	16	8 X dual Pentium III/933, SCALI

# Commodity Systems (CSx) II.

## Evaluation Hardware

<u>Systems</u>	<u>Location</u>	<u>CPUs</u>	<u>Configuration</u>
CS10	<i>Hull</i>	64	Pentium4 Xeon/2666 + Myrinet 2k ("eagle"), Streamline/SCORE
CS11	<i>Workstations UK</i>	32	Pentium4 Xeon/2400 + GbitEther, ScaMPI
CS12	<i>Essex</i>	48	Pentium4 Xeon/2400 + GbitEther ("sstream1"), Streamline/SCORE
CS13	<i>White Rose, Leeds</i>	256	Pentium4 Xeon/2200-2400 + Myrinet 2k, ("snowdon"), Streamline/SCORE
CS14	<i>NCSA</i>	1024	Pentium III Xeon/1000 + Myrinet 2k ("platinum")
CS15	<i>SDSC</i>	128	Pentium III Xeon/ 8000 + Myrinet 2k ("meteor")
CS16	<i>SDSC</i>	256	dual-Itanium2/1.3 GHz nodes + Myrinet 2k ("Teragrid")

# Applications Performance Overview

- **Serial (SPEC, DL) & Communication Benchmarks**
- **Parallel Applications Performance**

1. Computational Chemistry:  
**Molecular Simulation & Electronic Structure**
2. Computational Materials Science
3. Atomic & Molecular Physics
4. Computational Engineering
5. Environmental Modelling

- Capacity-based group solution
- Issues of Cost effectiveness
- On e.g. 128-256 CPU cluster, modal job size is ~ 32 CPUs
- Increasing trend to hierarchical clusters - Gbit network with HEC core (with e.g. myrinet)

**Capability and Capacity Computing  
Commodity vs. Proprietary Solutions**

# Performance Metrics: 1999-2001

Attempted to quantify delivered performance from the Commodity-based systems against current MPP (CSAR Cray T3E/1200E) and ASCI-style SMP-node platforms (e.g. SGI Origin 3800) i.e.

## Performance Metric (% 32-node Cray T3E)

$T$  (32-nodes Cray T3E/1200E) /  $T$  (32 CPUs) CSx

[  $T_{32\text{-node T3E}} / T_{32\text{-node CS1 Pentium III/450 + FE}$  ]

$T_{32\text{-node T3E}} / T_{32\text{-node CS6 Pentium III/800 + FE}$

$T_{32\text{-node T3E}} / T_{32\text{-CPU CS2 Alpha Linux Cluster + Quadrix}$

# Performance Metrics: 2002

## Performance Metric (% 32-node AlphaServer SC [PSC])

$T$  (32-CPU AlphaServer SC ES45/1000) /  $T$  (32 CPUs) CSx

$T_{32\text{-CPU AlphaServer ES45}} / T_{32\text{-CPU CS9 Pentium 4 Xeon / 2000 + Myrinet 2k}}$

# Commodity Comparisons with High-end Systems: Compaq AlphaServer SC ES45/1000 and the SGI O3800/R14k-500

CS9 - Pentium4/2000 Xeon + Myrinet 2k

% of 32-CPU AlphaServer SC Origin 3800

2002

## GAMESS-UK

SCF	95%	135%
DFT	70-73%	117-161%
DFT (Jfit)	59-70%	92-111%
DFT Gradient	69%	114%
MP2 Gradient	59%	78%
SCF Forces	92%	148%

**NWChem** (DFT Jfit) 65-88% 94-227%

**GAMESS / CHARMM** 89% 139%

## DL\_POLY

Ewald-based	44-53%	71-84%
bond constraints	41%	64%

**CHARMM** 95% 103%

**CASTEP** 47-67% 75-81%

**CPMD** 30% 53%


**ANGUS** 35-69% 51-147%

# DL\_POLY V2: High-end and Commodity-based Systems

**Bench 4. NaCl; 27,000 ions, Ewald, 75 time steps, Cutoff=24Å**

**Bench 5: NaK-disilicate glass; 8,640 atoms, MTS + Ewald: 270 time steps**

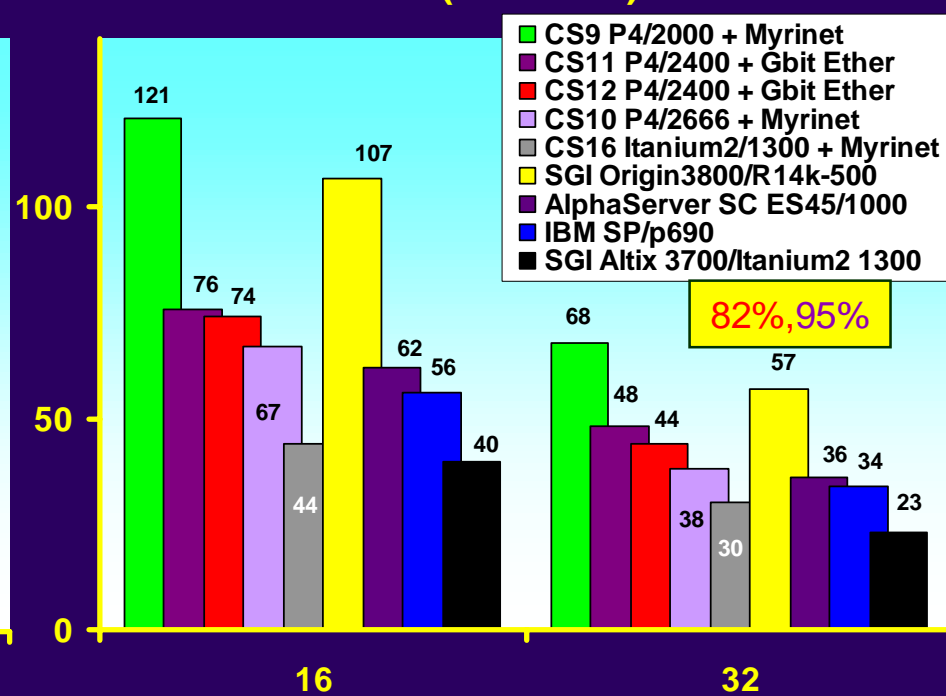
Measured Time (seconds)



Number of CPUs

$T3E_{128} = 94$

Measured Time (seconds)



Number of CPUs

$T3E_{128} = 75$

# Commodity Comparisons with High-end Systems: The Compaq AlphaServer SC ES45/1000

Application	% of 32 CPUs of Compaq AlphaServer SC ES45/1000	
	Cluster CS10 Pentium4/2666 Xeon + Myrinet 2k	CS12 Pentium4/2400 Xeon + Gbit Ethernet
<b>GAMESS-UK</b>		
SCF	91%	85%
DFT	85-103%	78 - 95%
DFT (Jfit)	75%	70%
DFT Gradient	85%	85%
MP2 Gradient		59% †
SCF Forces		92% †
<b>DL_POLY- 2 (and - 3)</b>		
Ionic Systems	77- 95% ( 91%)	66-82% ( 84%)
Macromolecular	75% ( 92%)	61% ( 86%)
<b>CHARMM</b>	139%	104%
<b>CASTEP</b>	60-83%	53- 66%
<b>CPMD</b>	60-65%	55%
<b>ANGUS</b>	36-73%	34-72%
<b>SBLI</b>	53%	51%

# SUMMARY

- Processor Performance Overview
  - Single-processor performance, Performance Metrics
- SPECfp and Computational Chemistry Benchmark (serial)
  - Comparison involves 200+ computers
  - Matrix'89 and Matrix'97 kernels (MMO, diagonalisation), Application kernels" (SCF, MD, QMC and JACOBI + STREAM), Application packages (GAMESS-UK, DL\_POLY)
- Increasing importance of Rate-based benchmarks:
  - SPECfp\_rate and Chemistry Rate Benchmark
- Machine COSTS vs. Performance: URLs:
  - Powerpoint presentation and Paper:  
*<http://www.cse.clrc.ac.uk/disco/hw-perf.shtml>*
- Benchmarking & Evaluation of Commodity-based Systems
  - CS10 Xeon / 2666 cluster (Myrinet-2k) ca. 82% of Alpha ES45/1000.
  - Cost effectiveness of the CS12 Xeon / 2400 Cluster with Gbit Ethernet (Streamline, 75% of the Alpha ES45/1000).