PERFORMANCE of VARIOUS COMPUTERS in COMPUTATIONAL CHEMISTRY

13th 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),
H. Richardson (Sun),
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

MEW12 (2001) - Summary PI relative to the Compaq Alpha ES40/833

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

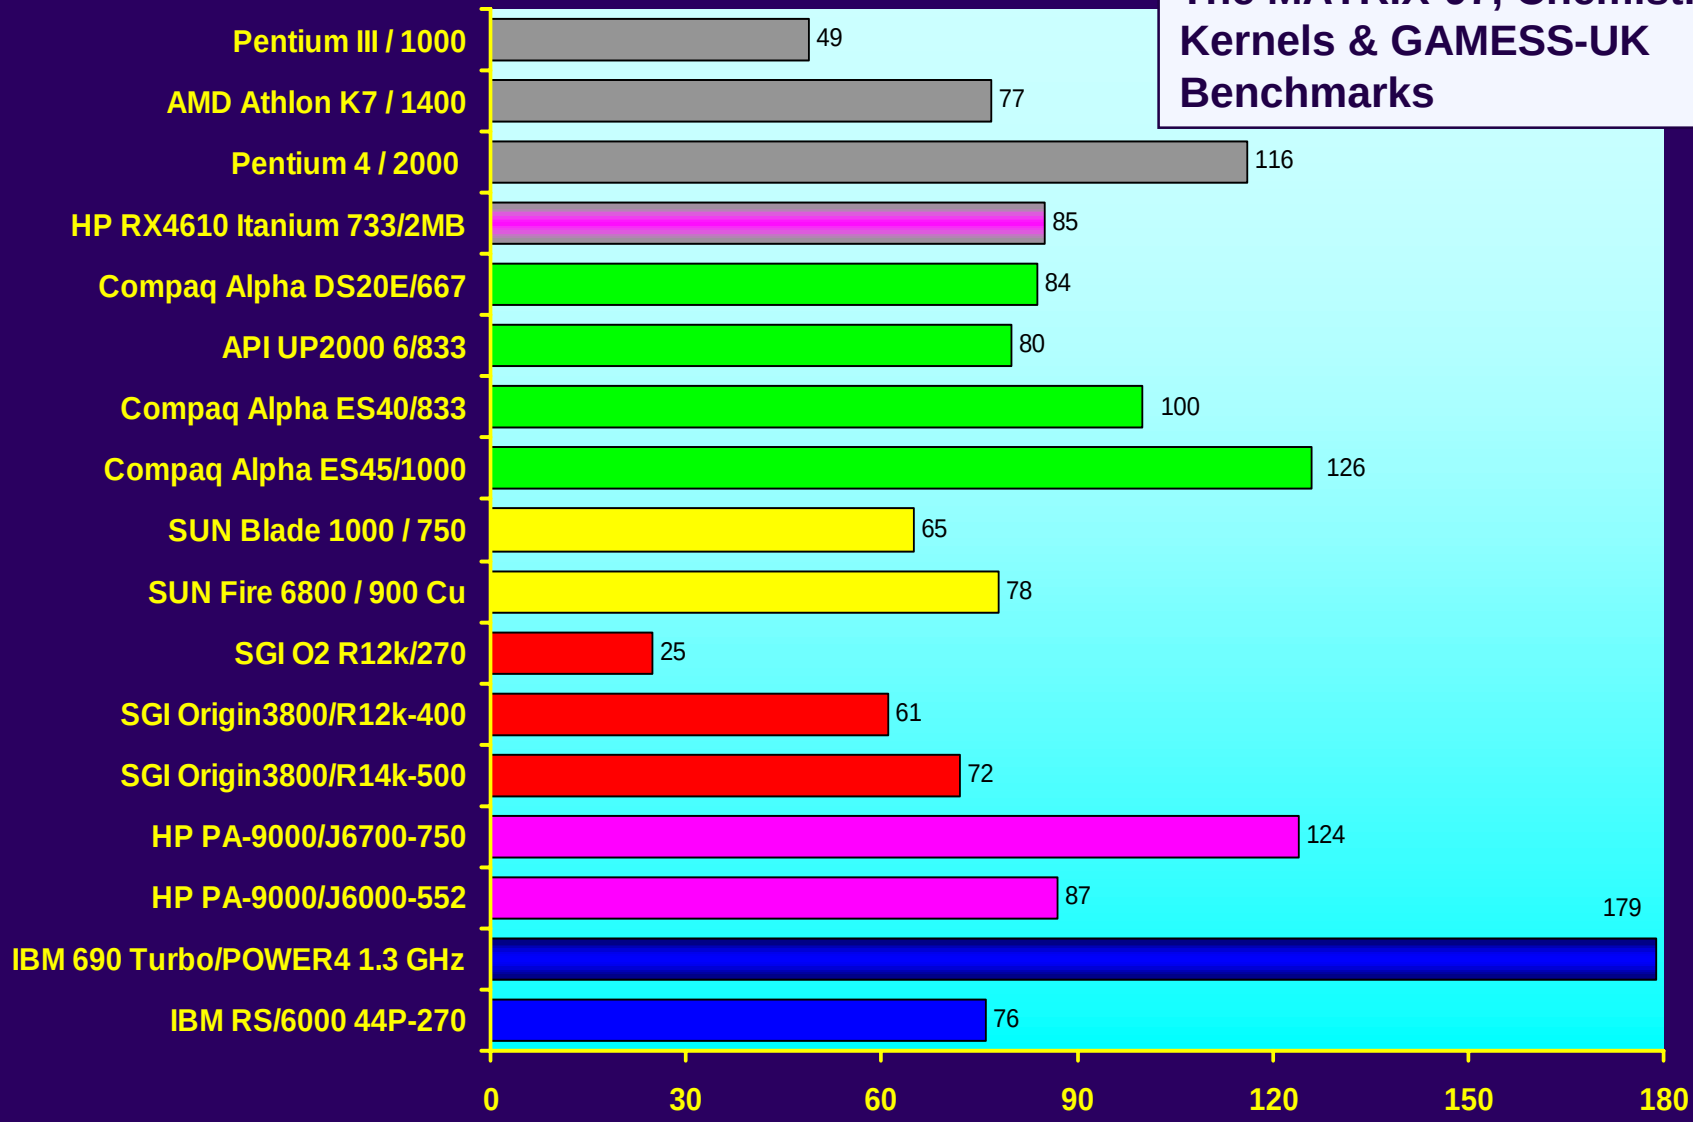

Fortran Compilers - Performance

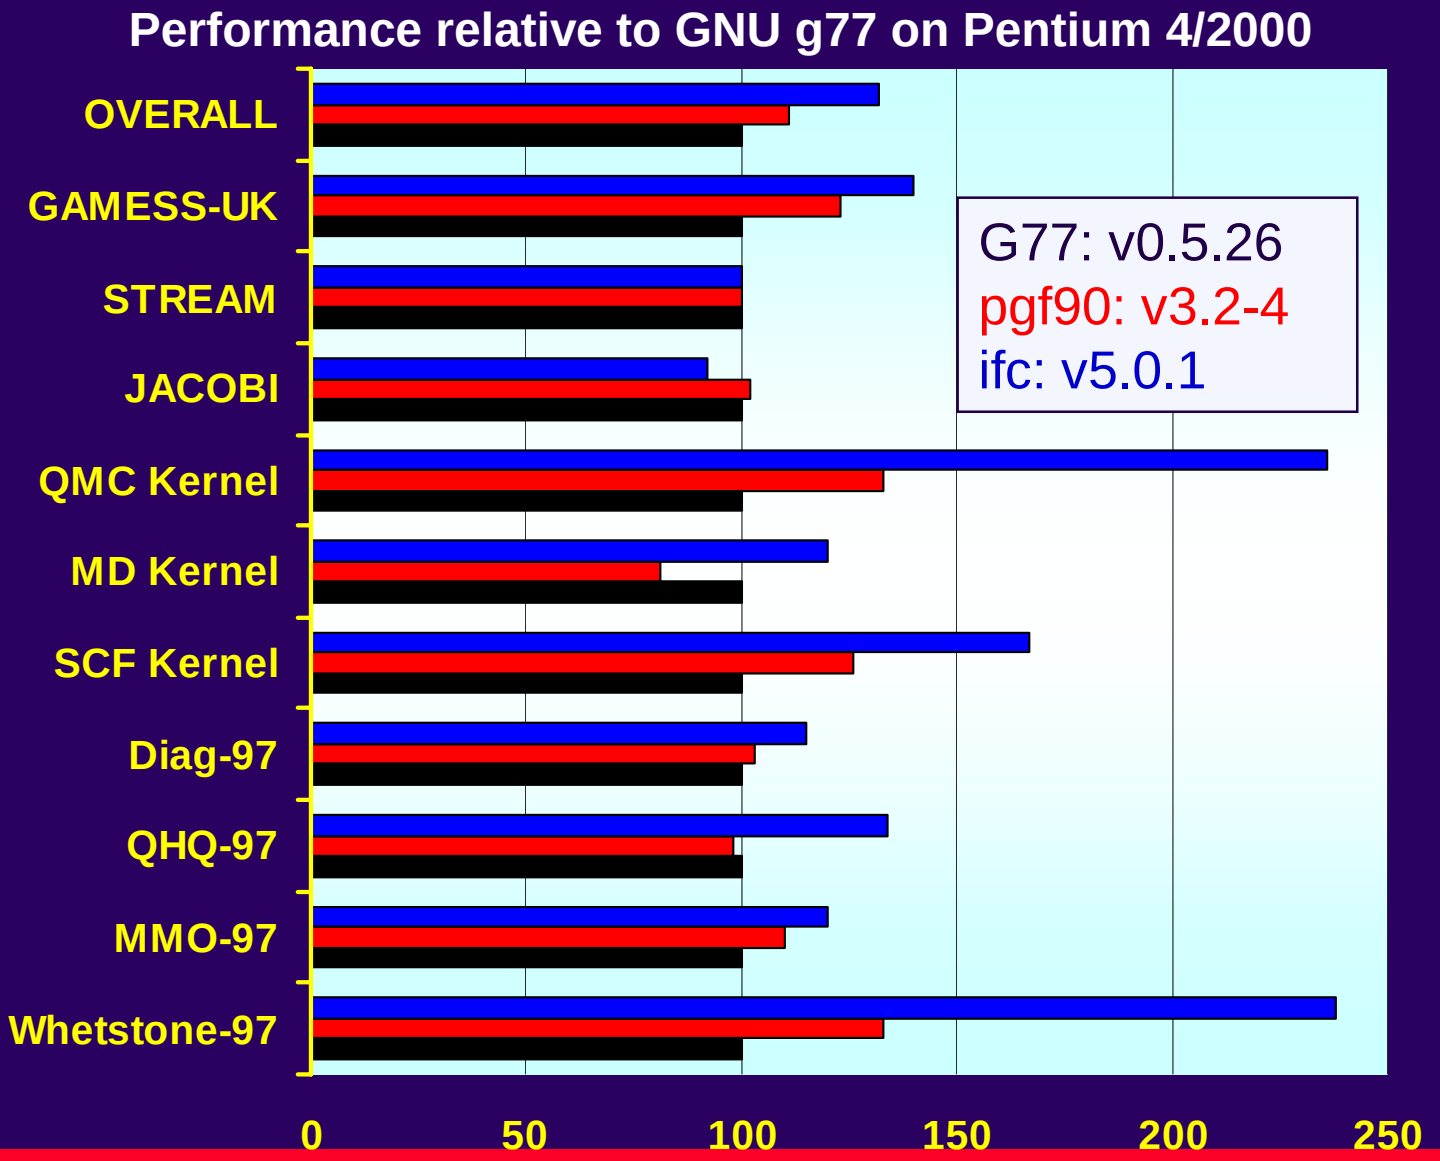

MACHINES UNDER EVALUATION

Machine	Processor
AMD Athlon	*MP* 1.533 , *1.667 GHz, 2 GHz*
Pentium - PC	Pentium 4- 1.5, 2.0 GHz
SUN Blade 2000 / 1056 Cu	*UltraSPARC-3 / 1056 MHz*
SUN Fire V880 / 900 Cu	*UltraSPARC-3 / 900 MHz*
HP PA-9000 / J6700	PA8700 / 750 MHz
HP PA-9000 / RP7410	*PA8700+ / 875MHz*
HP RX4610	Itanium / 733 MHz (2 MB L3)
HP RX2600 (2-way)	*Itanium 2 / 1000 MHz (3 MB L3)*
HP RX5670 (4-way)	*Itanium 2 / 1000 MHz (3 MB L3)*
HP ZX6000	*Itanium 2 / 900 MHz (1.5 MB L3)*
Intel Tiger (4-way)	*Itanium 2 / 1000 MHz (3 MB L3)*
Compaq AlphaServer ES45	AXP A21264C / 1000 MHz
Compaq AlphaServer ES45	*AXP A21264C / 1250 MHz*
Compaq Alpha Marvel	*EV7 1000 MHz*

Machine	Processor
SGI Origin3800/R14k	*R14000/R14010 600 MHz*
SGI Origin3800/R14k	R14000/R14010 500 MHz
SGI Origin300/R14k	R14000/R14010 500 MHz
SGI O2/R12k-SC	R12000/R12010 270 MHz
IBM p-Series 630	*RS/6000 /POWER4 1.0GHz*
IBM p-Series 690	*RS/6000 /POWER4 1.3GHz*
Intellisation Z Pro	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
	RS/6000 NH - 222 MHz

SPEC Benchmarks

NOT:
- Graphics
- Network
- I/O

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

http://www.specbench.org

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

- SPECratio - T(measured system) / T (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.equake	C	Seismic wave propagation simulation
187.facerec	F90	Image processing: Face recognition
188.ammp	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 HP RX2600 Itanium2/1000

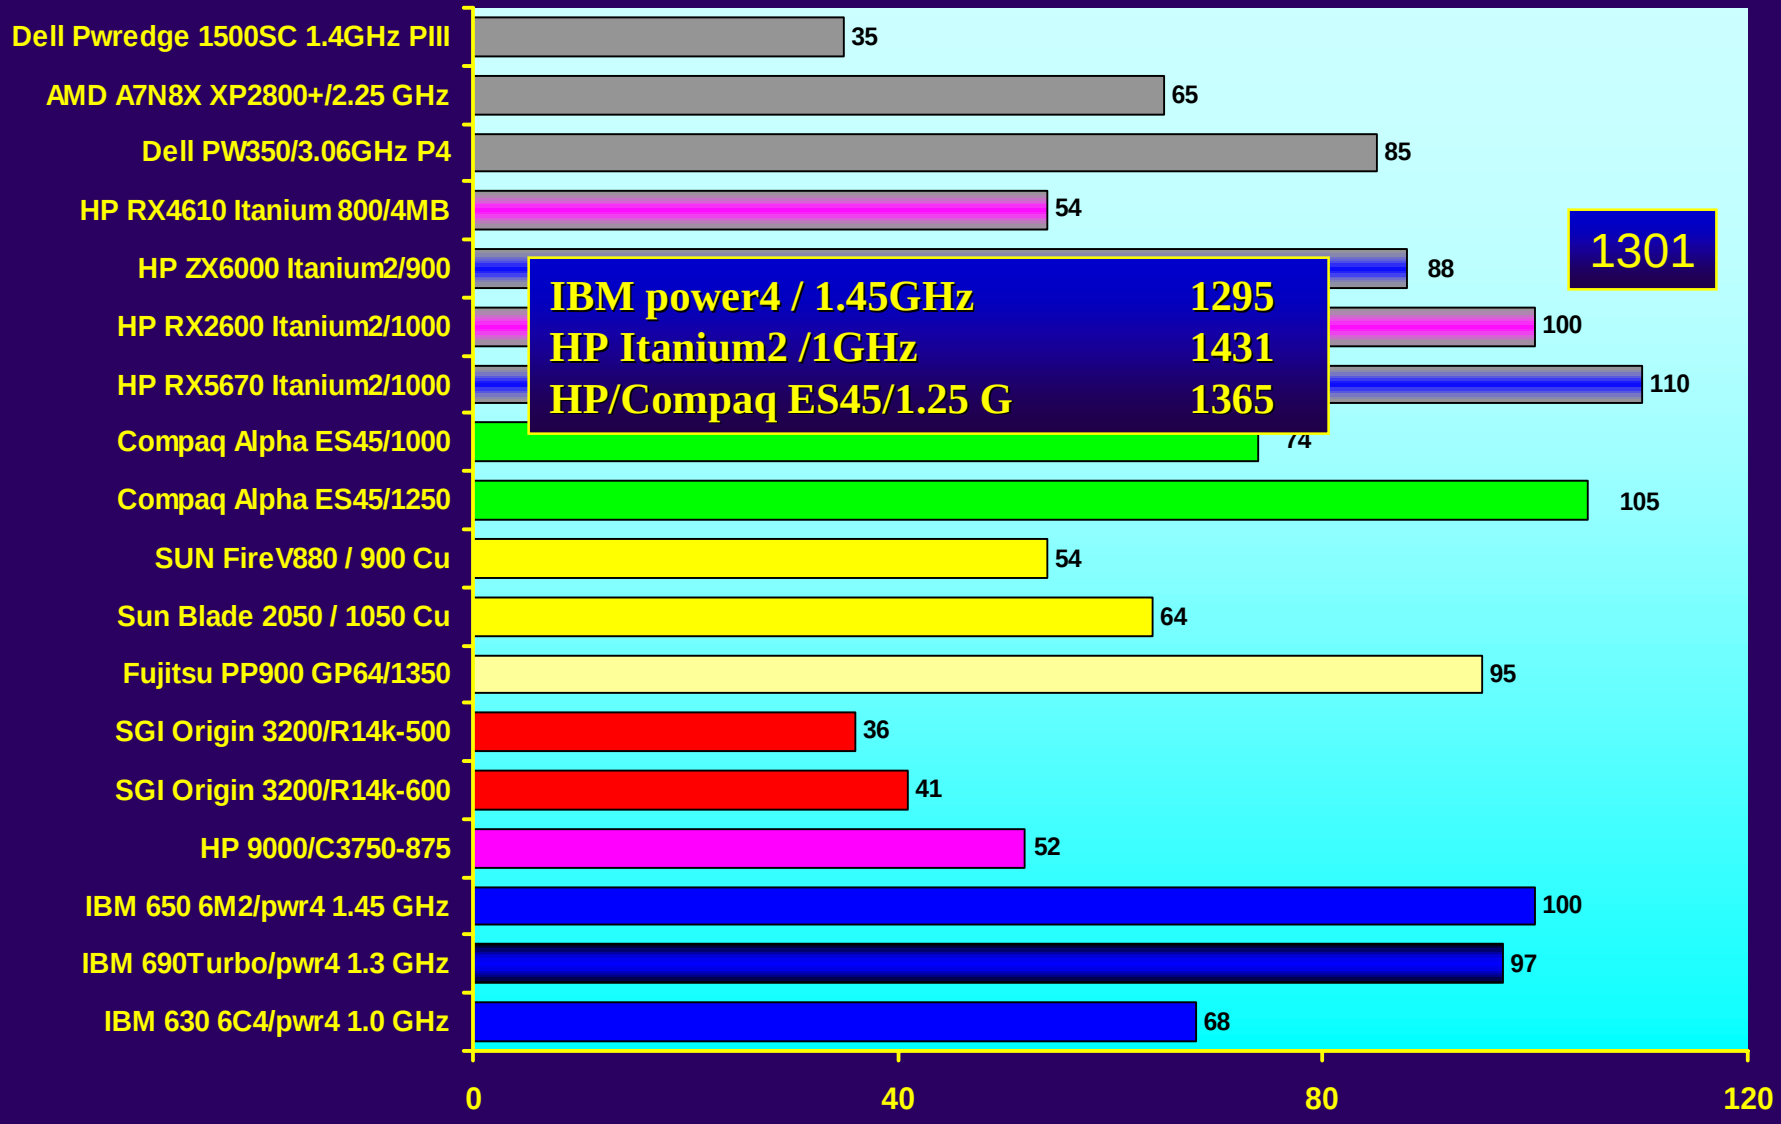

SPEC CPU 2000 - SPECfp2000_base

Values relative to HP RX2600 Itanium2/1000

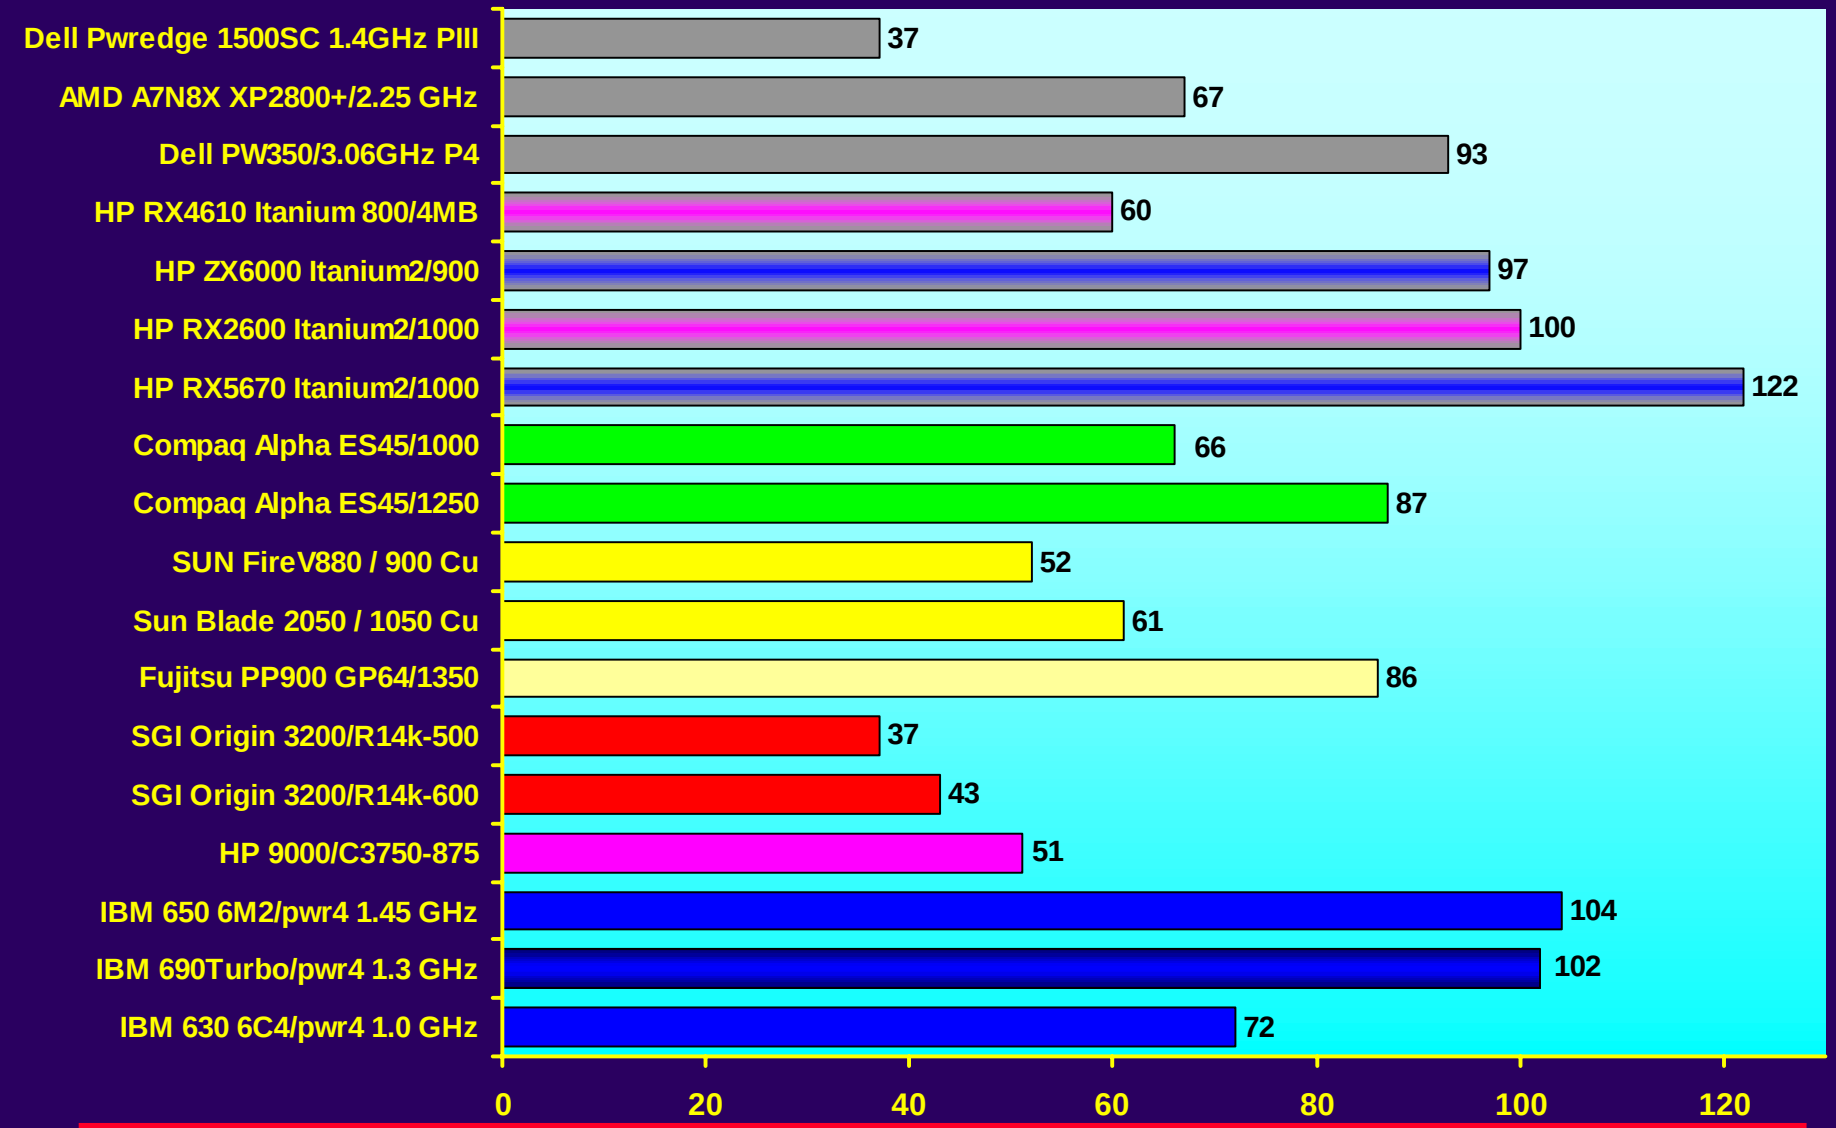

SPEC CPU 2000 - SPECint2000

Values relative to HP RX2600 Itanium2/1000

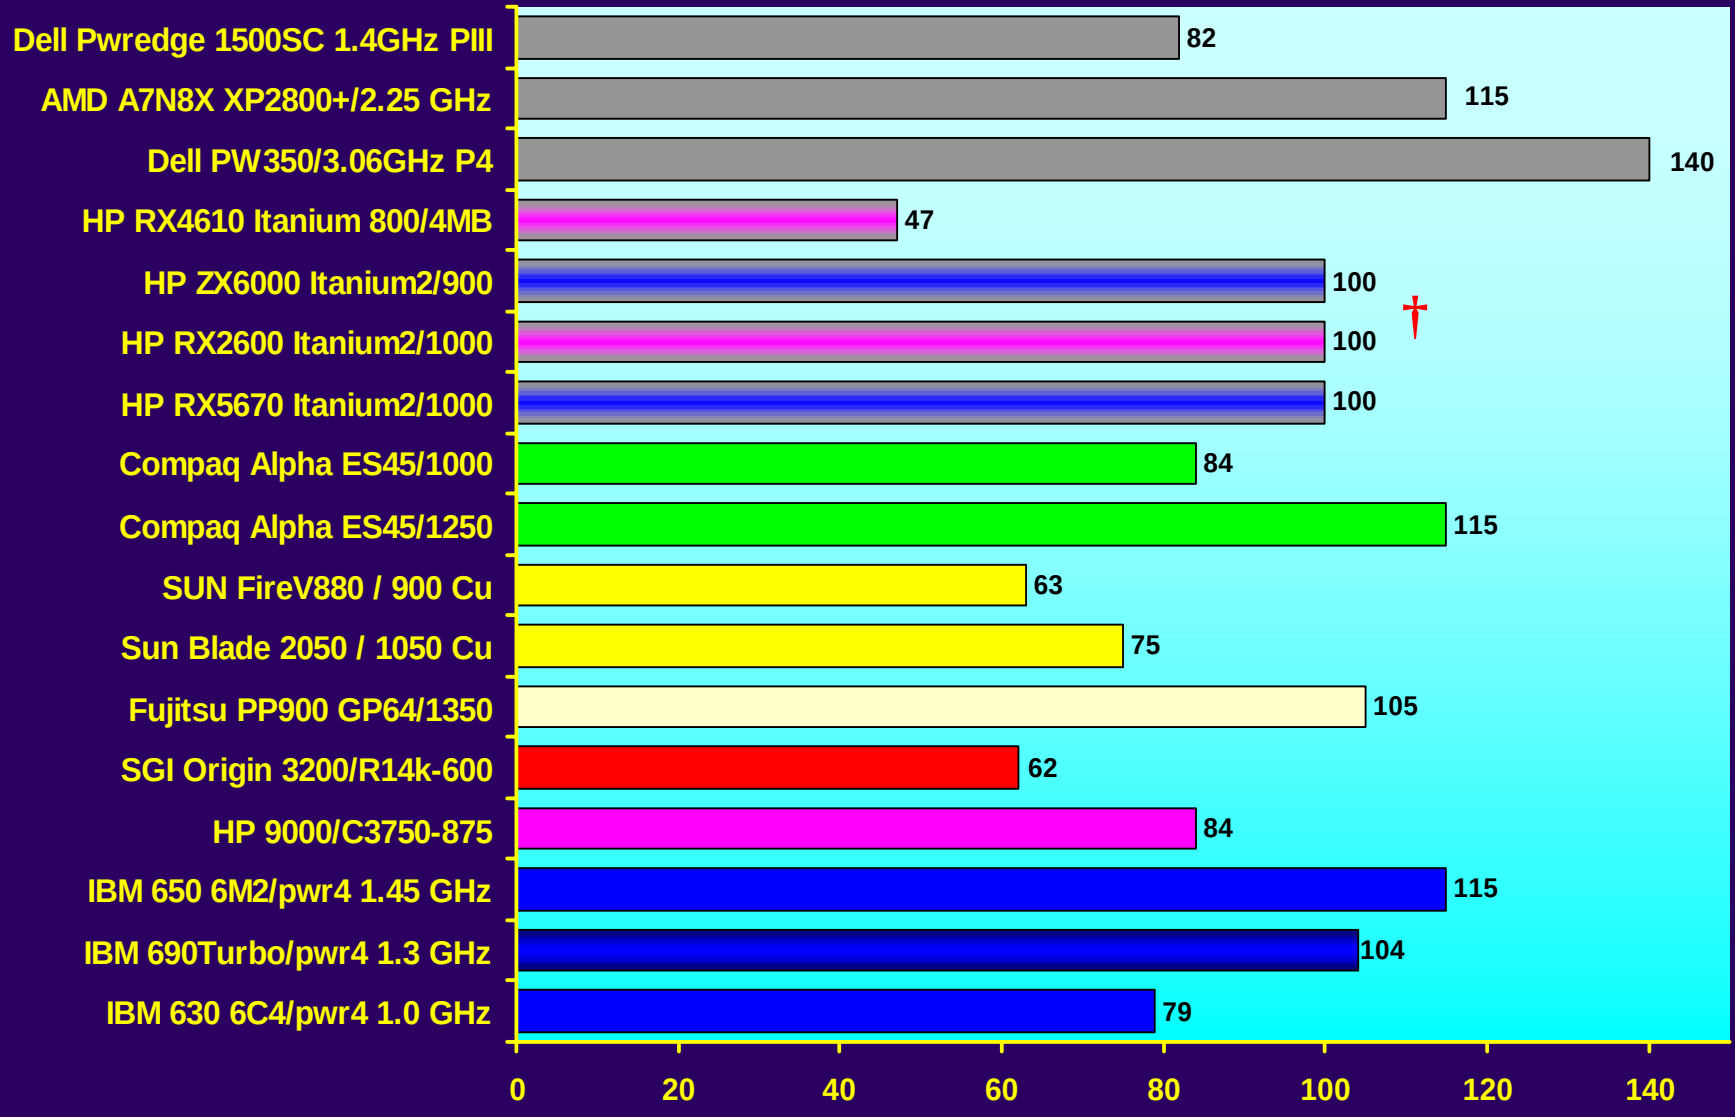

IBM p-series 690Turbo:Multi-chip Module (MCM)

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

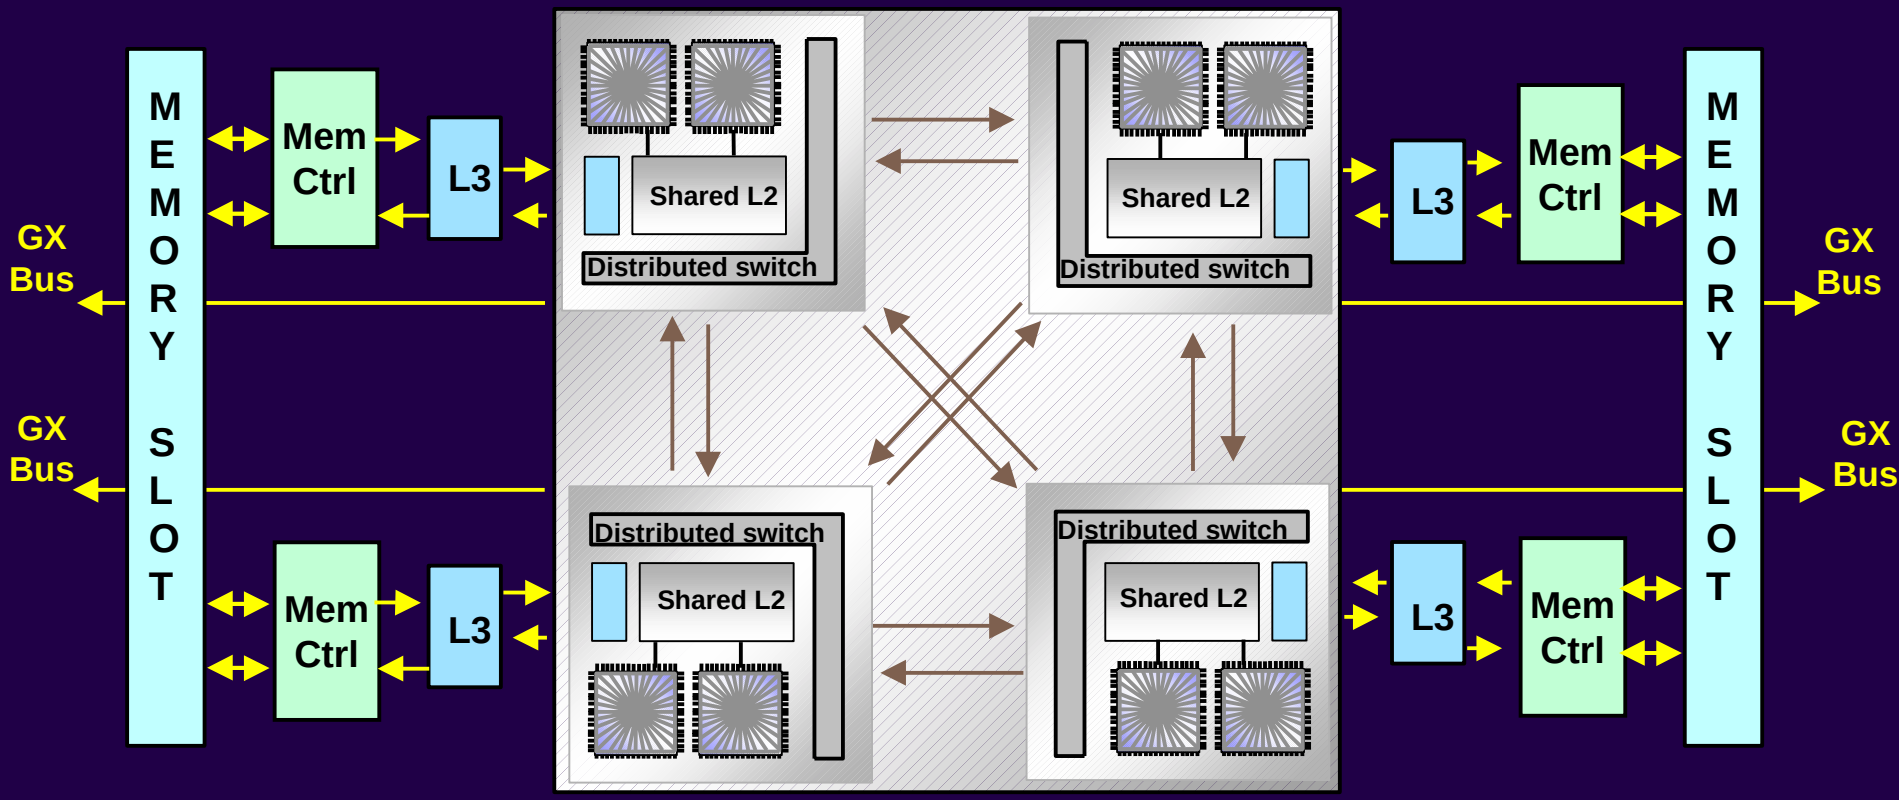

L3 cache shared across all processors

4 GX Bus links for external connections

SPEC CPU 2000 - SPECfp2000

Units of SPECfp as a function of Peak Performance

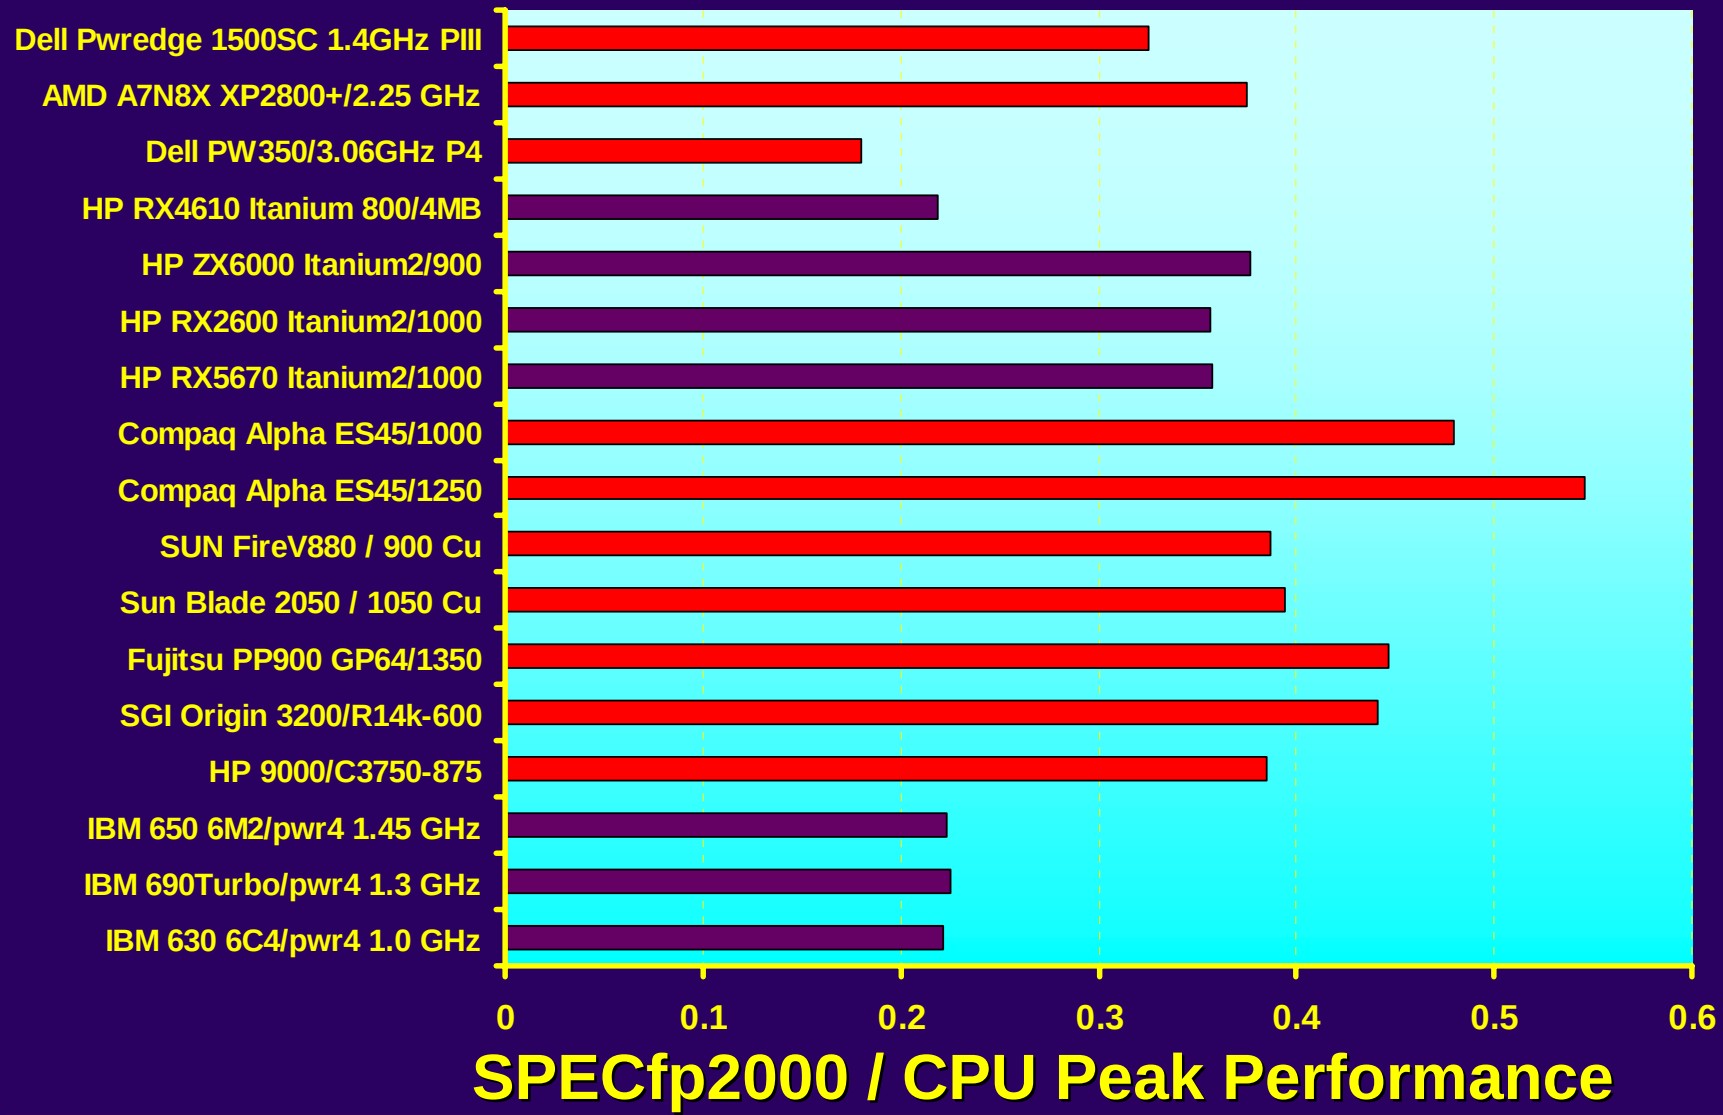

SPEC CPU 2000 - SPECfp vs SPECfp_rate (32 CPUs)

Values relative to IBM 690 Turbo 1.3 GHz

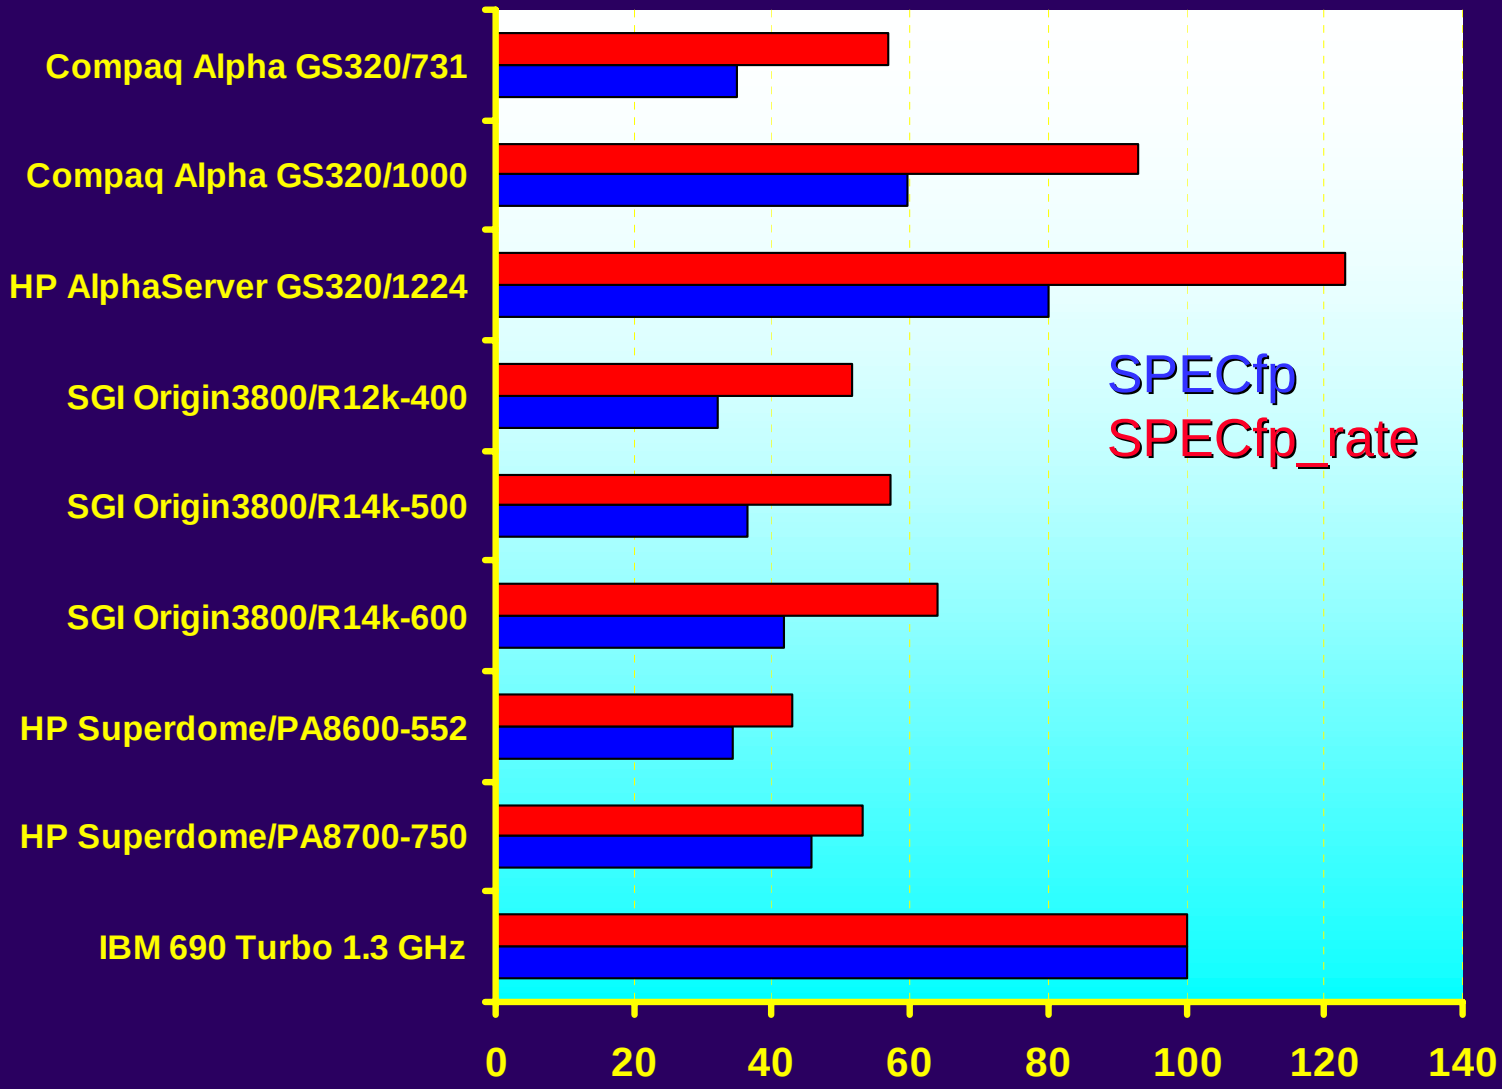

Computational Chemistry Benchmark Suite

- 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 X 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)

Memory bandwidth

- 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
- $Q^{\dagger}HQ$ 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

-lblas

Matrix-97: SPARSE MMO Benchmark.

Performance relative to the HP RX2600 Itanium2/1000

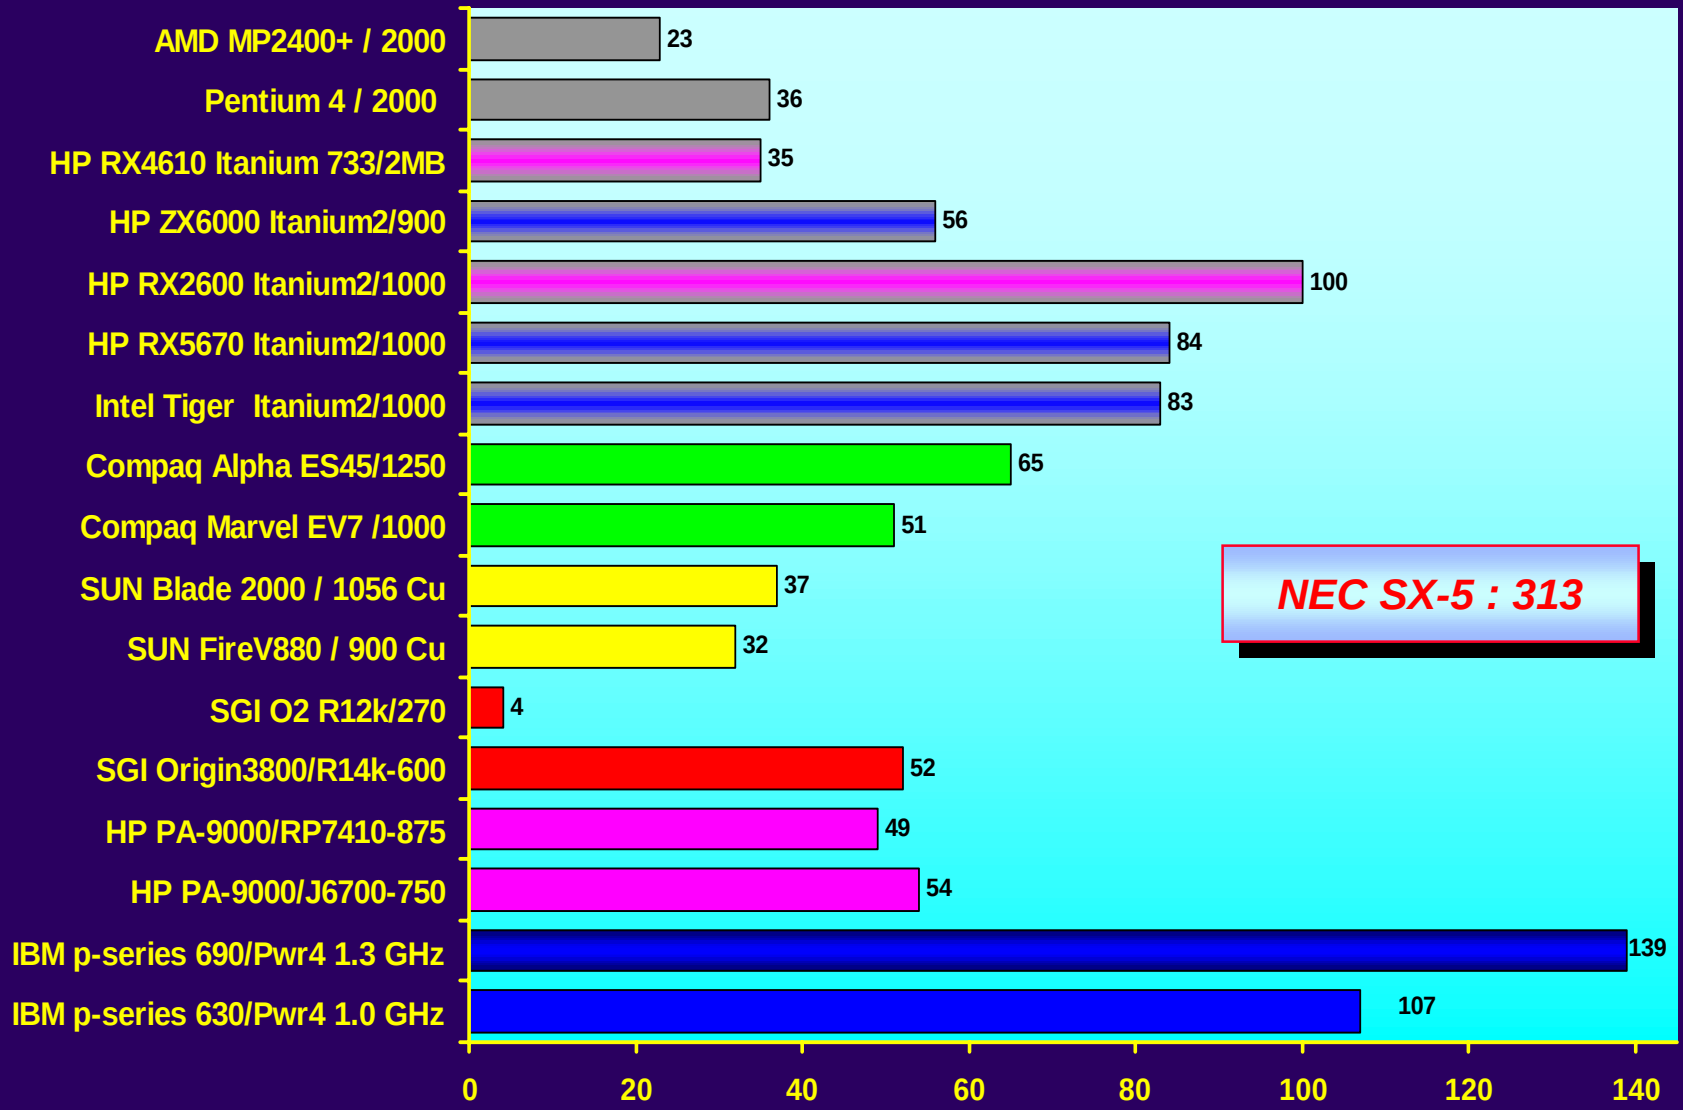

Matrix-97: Q†HQ MMO Benchmark.

Performance relative to the HP RX2600 Itanium2/1000

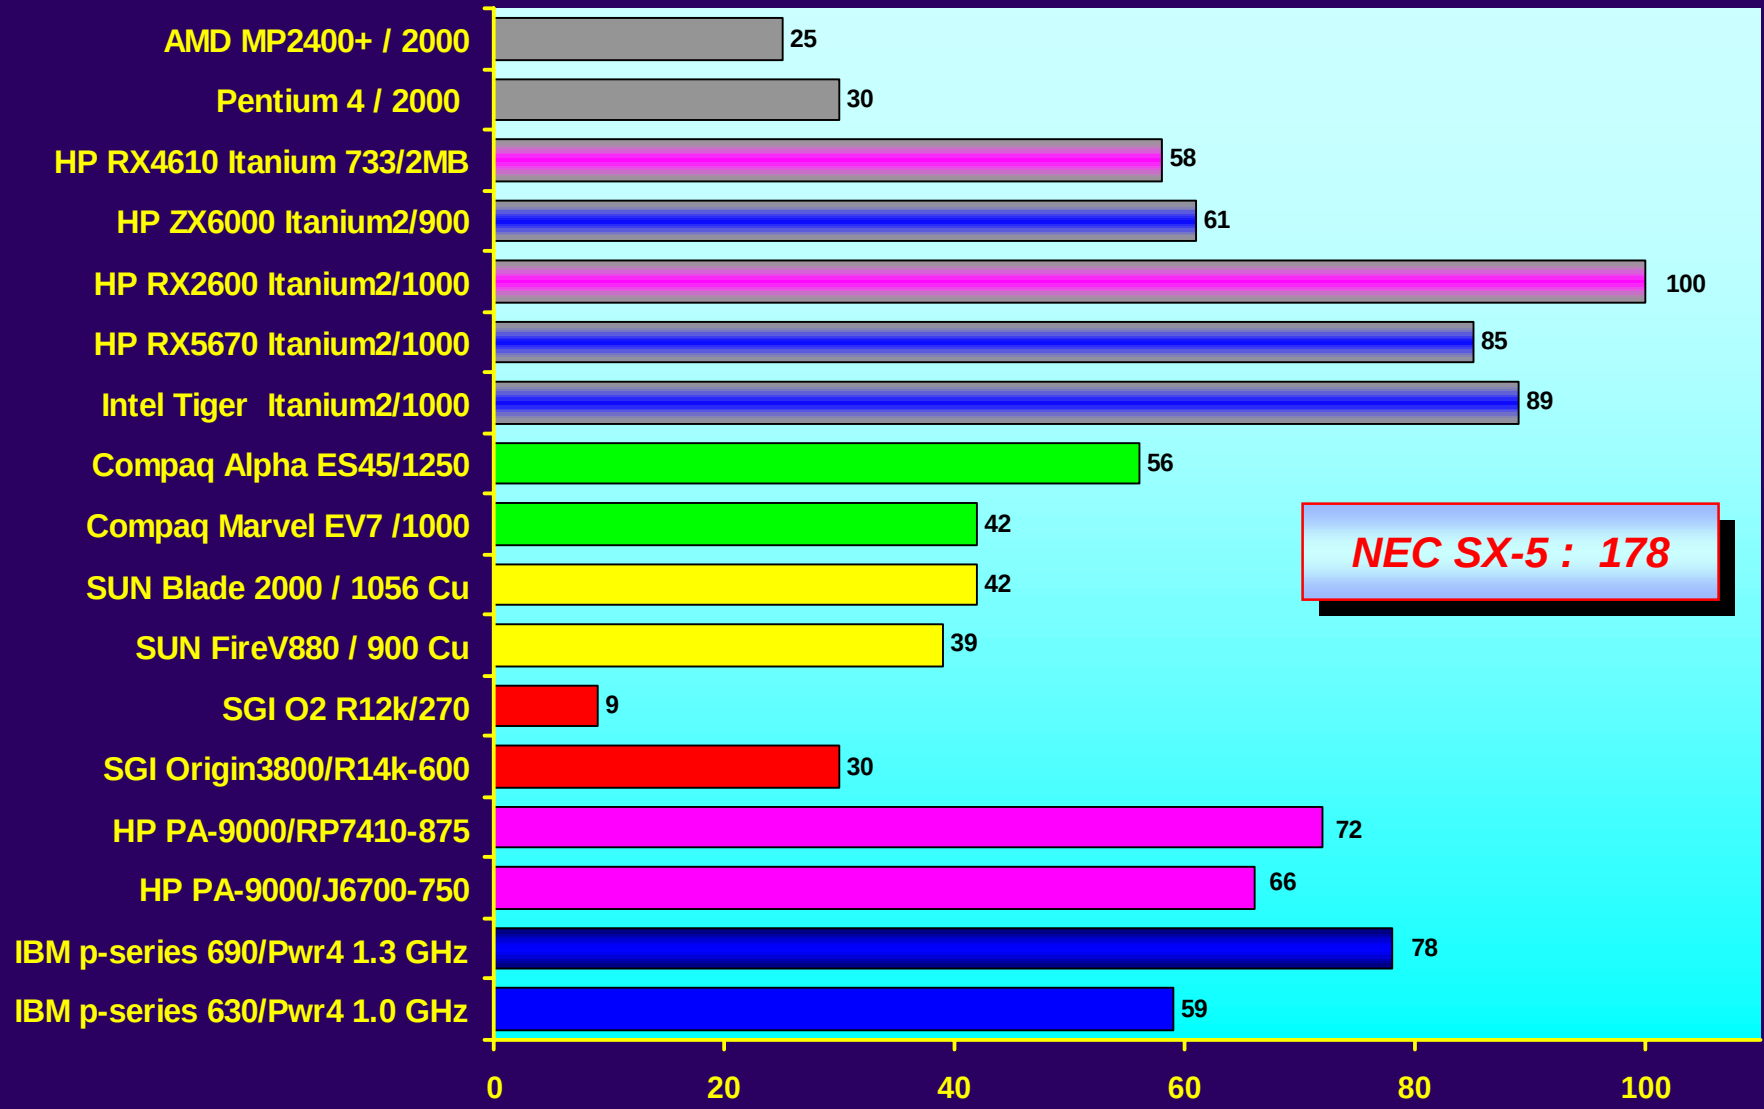

Matrix-97: Diagonalisation Benchmark

Performance relative to the HP RX2600 Itanium2/1000

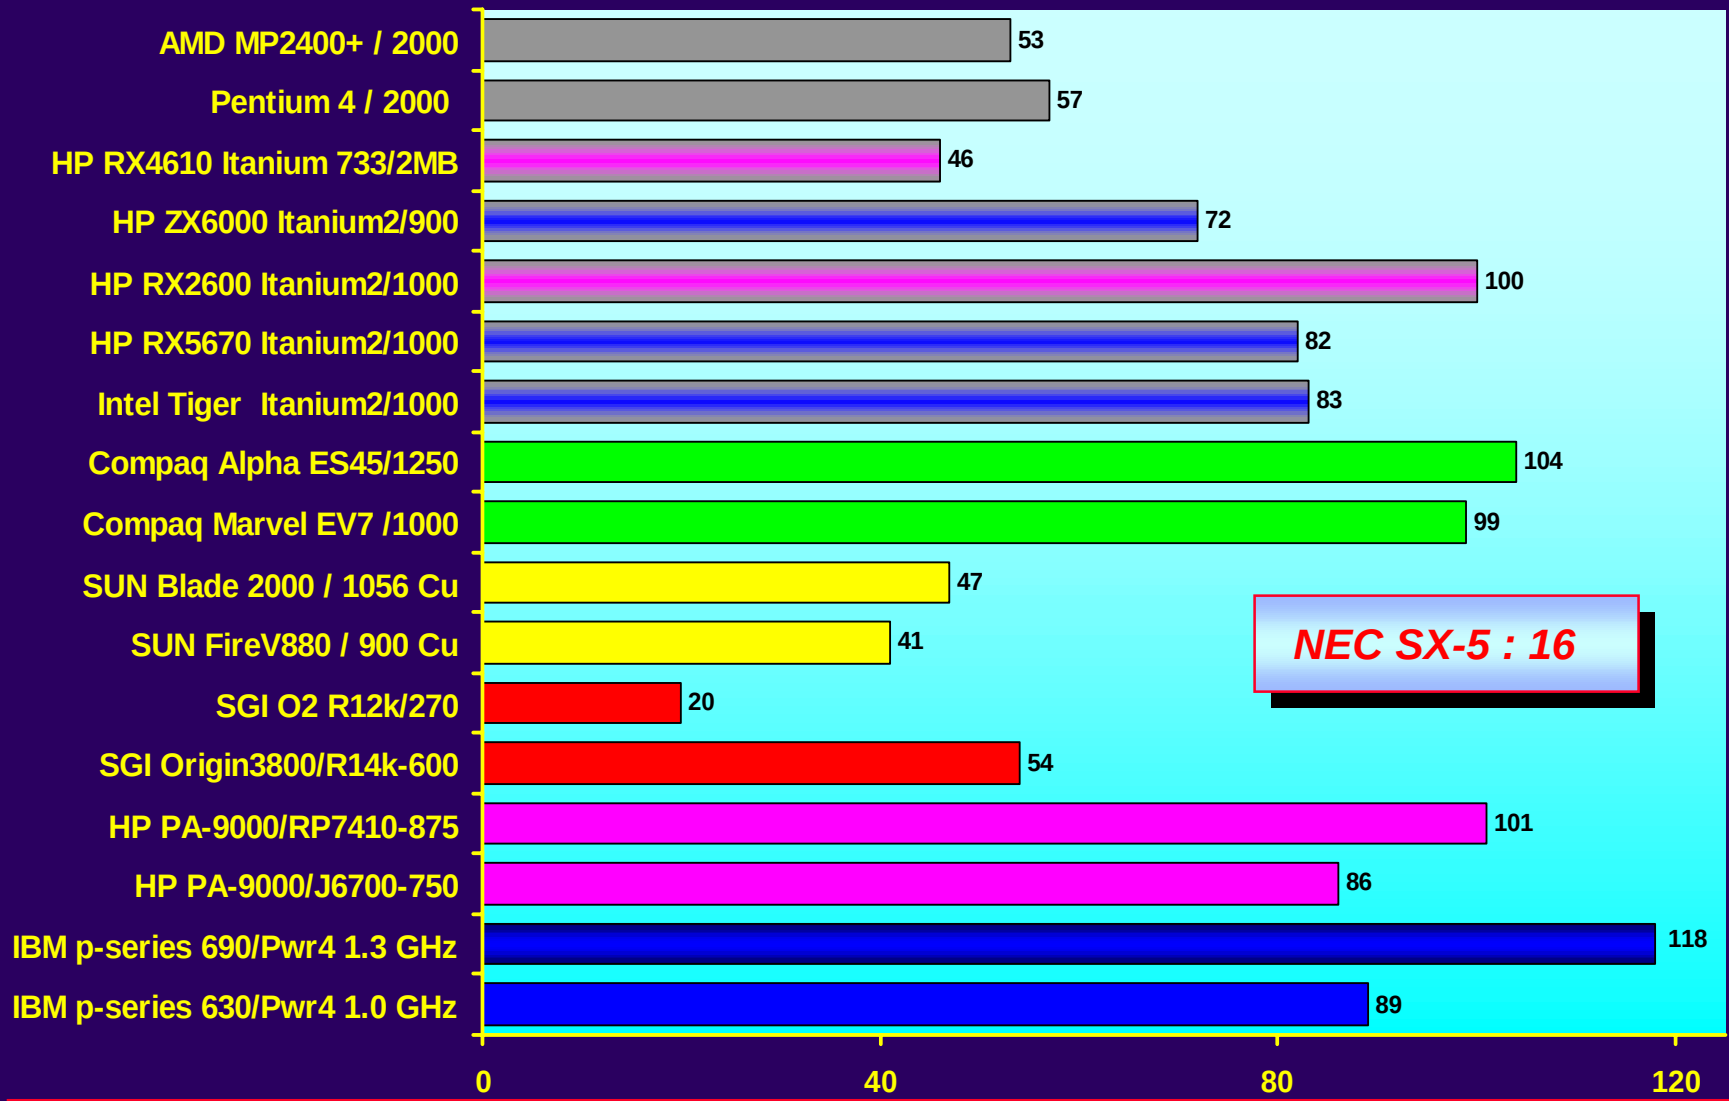

The Matrix-97 Benchmarks.

Performance relative to the HP RX2600 Itanium2/1000

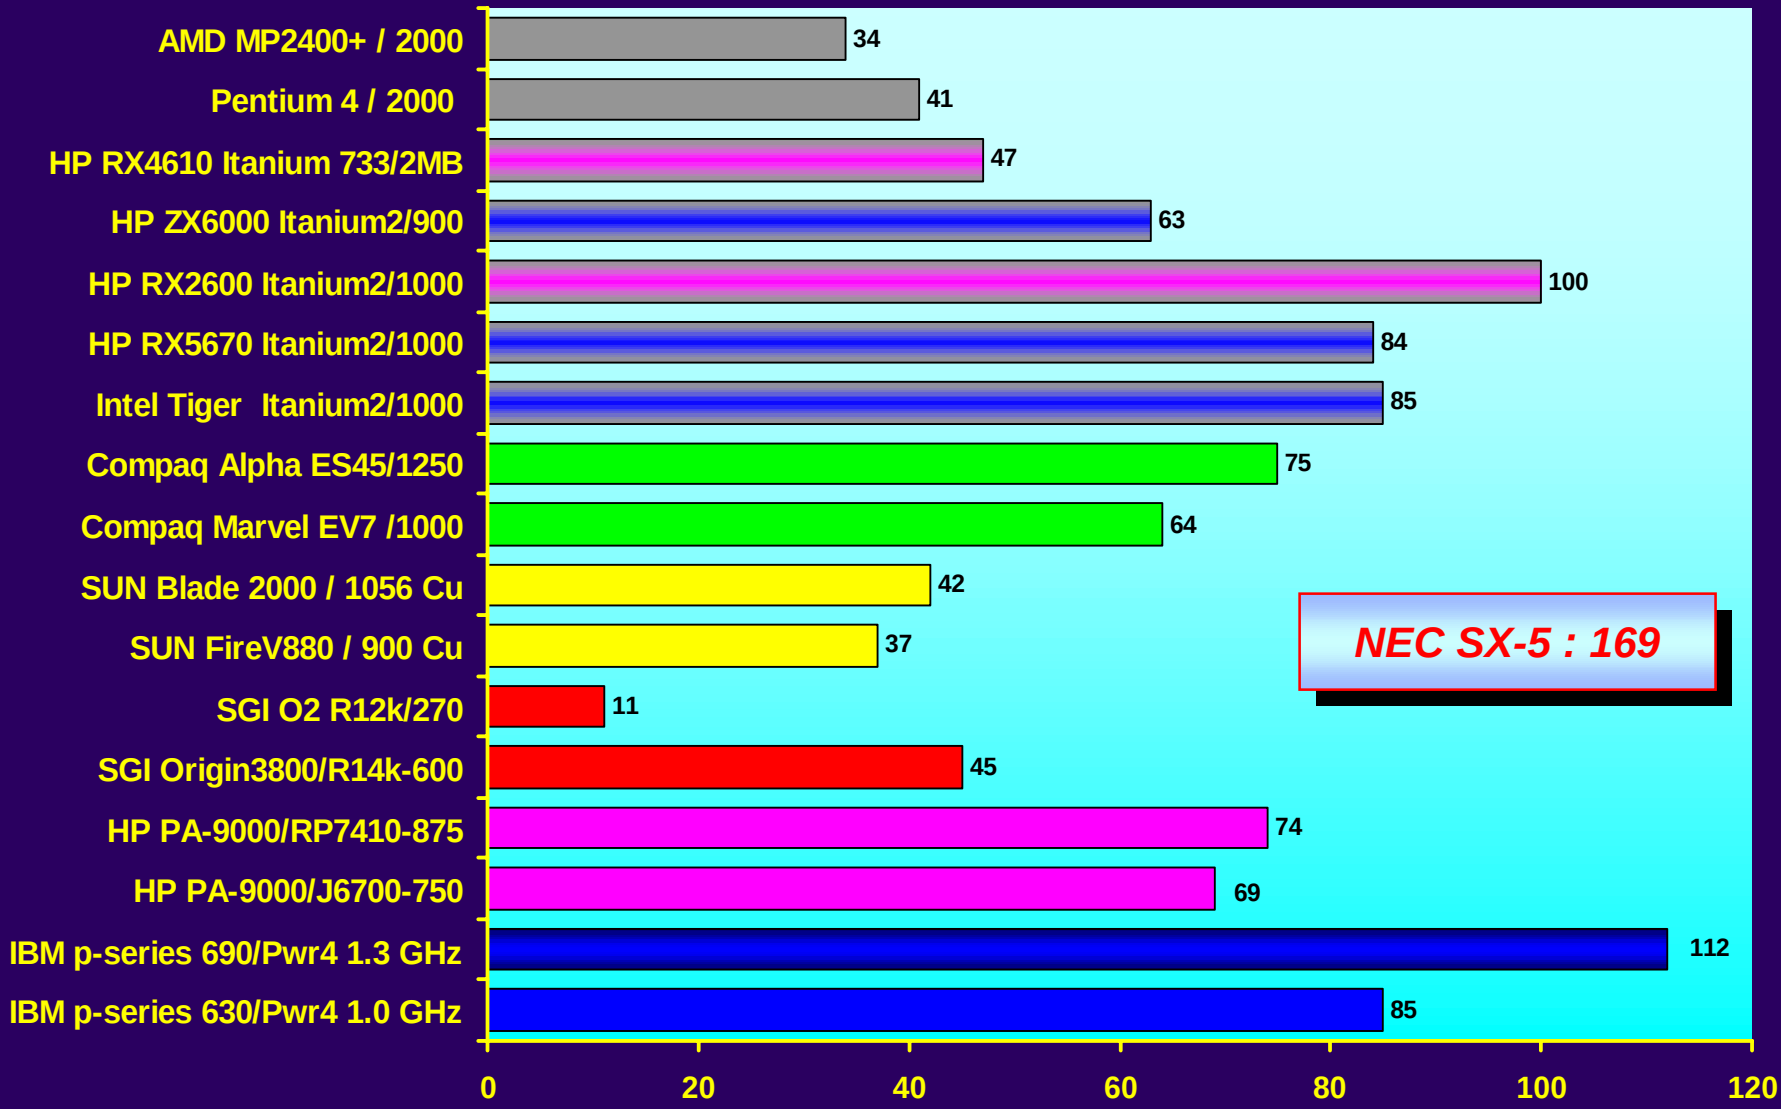

COMPUTATIONAL CHEMISTRY KERNELS

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_4 (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 HP RX2600 Itanium2/1000

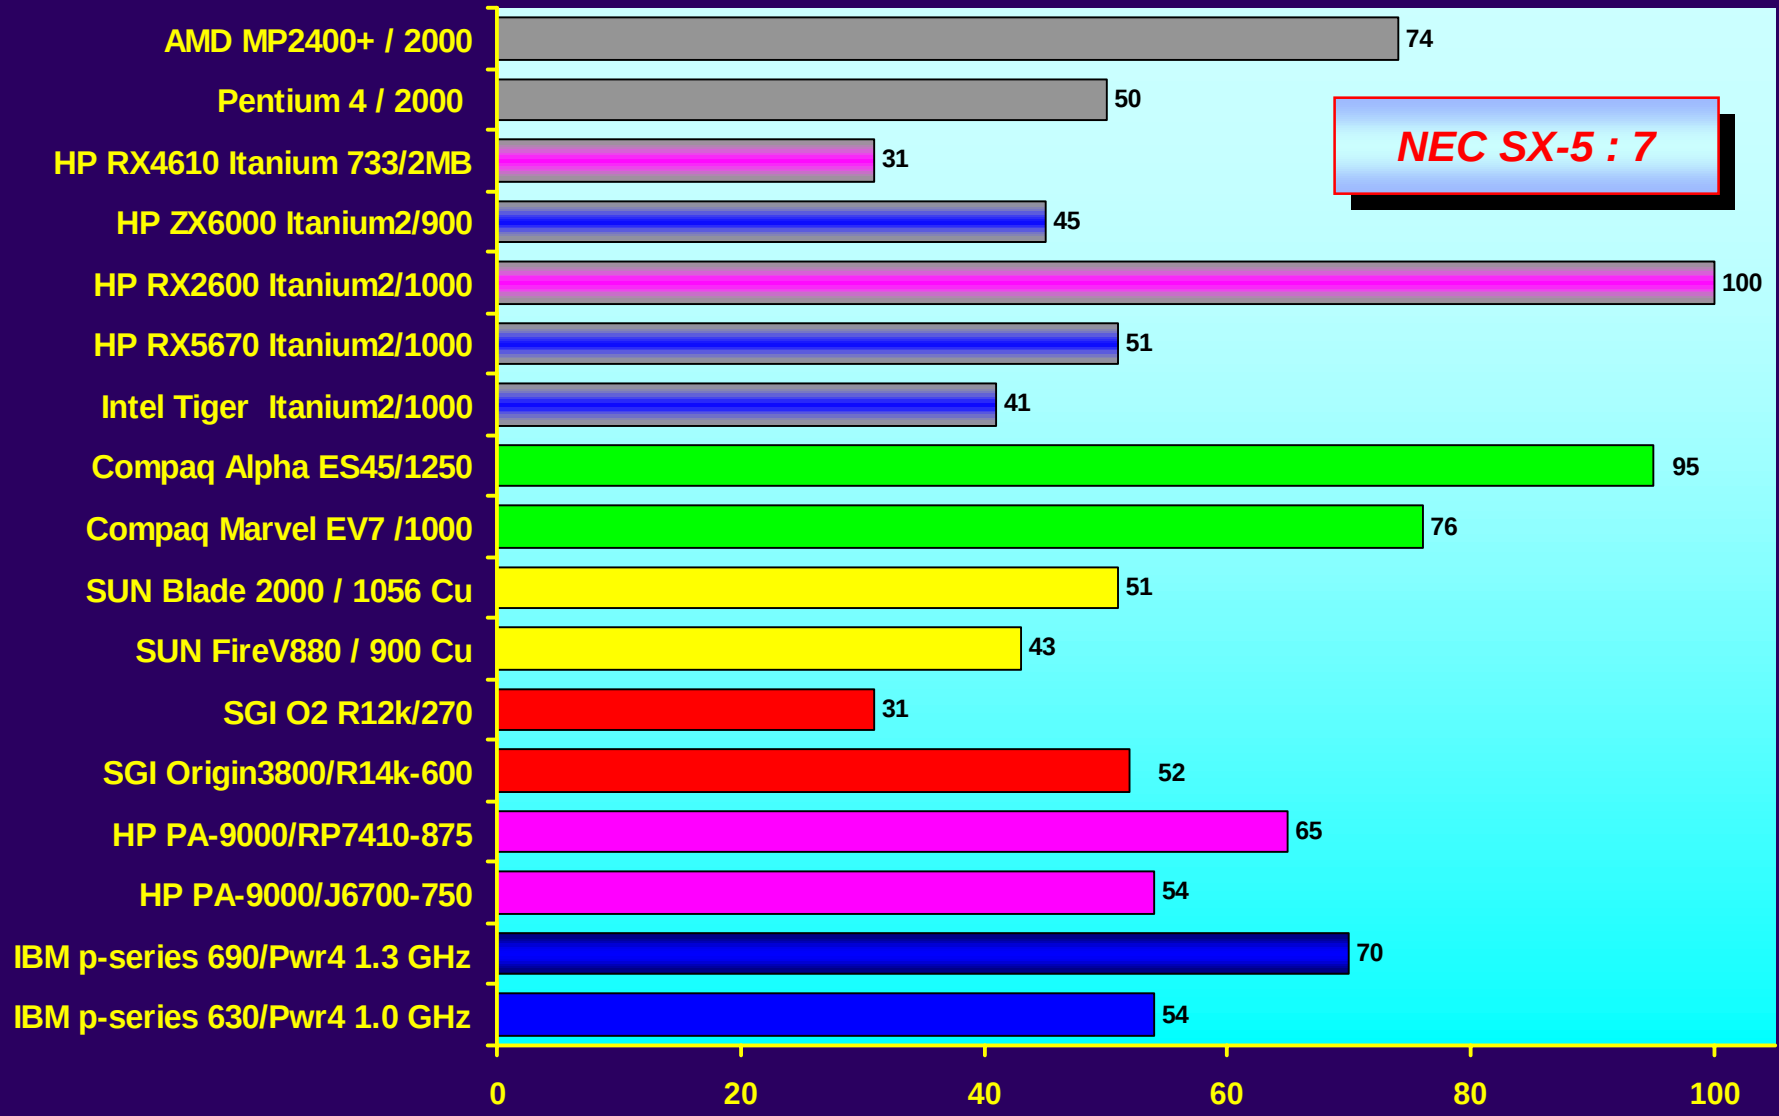

Chemistry Kernels - Molecular Dynamics.

Performance relative to the HP RX2600 Itanium2/1000

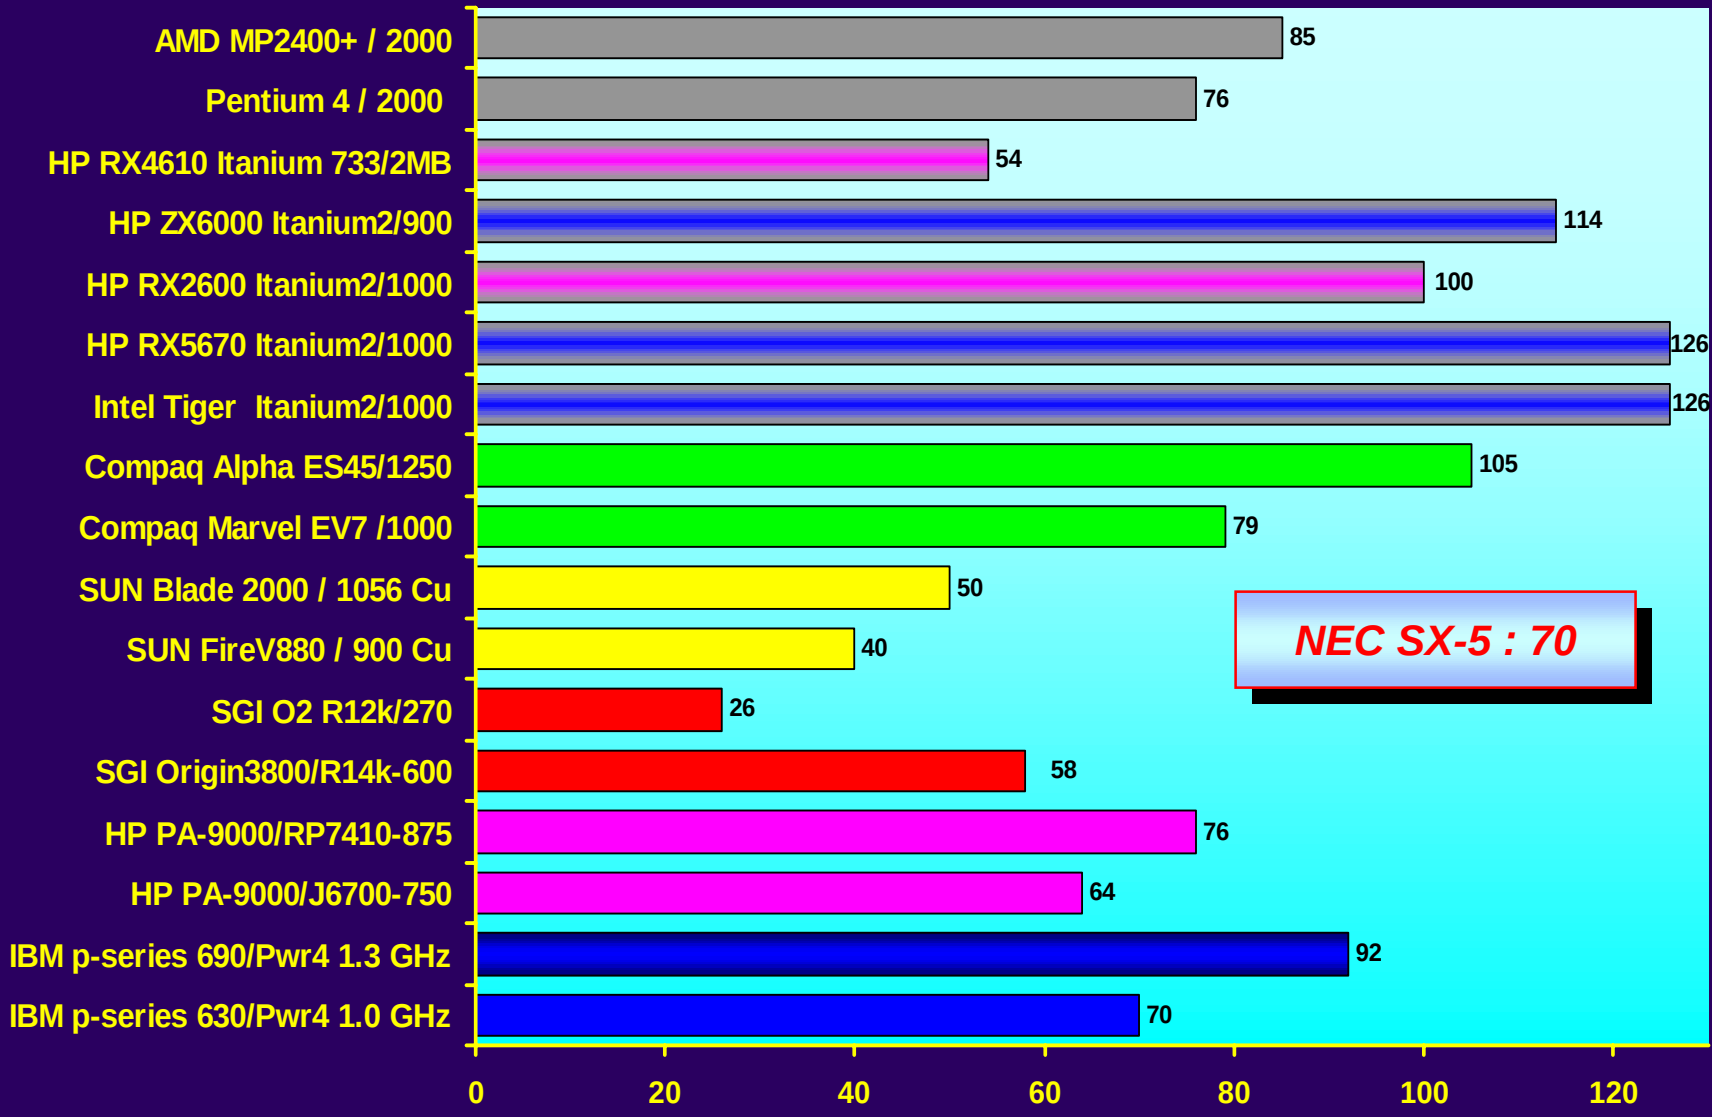

Chemistry Kernels - Quantum Monte Carlo.

Performance relative to the HP RX2600 Itanium2/1000

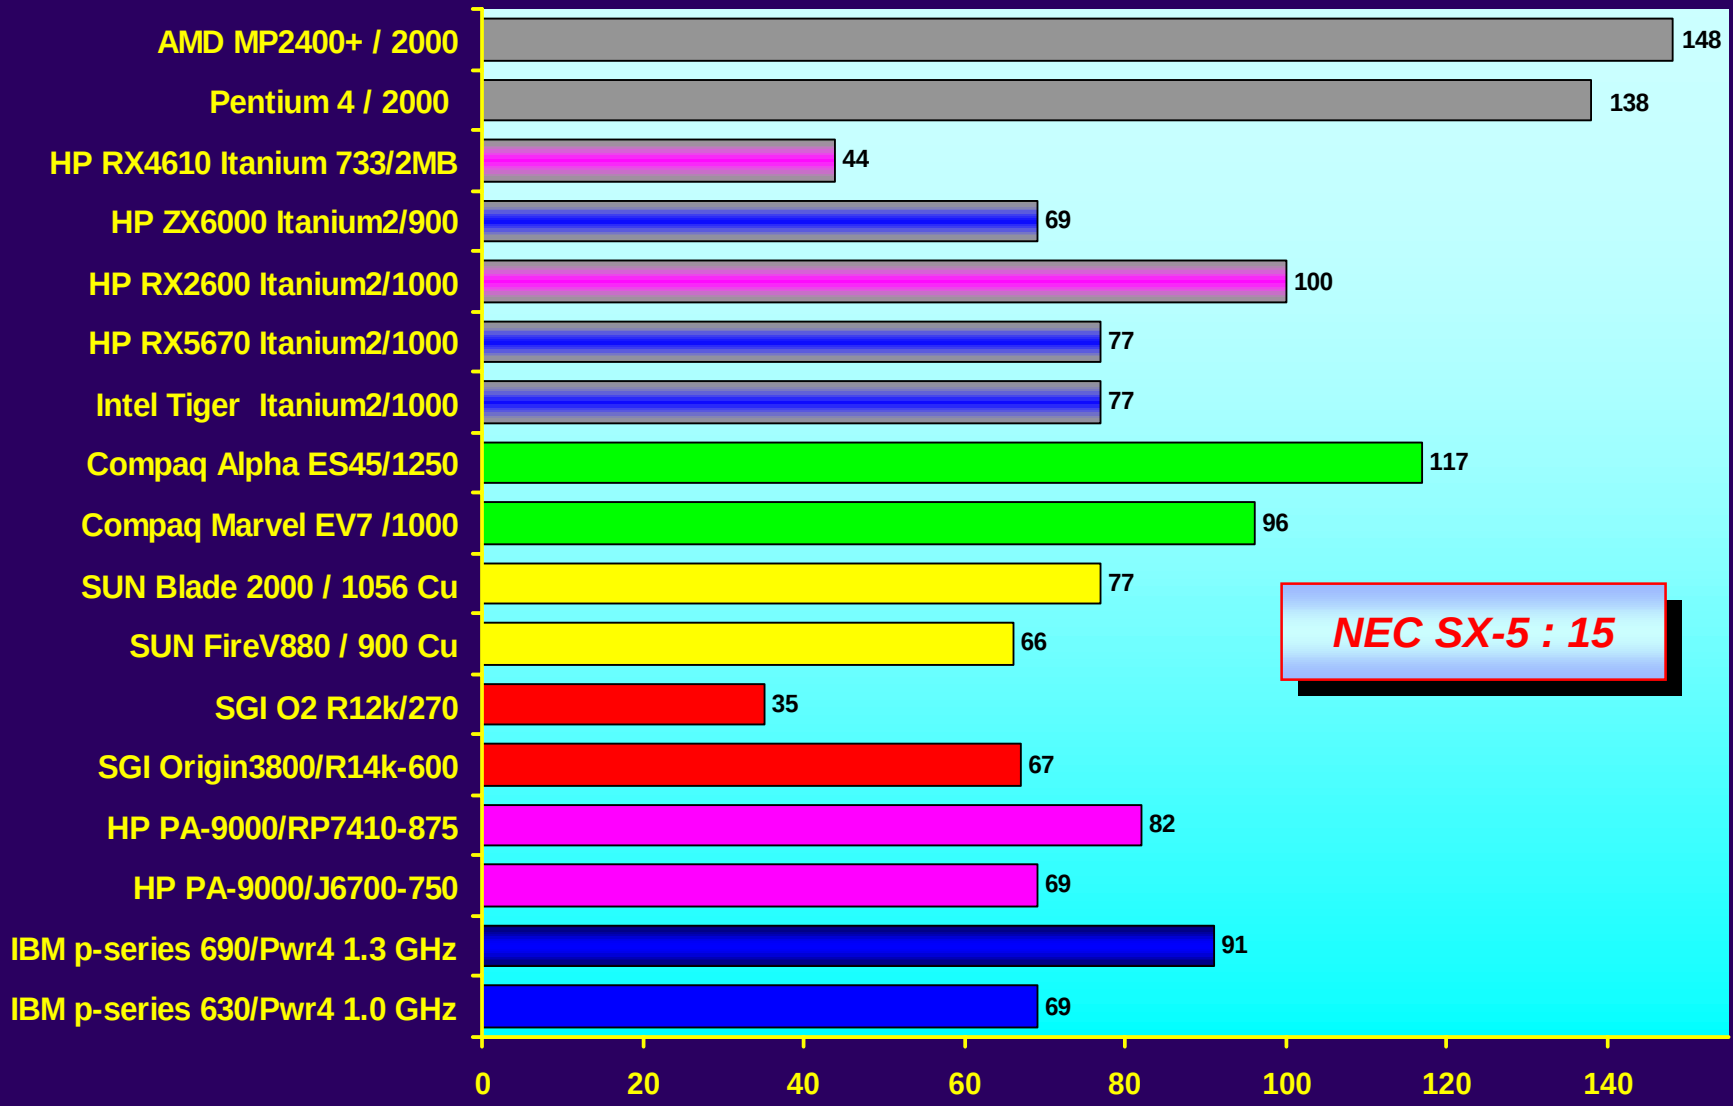

Chemistry Kernels - Jacobi Solver.

Performance relative to the HP RX2600 Itanium2/1000

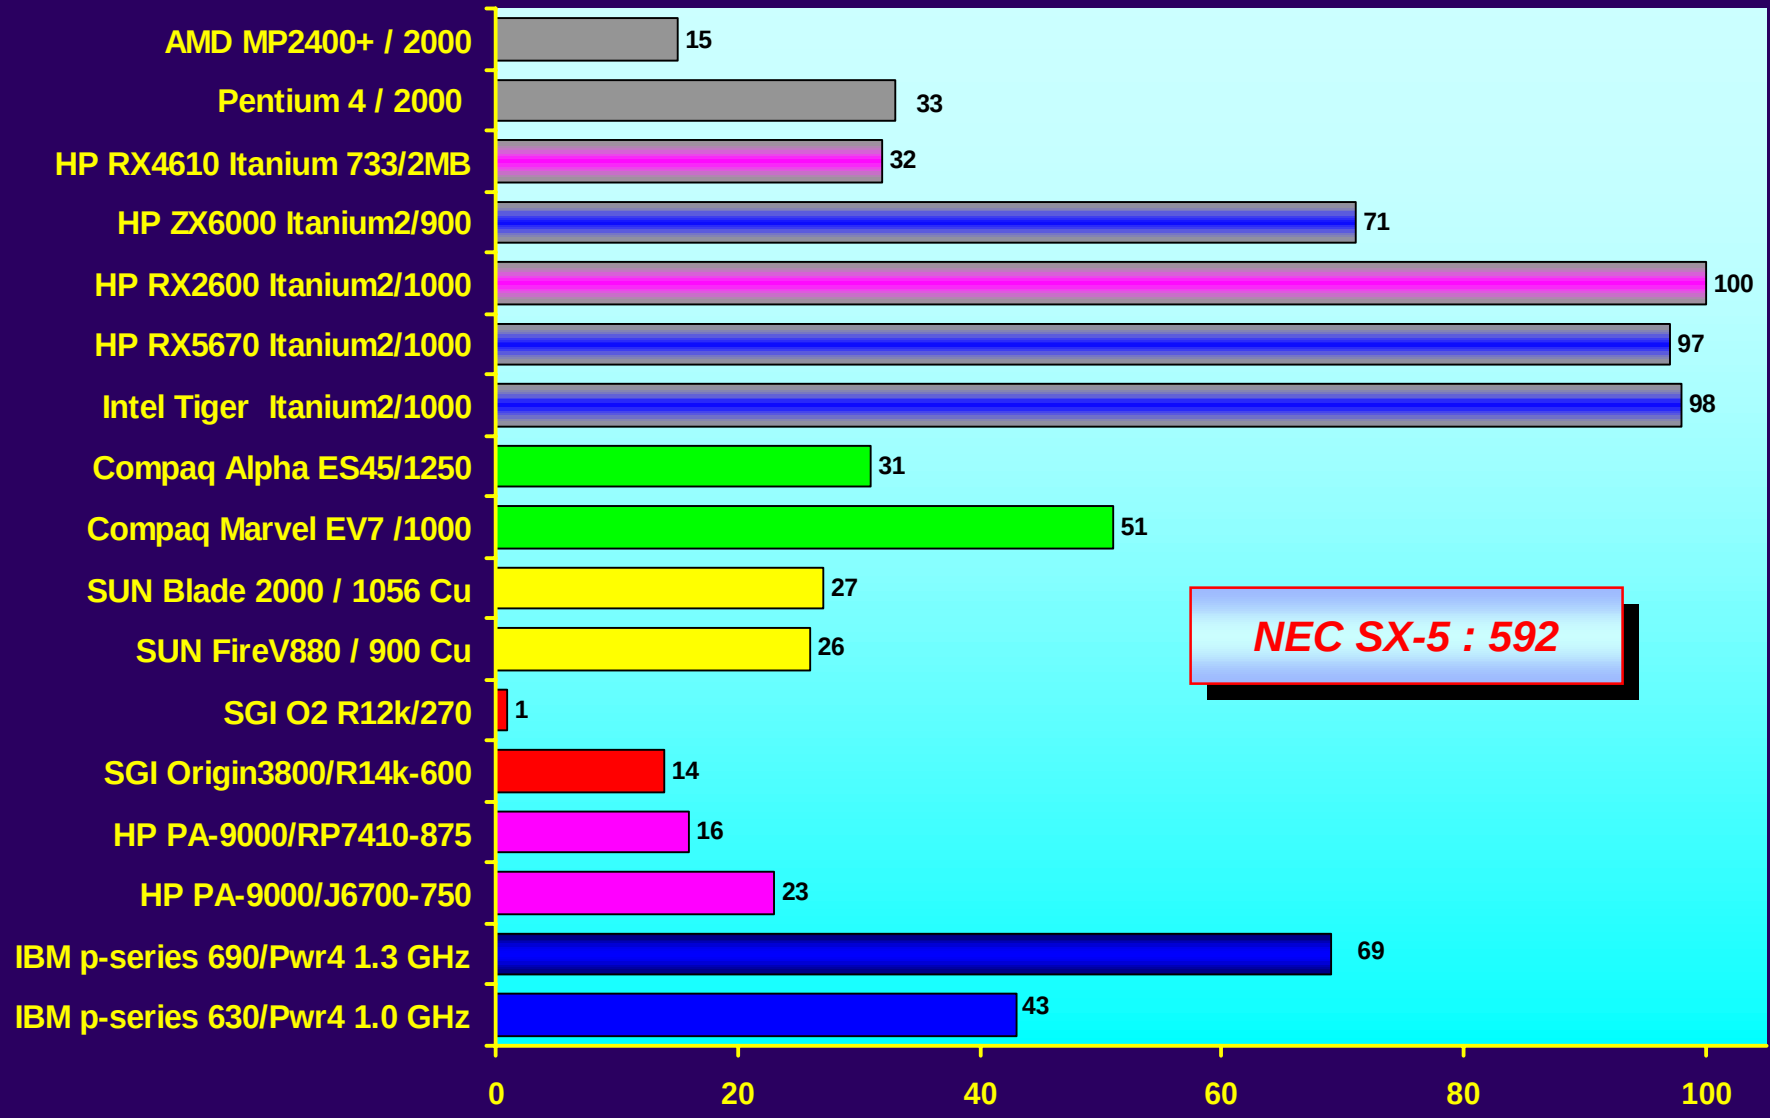

STREAM: Measured Sustainable Memory Bandwidth in HPC (TRIAD)

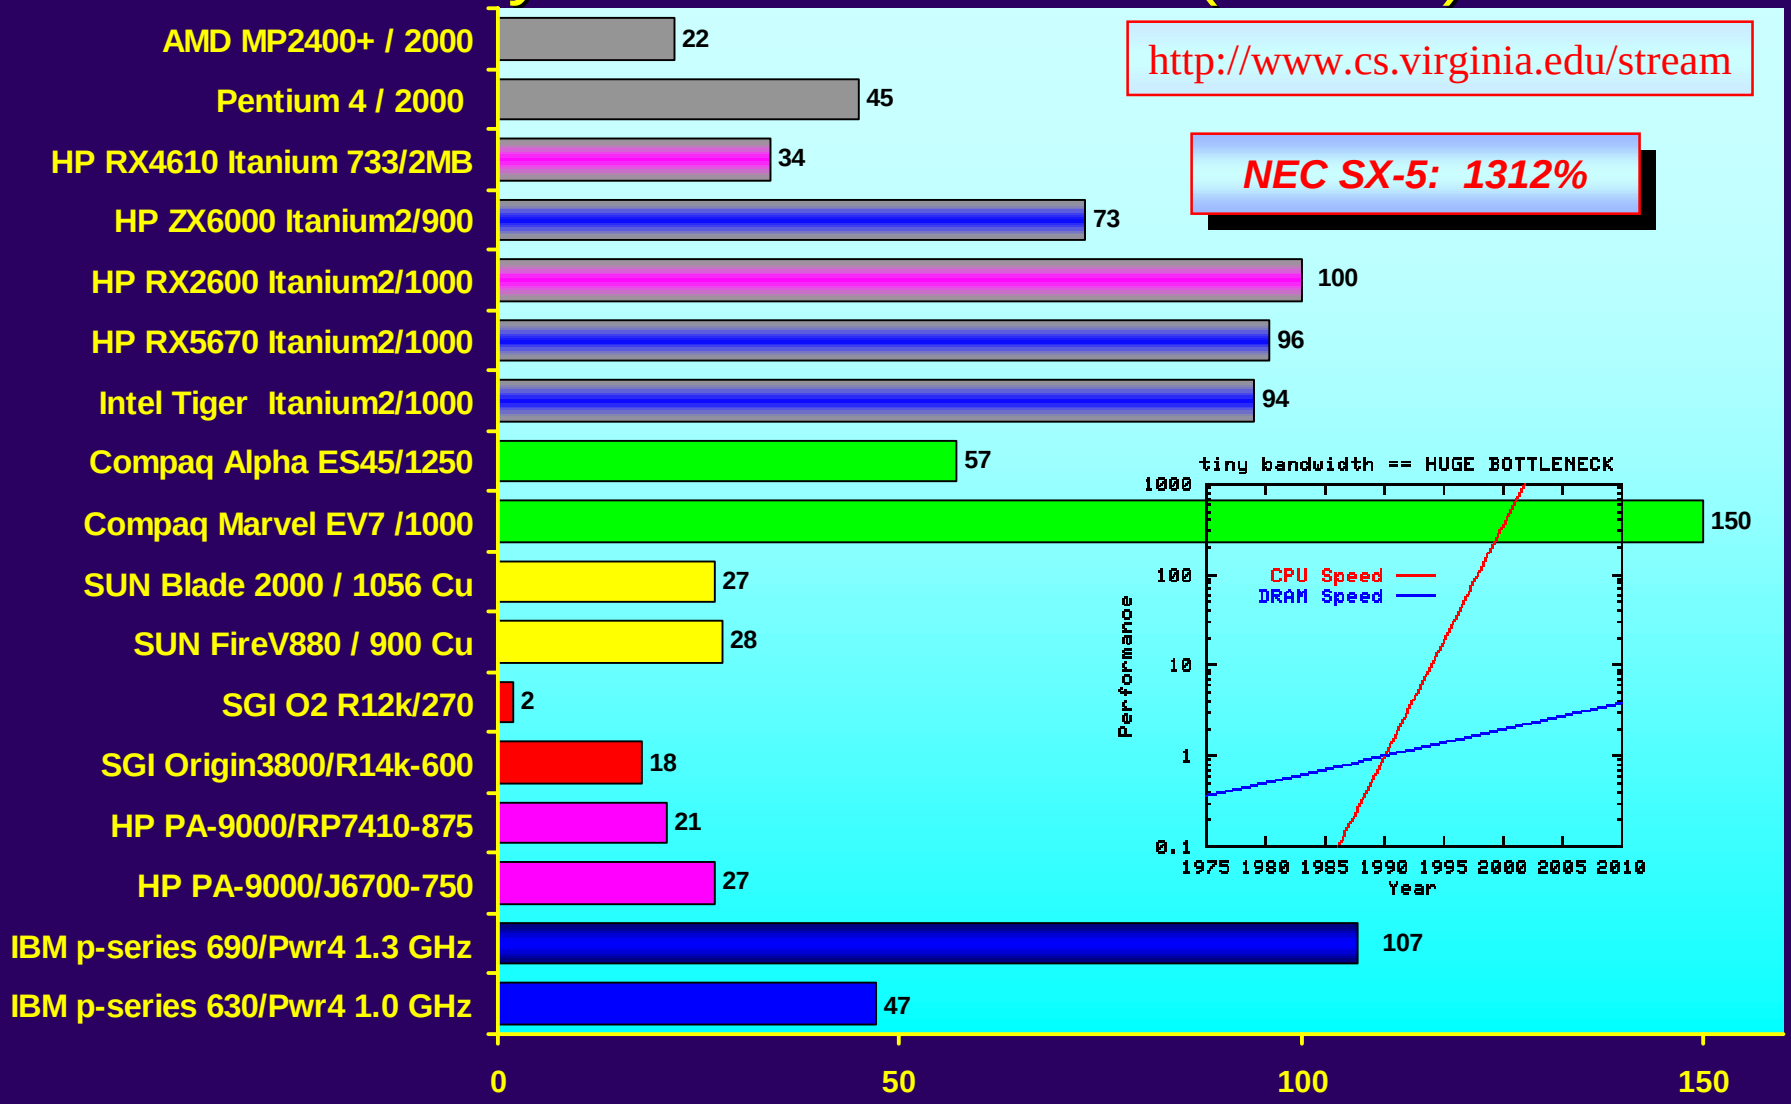

STREAM: Sustainable Memory Bandwidth (TRIAD) per process

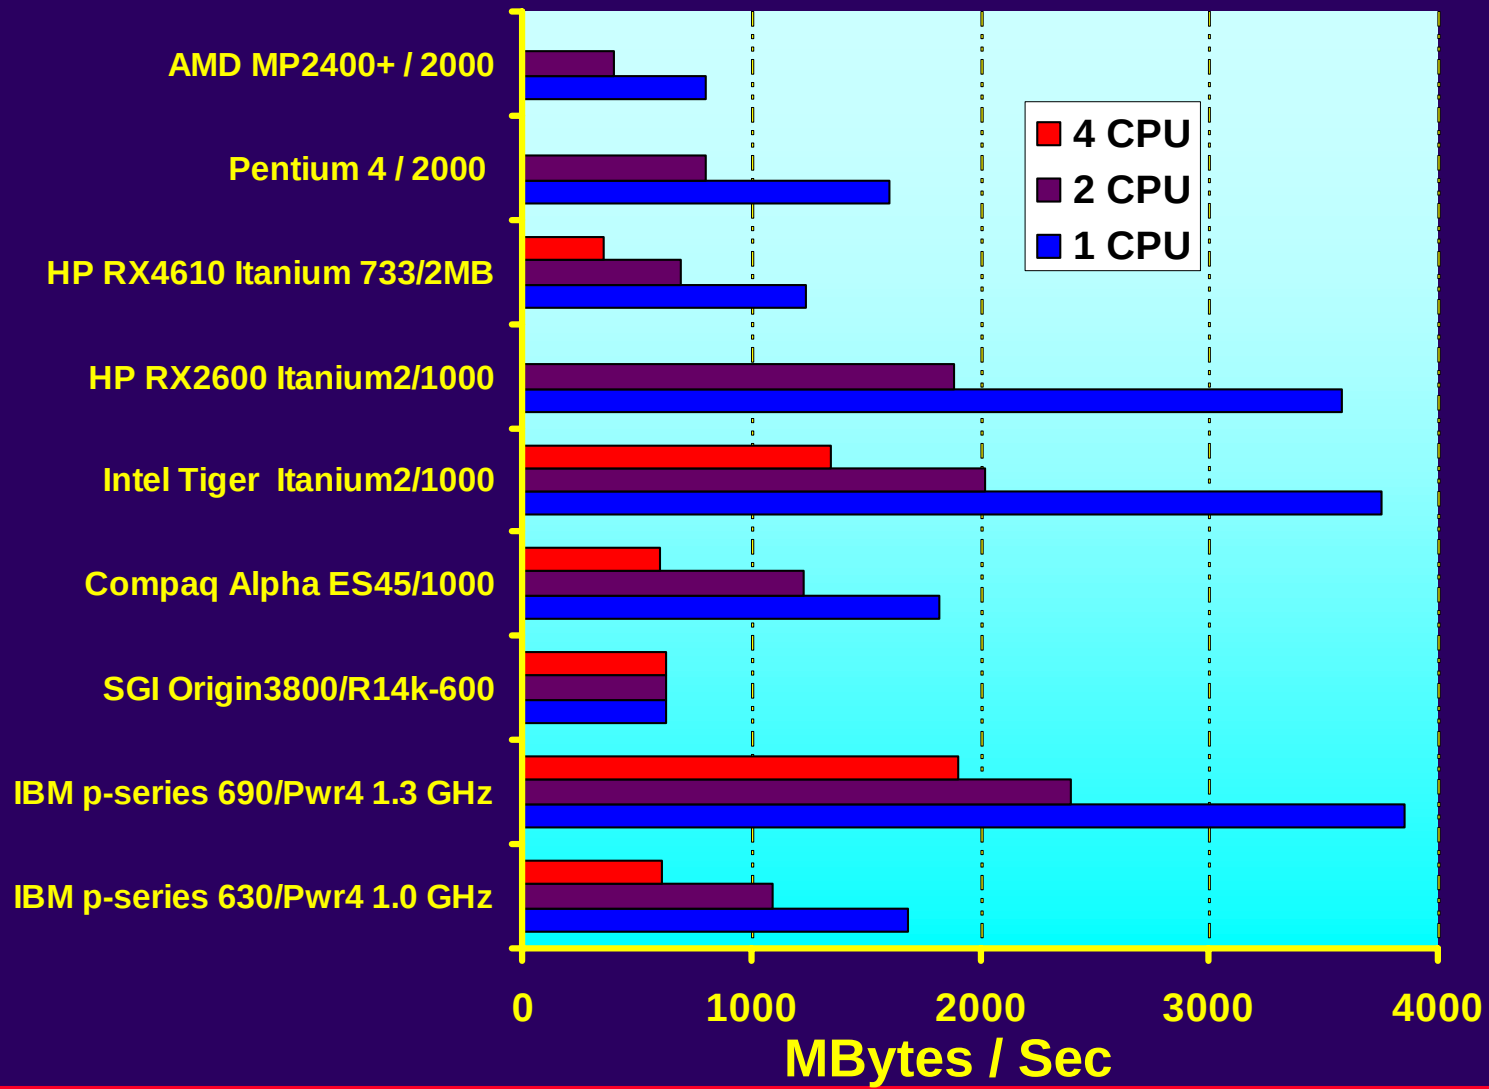

Computational Chemistry Kernels.

Performance relative to the HP RX2600 Itanium2/1000

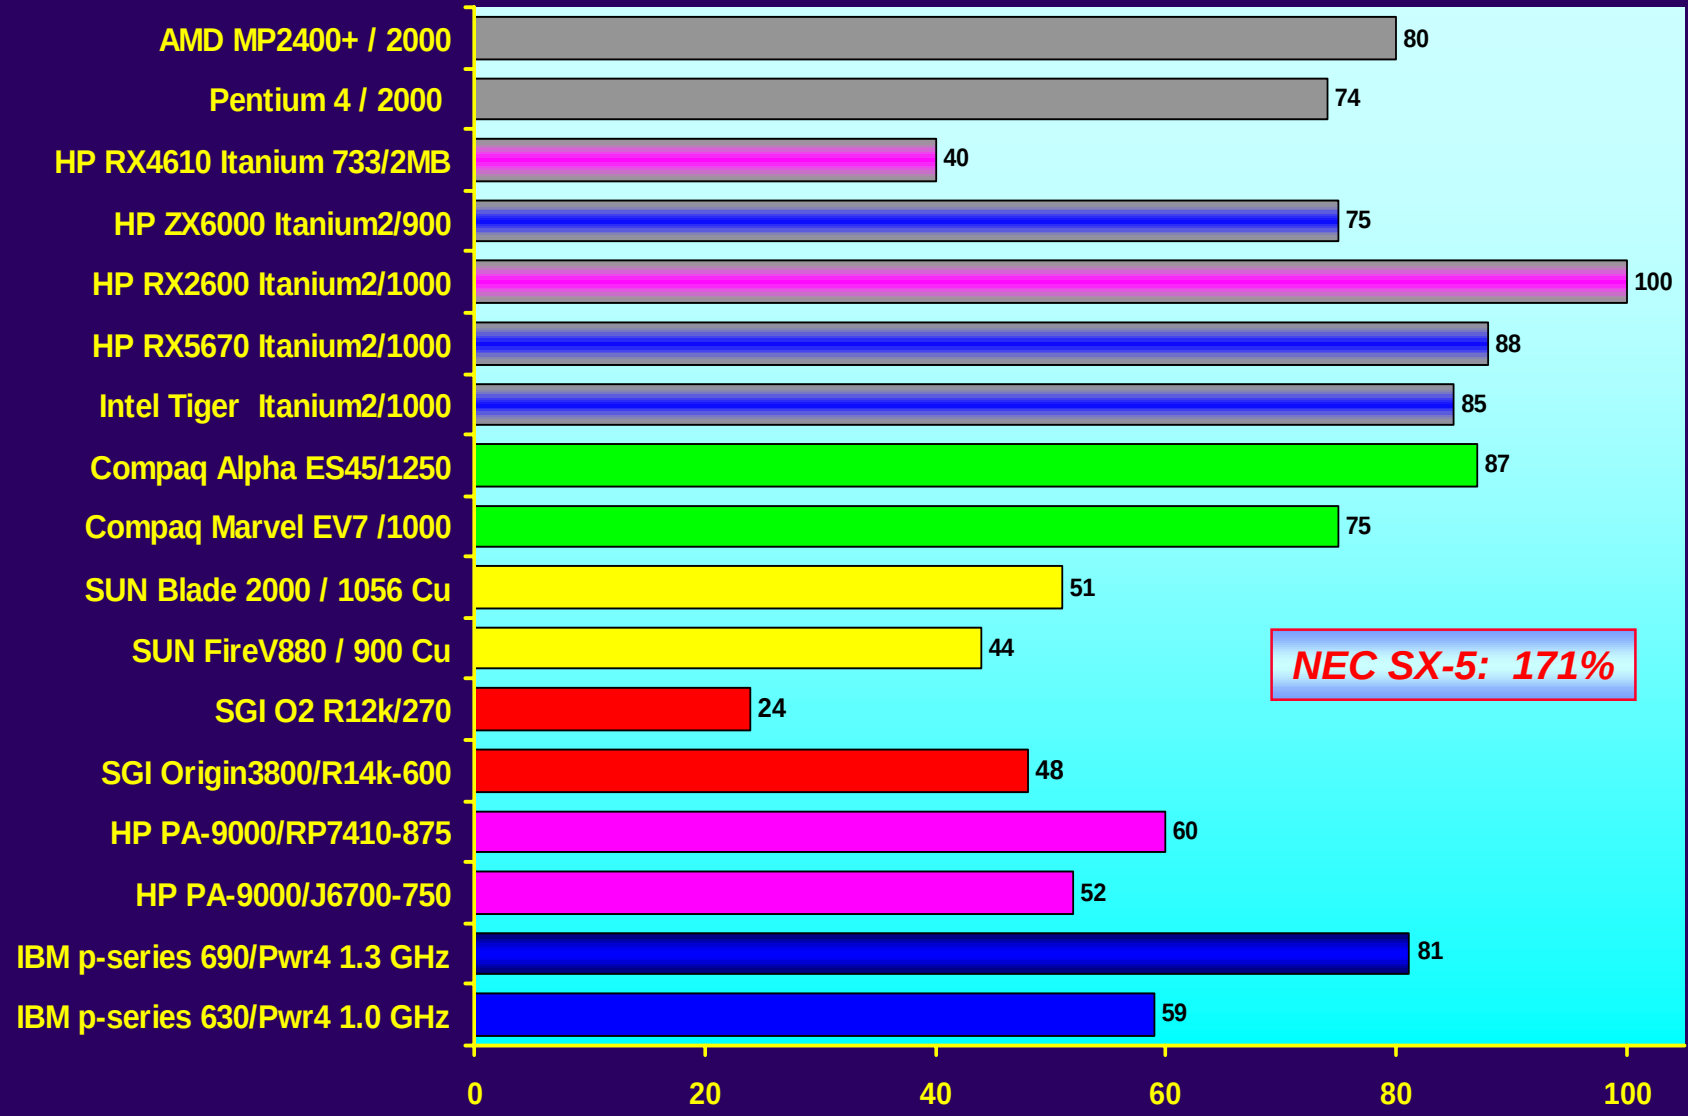

GAMESS-UK and DL_POLY Benchmarks

GAMESS-UK

DL_POLY

12 Typical QC Calculations

Module	Basis (GTOs)	Species
1. SCF	STO-3G (124)	Morphine
2. SCF	6-31G (154)	$C_6H_3(NO_2)_3$
3. ECP Geometry	ECPDZ (70)	Na_7Mg^+
4. Direct-SCF	6-31G (82)	Cytosine
5. CAS-geometry	TZVP (52)	H_2CO
6. MCSCF	EXT1 (74)	H_2CO
7. Direct-CI	EXT2 (64)	H_2CO/H_2+CO
8. MRD-CI (26M)	ECP (59)	$TiCl_4$
9. MP2-geometry	6-31G* (70)	H_3SiNCO
10. SCF 2nd derivs.	6-31G (64)	C_5H_5N
11. MP2 2nd derivs.	6-31G* (60)	C_4
12. Direct-MP2	DZP (76)	C_5H_5N

Six Typical Simulations

Simulation	Atoms	Time steps
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 MgCl2 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 HP RX2600 Itanium2/1000

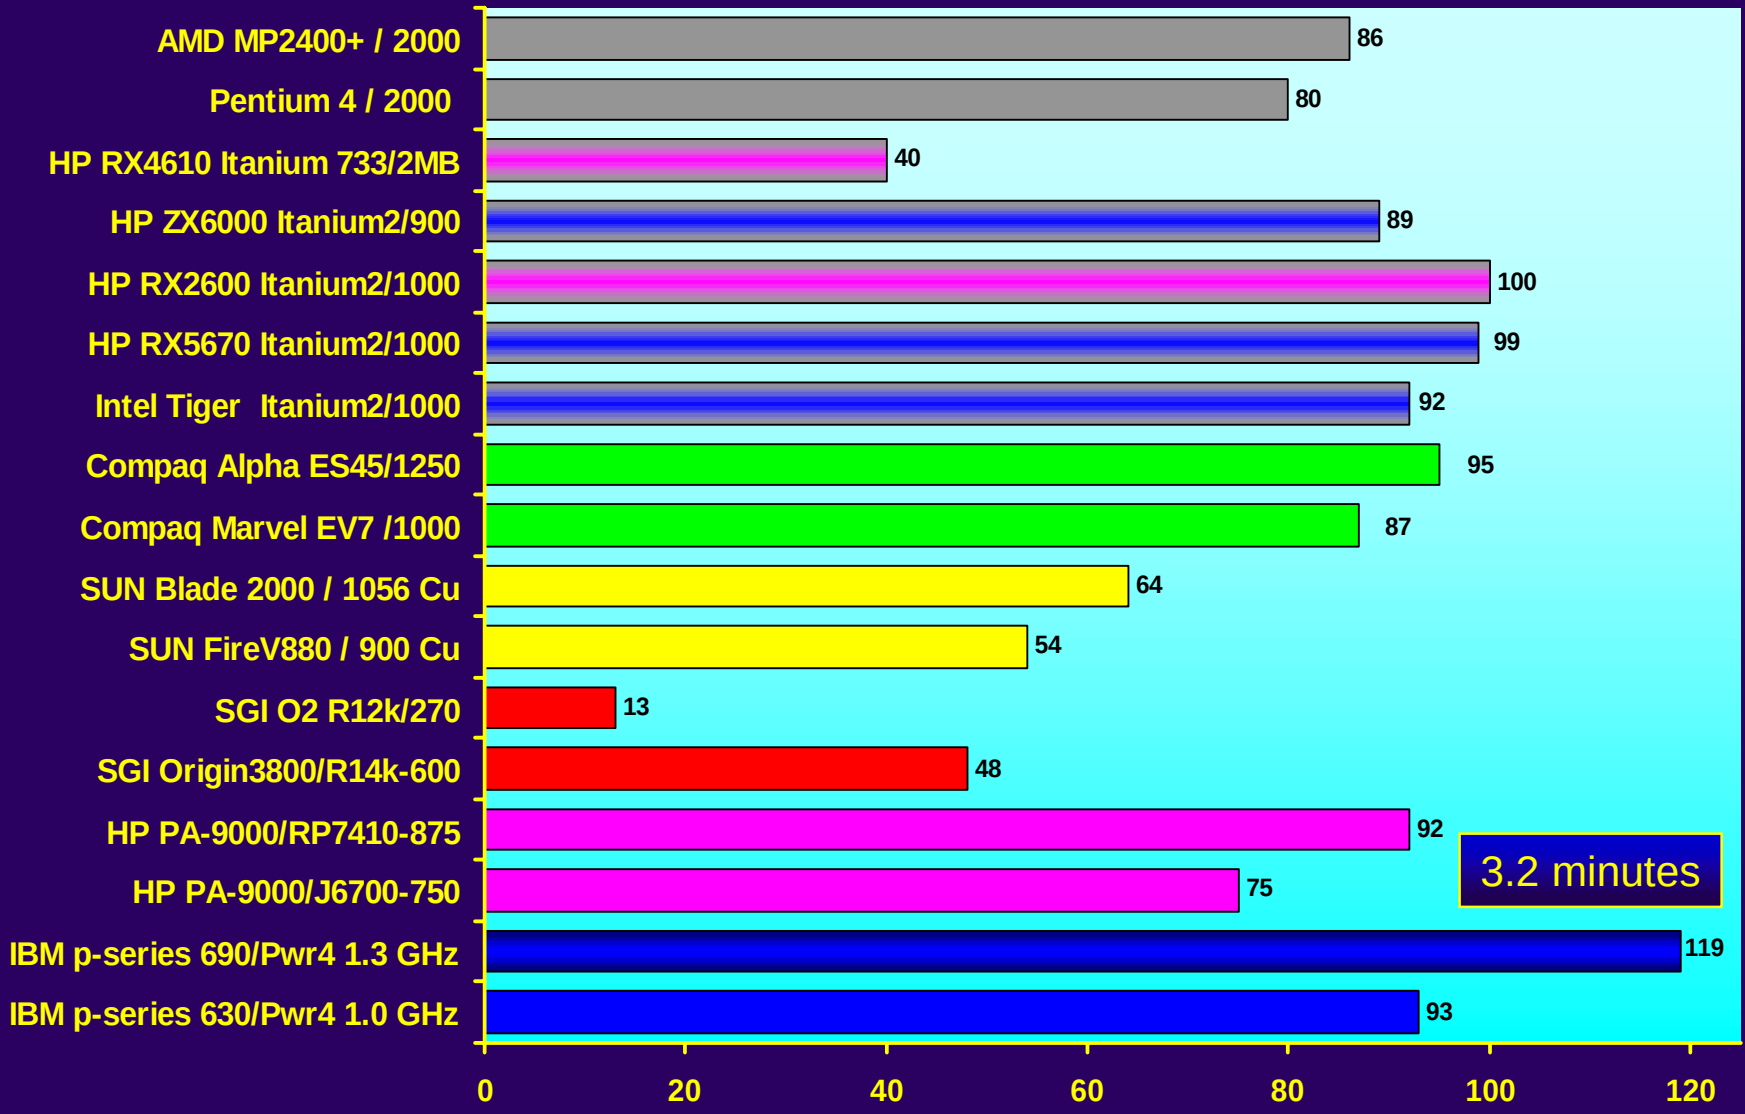

SPECfp2000 and the GAMESS-UK Benchmark

The GAMESS-UK-99 Benchmark

10 Typical QC Calculations

Module	Basis (GTOs)	Species
1. Direct- SCF	6-31G (227)	Morphine
2. SCF	6-31G** (265)	$C_6H_3(NO_2)_3$
3. DFT B3LYP	6-311G* (167)	Cytosine
4. MCSCF	CC-PVTZ (100)	H_2CO
5. Direct-CI	CC-PVTZ (100)	H_2CO/H_2+CO
6. CCSD(T)	TZV+2d+1f (144)	C_4
7. MP2-geometry	TZVP (105)	H_3SiNCO
8. SCF 2nd derivs.	6-311G** (144)	C_5H_5N
9. MP2 2nd derivs.	TZVP(C2d) (104)	C_4
10. Direct-MP2	6-31G* (130)	Cytosine

Benchmark Time: 29.8 minutes on HP RX2600 Itanium2/1000

The GAMESS-UK-99 Benchmark

Performance relative to the HP RX2600 Itanium2/1000

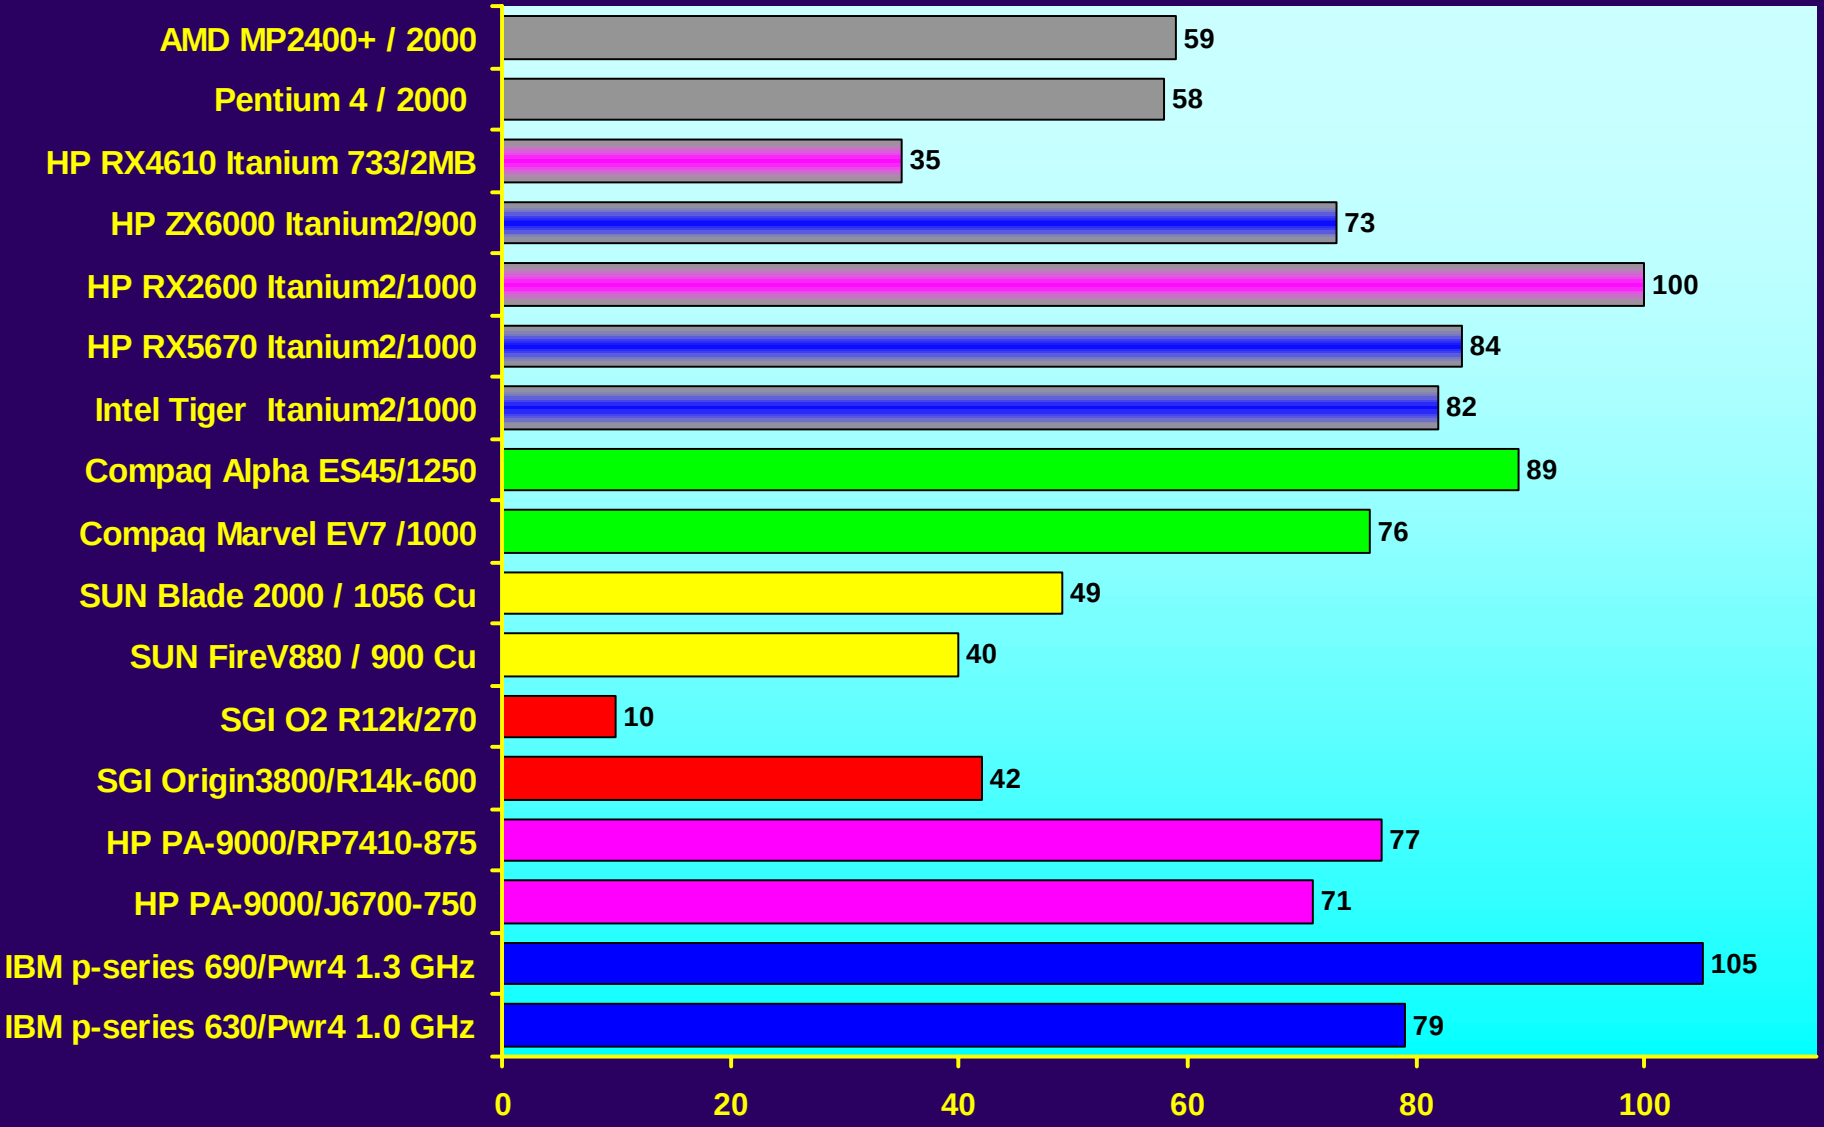

The GAMESS-UK and GAMESS-UK-99 Benchmark

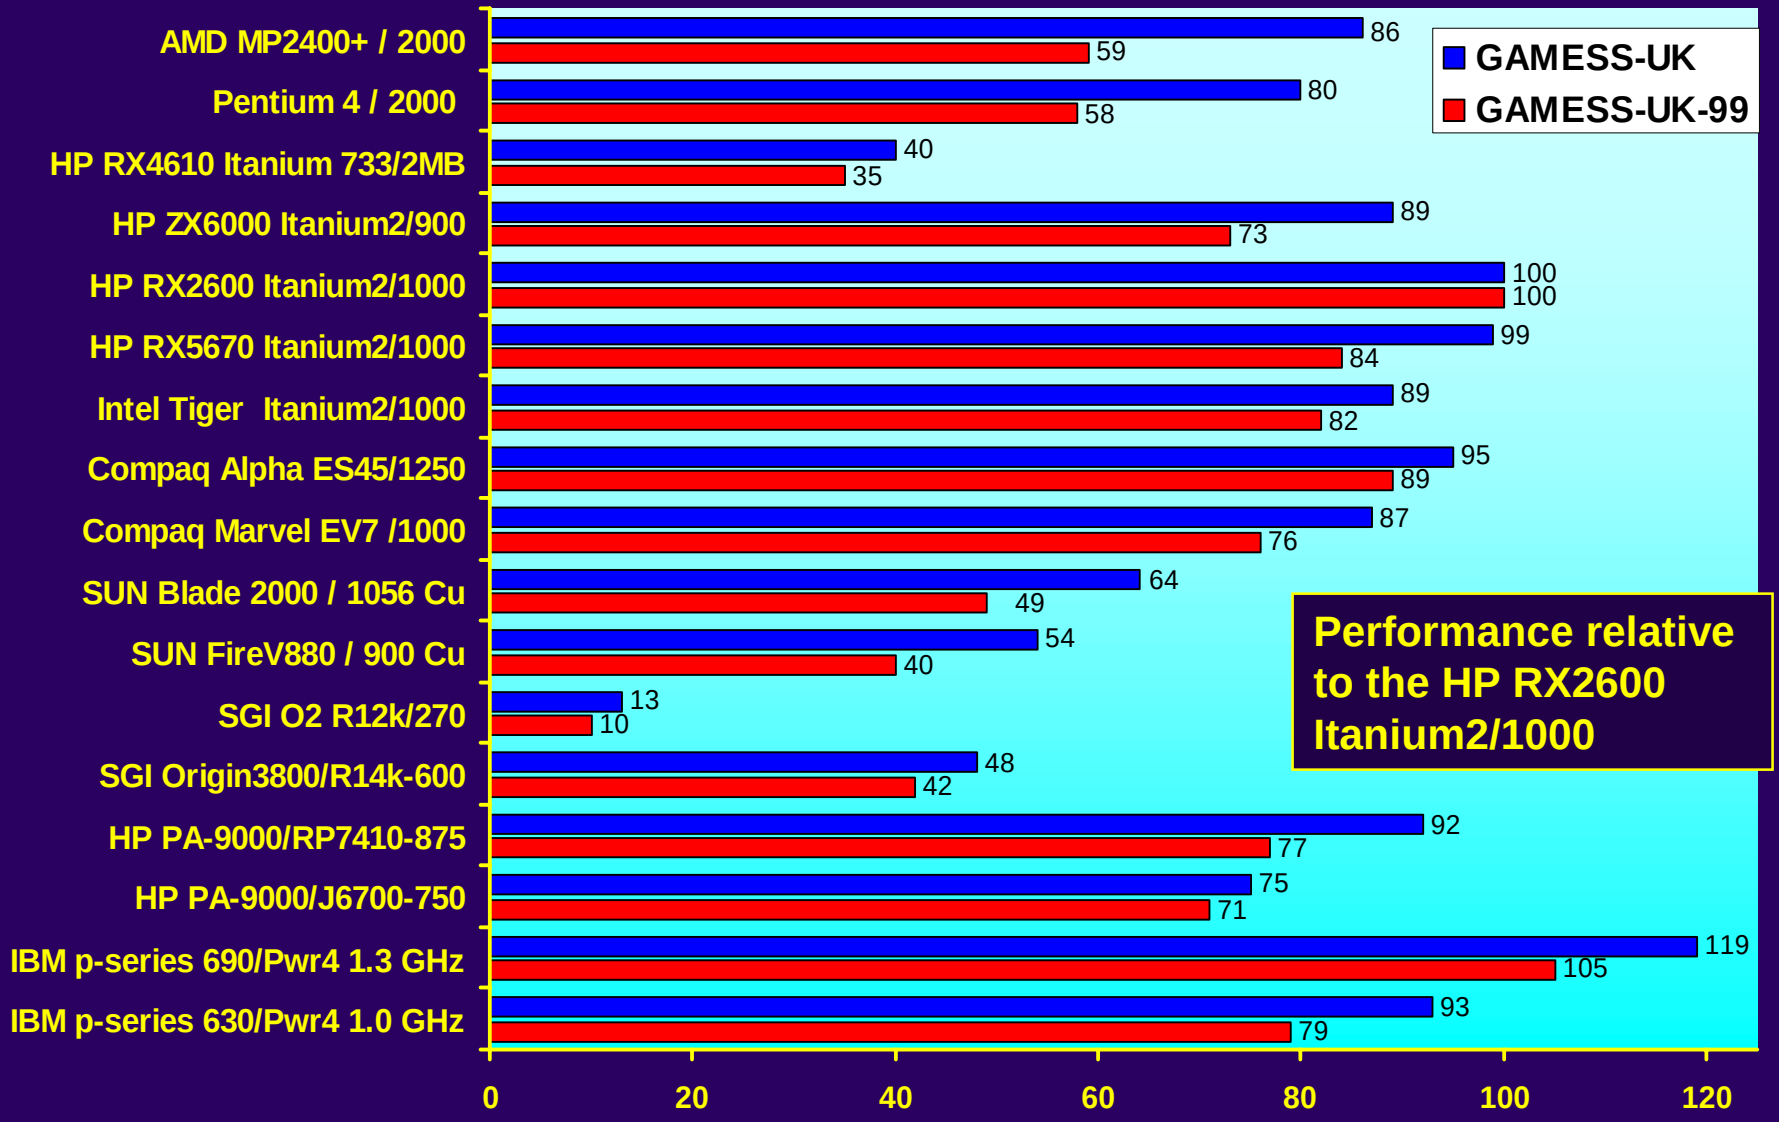

The DL_POLY Benchmark.

Performance relative to the HP RX2600 Itanium2/1000

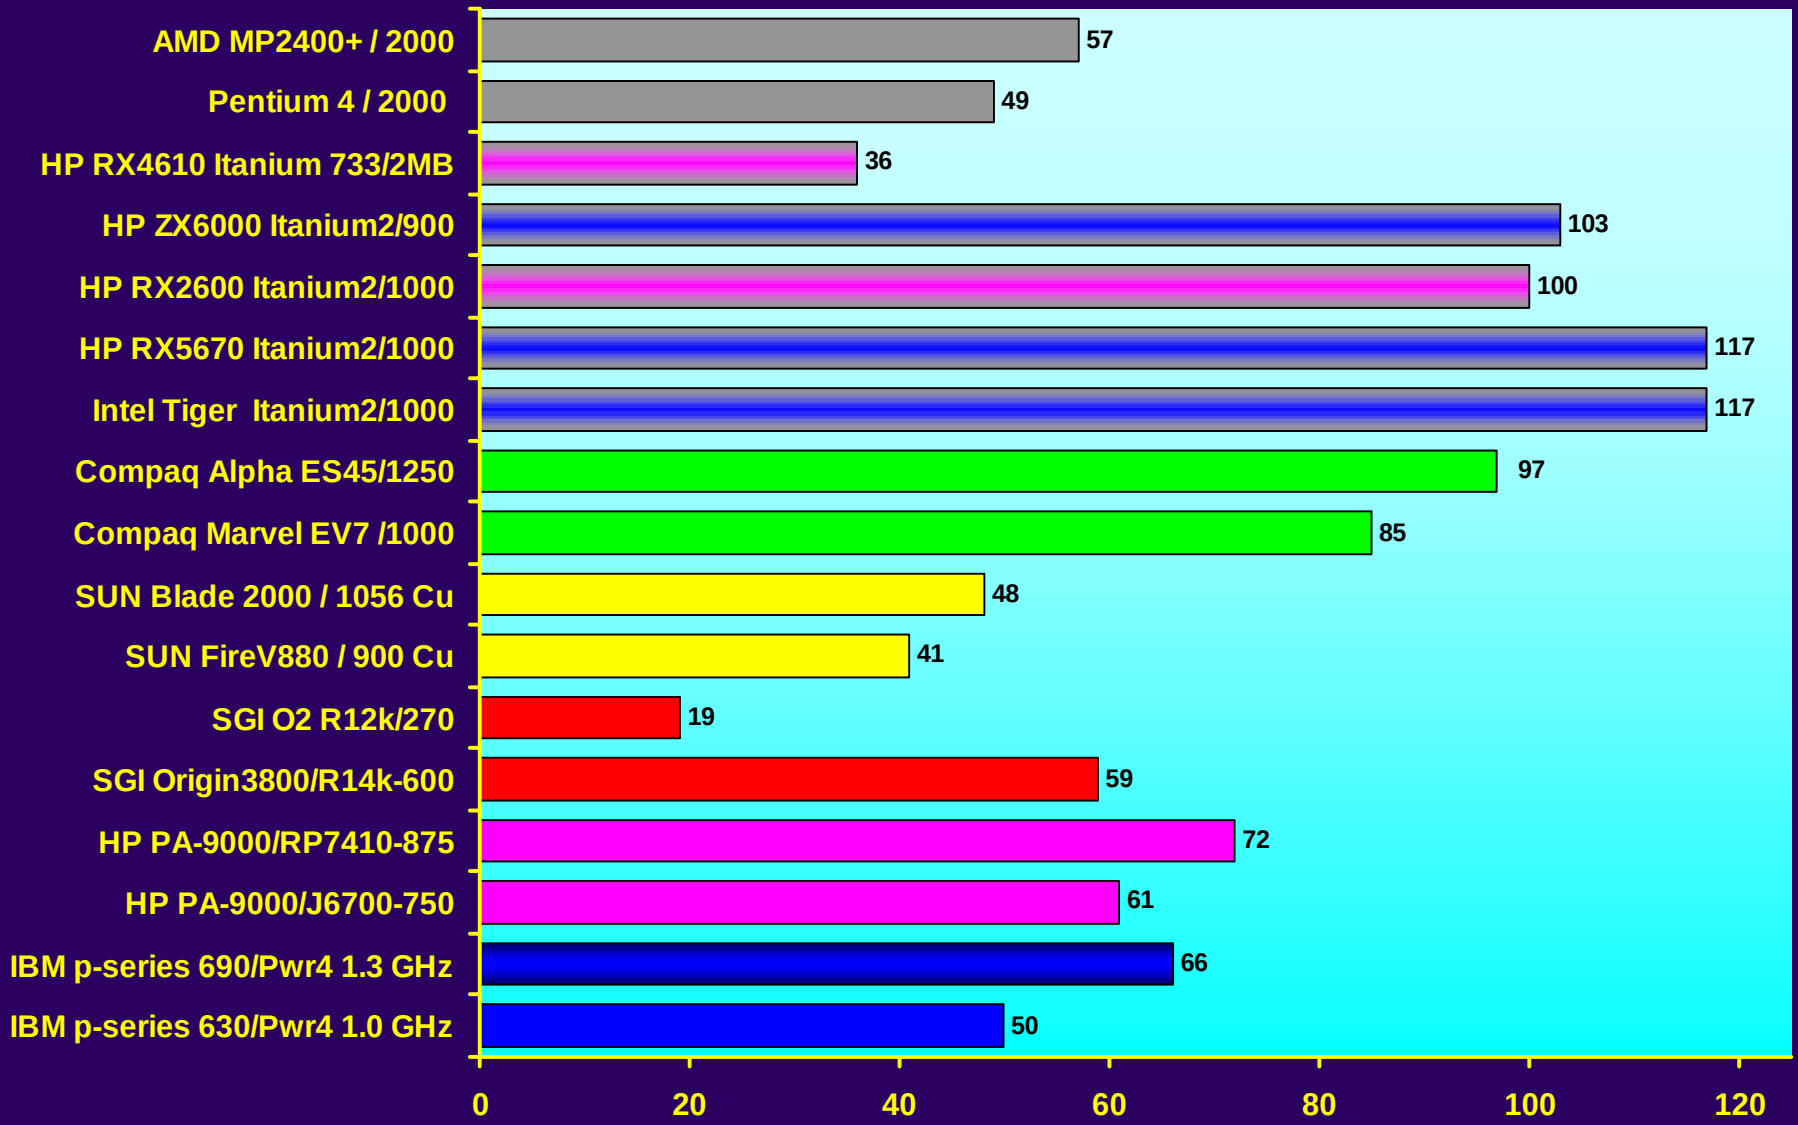

SPECfp2000 and the DL_POLY Benchmark

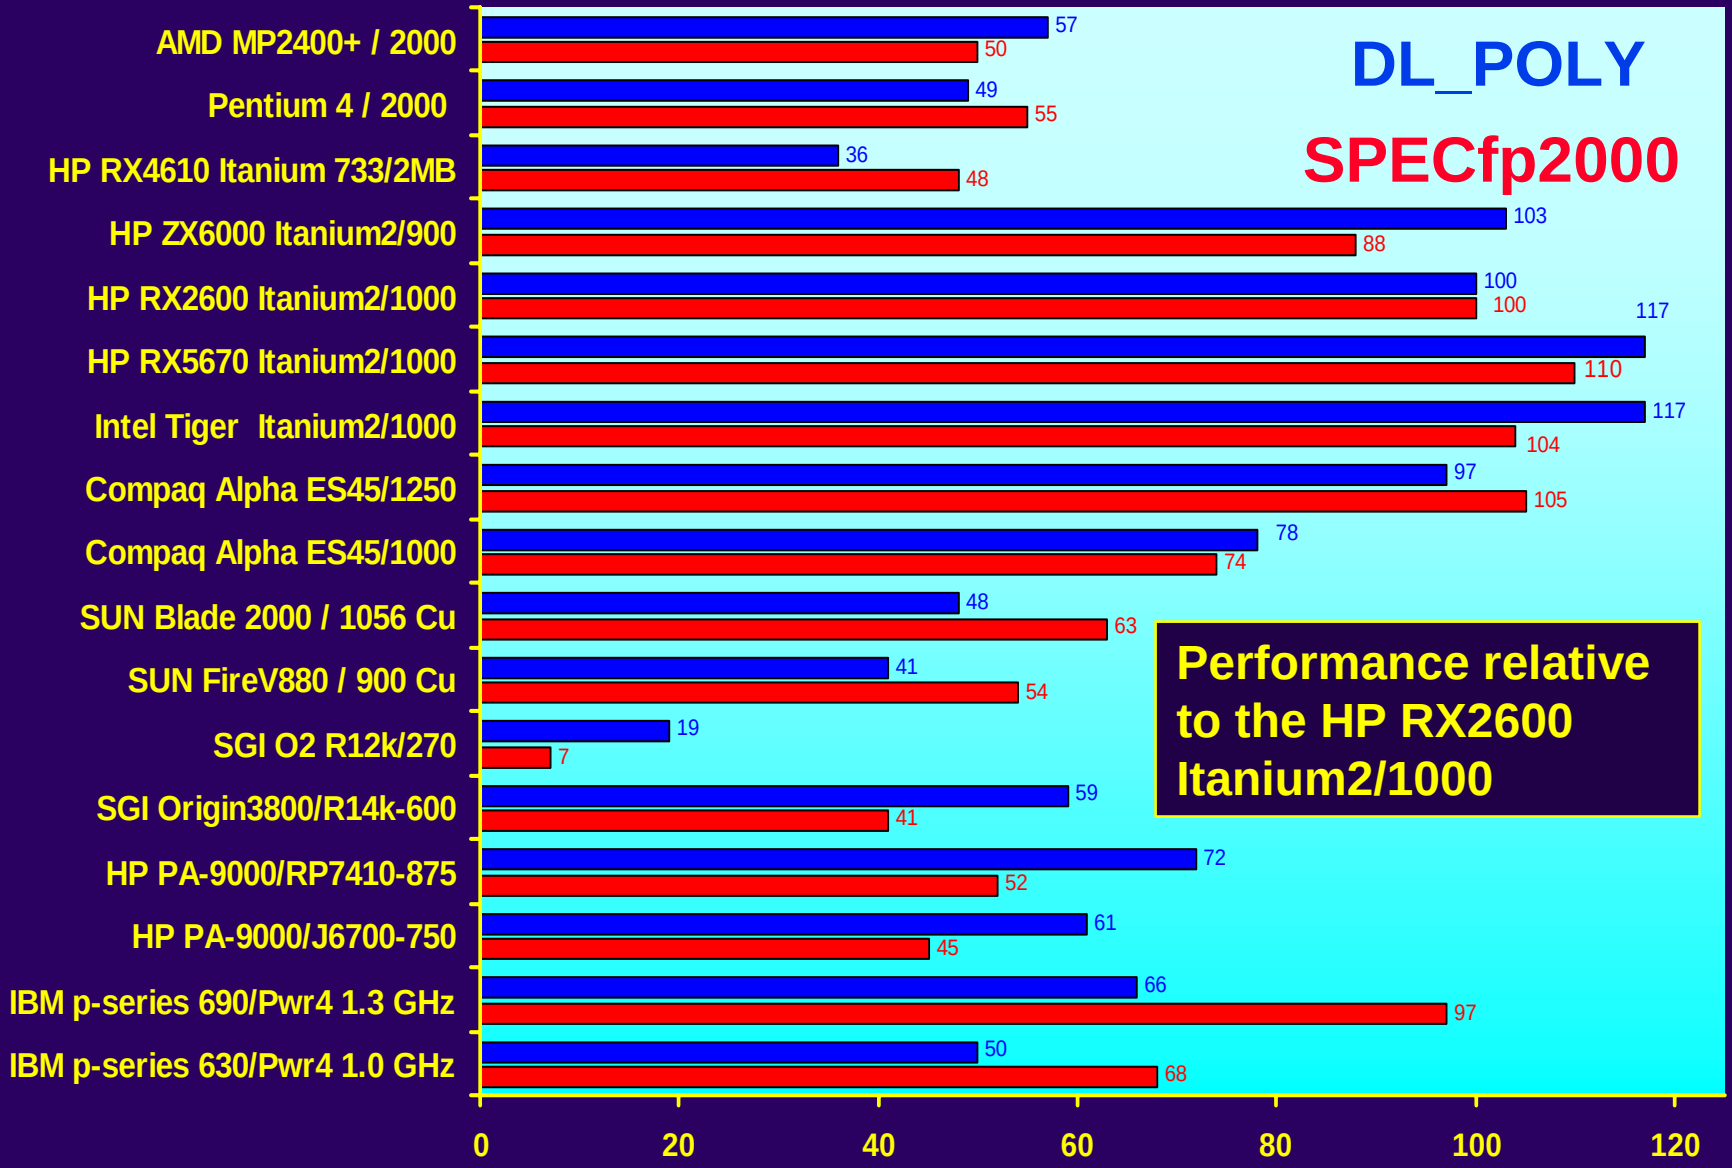

Summary Index relative to the HP RX2600 Itanium2/1000

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

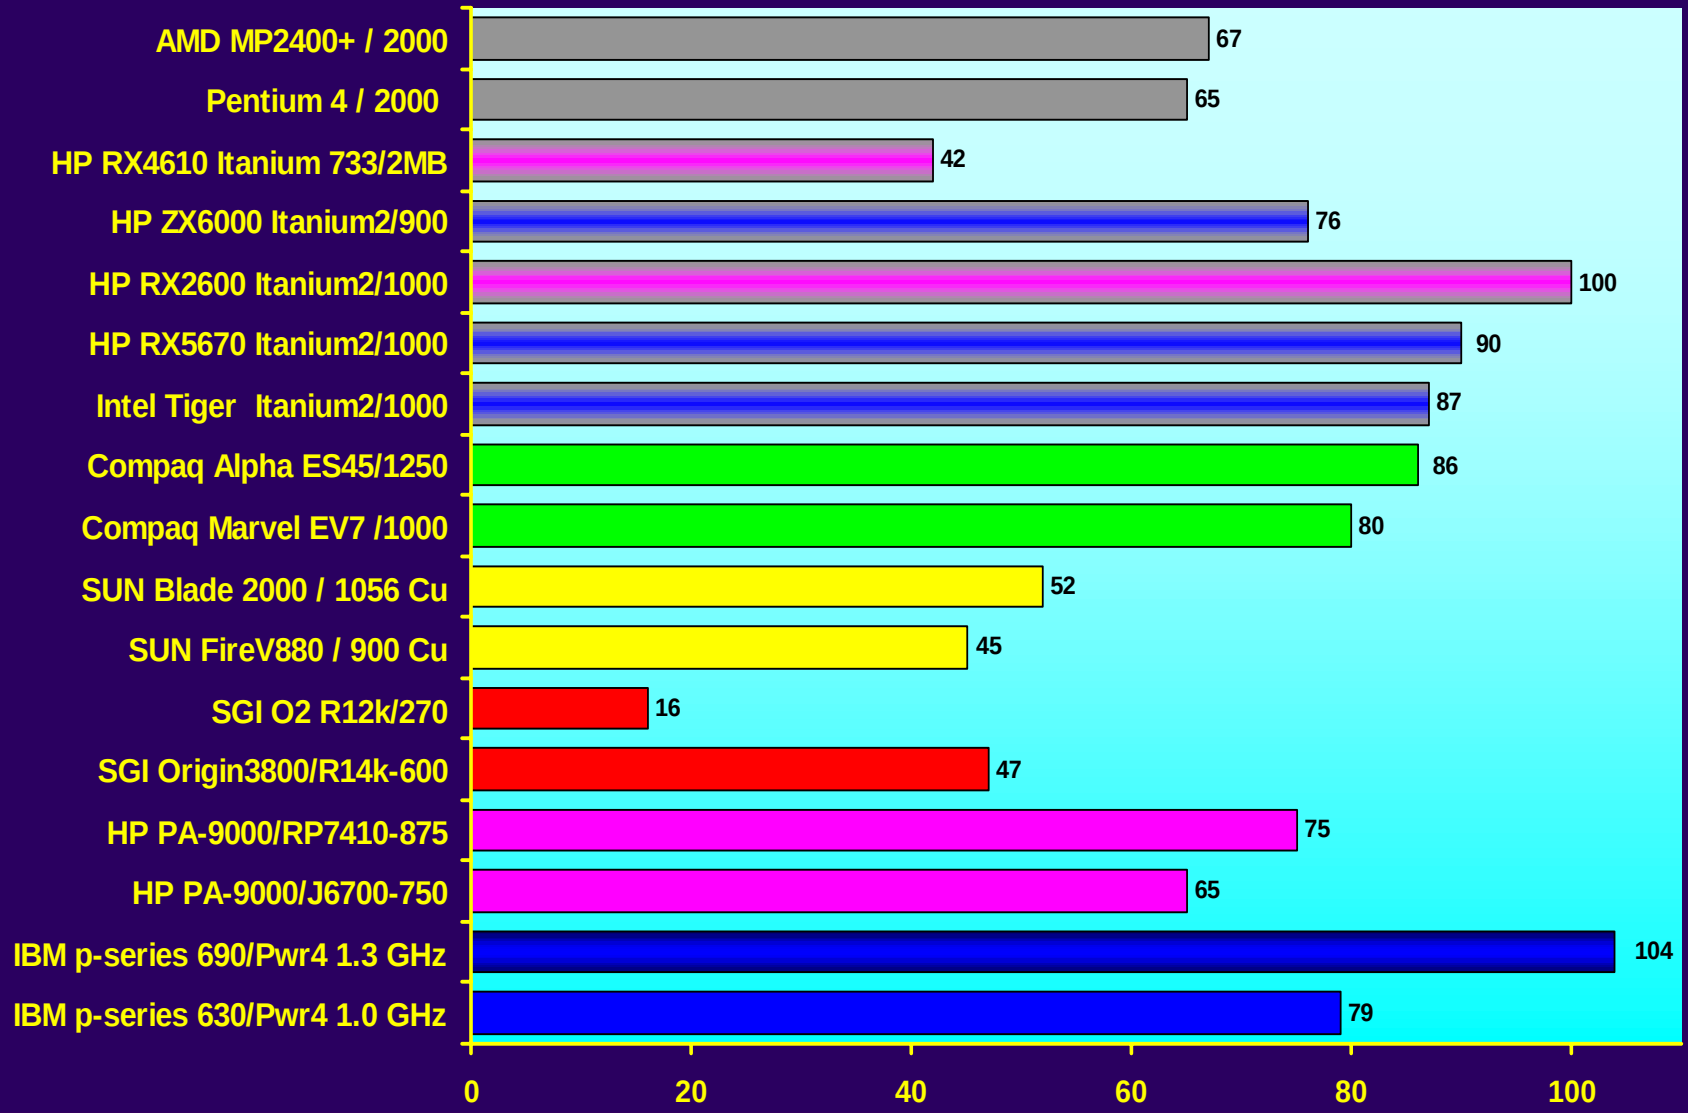

Computational Chemistry Rate Benchmark

- "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* X T_ref / T_i
  - T_ref is the elapsed time on a reference system. (*HP RX/2600 Itanium2/1000* 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 X B) are performed involving matrices of increasing order:
    - MATRIX-97+: 100, 200, … , 1200 (B is sparse)

Memory bandwidth

- Diagonalisation Benchmark
  - performance of 7 routines from mathematical libraries and QC codes:
    - MATRIX-97+: 100, 200, …, 600
- $Q^{\dagger}HQ$ Benchmark
  - involves use of library routines e.g. BLAS. Uses both a scalar and vector algorithm (DGEMM):
    - MATRIX-97+: 100, 200, … , 1000

-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^{nd} derivs., MP2 2^{nd} derivs., Direct-MP2)

Rate Benchmark: DL_POLY Component

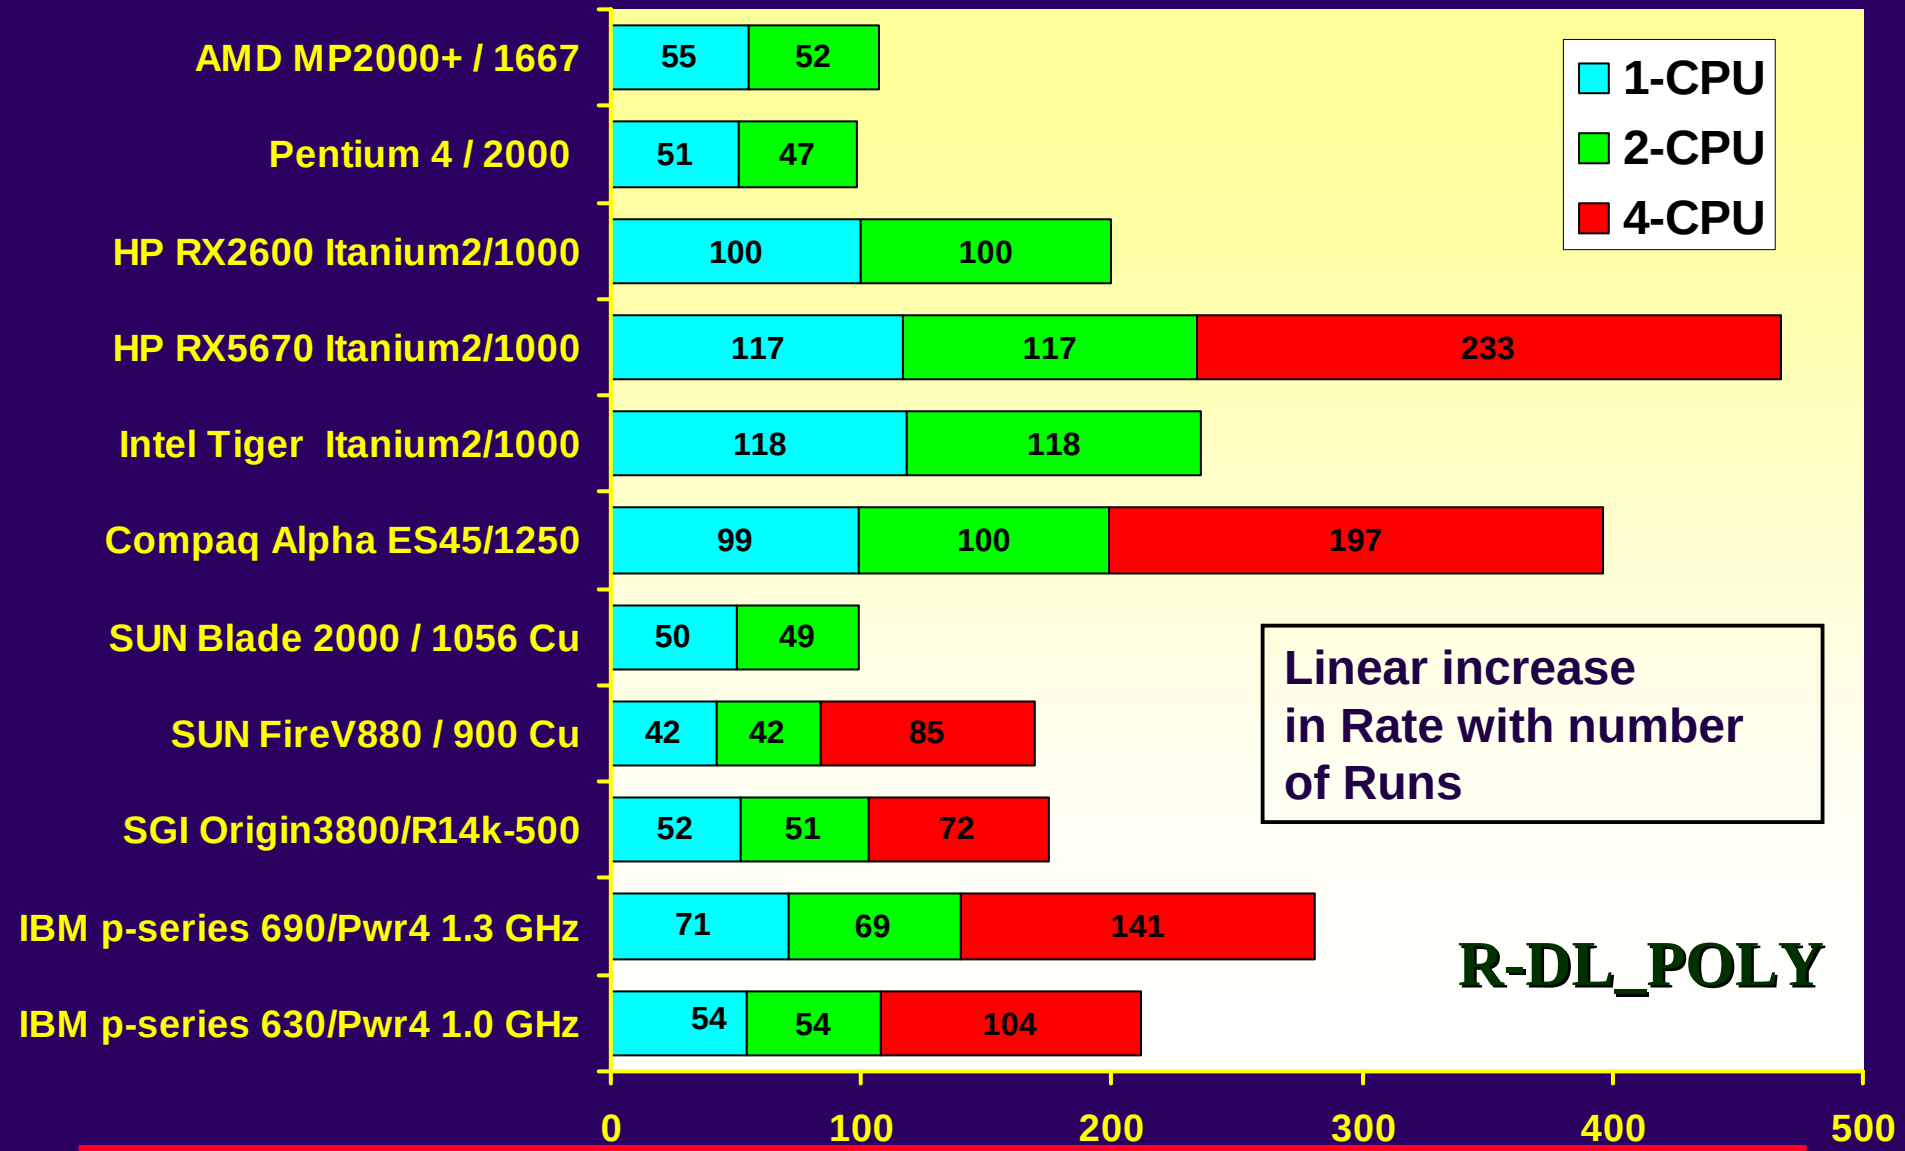

Computational Chemistry Rate Metric

- 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* X T_ref / T_i
  - T_ref is the elapsed time on the HP RX2600 / Itanium 2 scaled to a single CPU time of 100 units.

R_*CHEM* = *0.1* X R_***DGEMM*** +

0.1 X R_***Diagonalisation*** +

0.4 X R_*DLPOLY* +

0.4 X R_*GAMESS-UK*

Chemistry Rate Benchmark

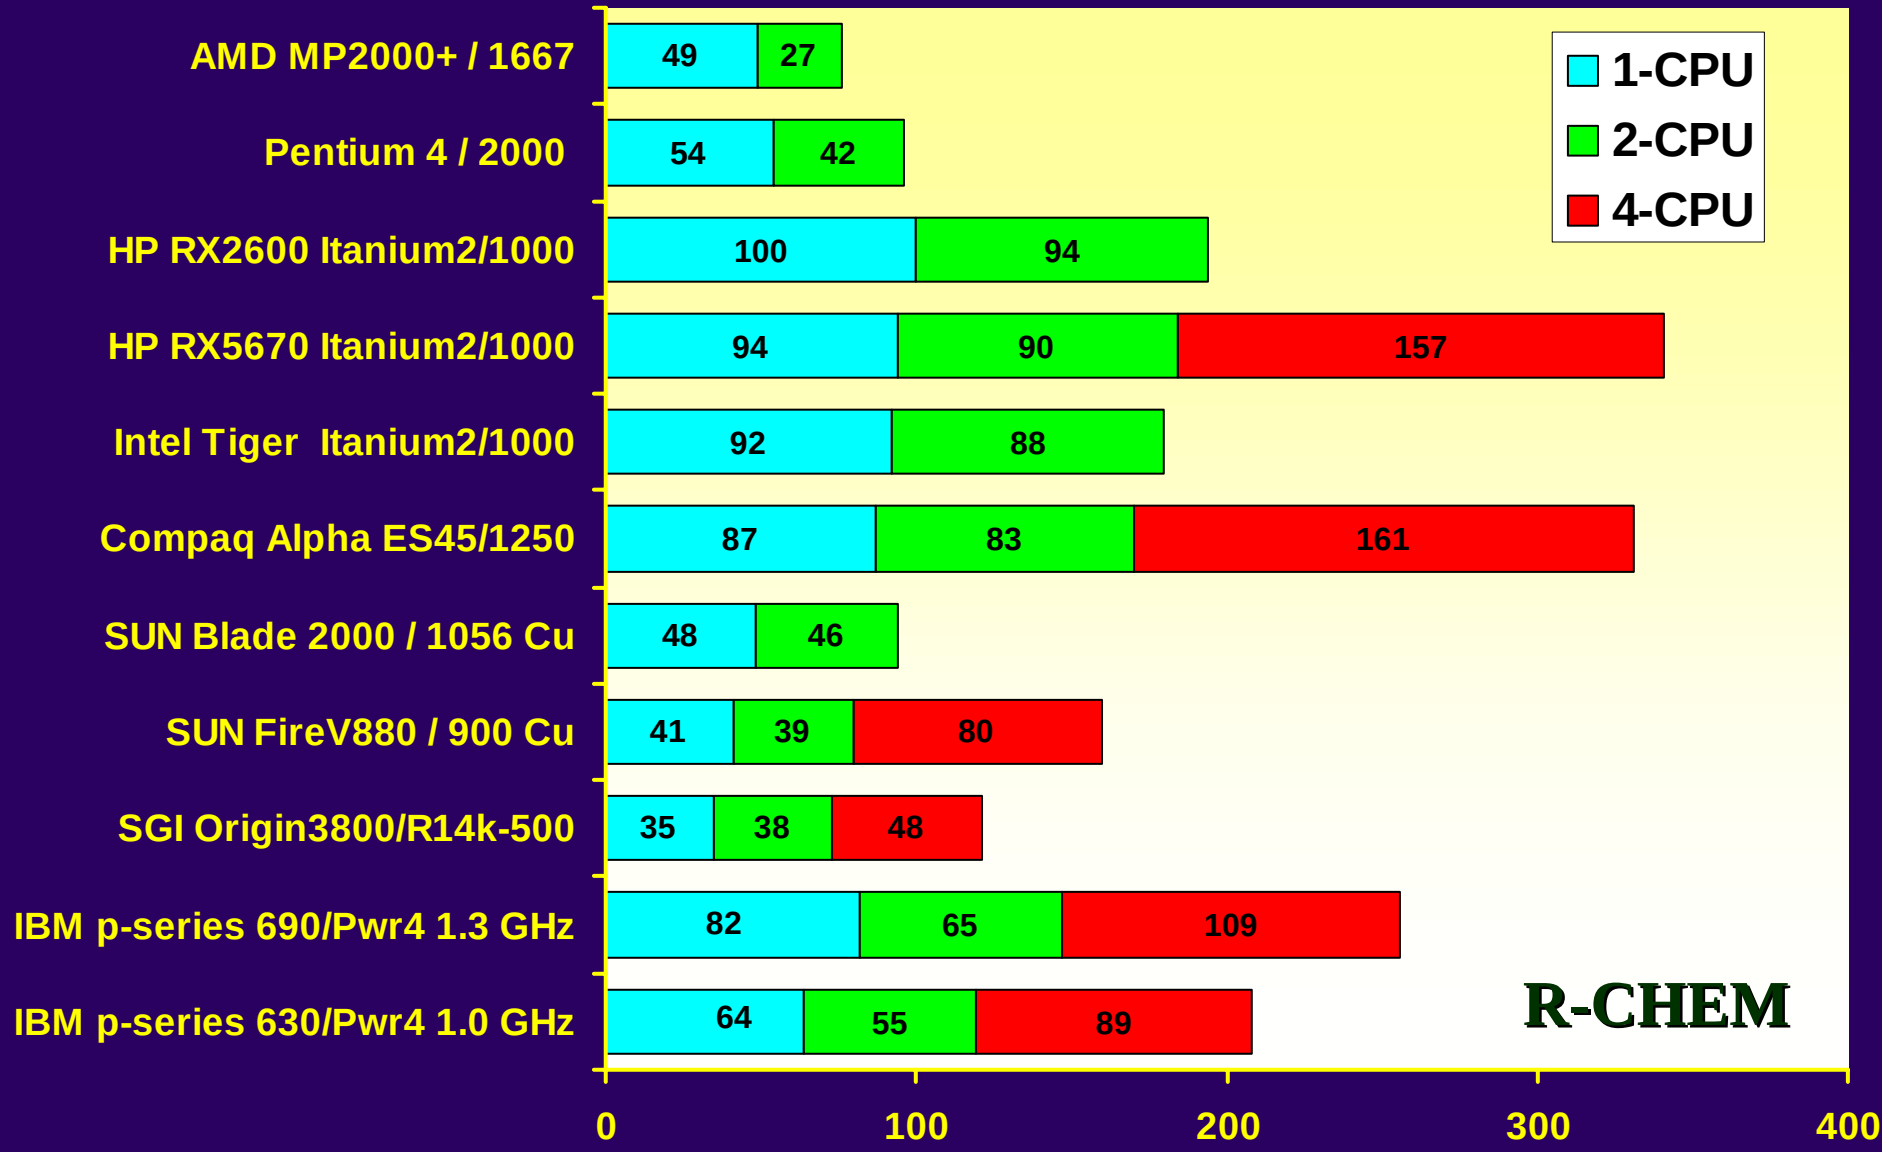

SUMMARY

- Processor Performance Overview
  - Single-processor performance, Performance Metrics
- SPECfp and Computational Chemistry Benchmark (serial)
  - Comparison involves 200+ computers (supercomputers, workstations, PCs and MPP nodes)
  - 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