Performance of Parallel Chemistry Codes on Linux Clusters

Martyn F. Guest
CCLRC Daresbury Laboratory

m.f.guest@daresbury.ac.uk

http://www.dl.ac.uk/CFS/parallel/chemistry

Outline

- Computational Chemistry - Background
  - Molecular Electronic Structure & Molecular Simulation
- Commodity-based and High-end Systems
  - Prototype Commodity Systems; CS1 - CS7
  - High-end systems from Cray, SGI, IBM and Compaq
  - Performance Metrics
- Application performance
  - Molecular Simulation
    - DLPOLY and CHARMM
  - Electronic Structure
    - Distributed data: Global Arrays (GAs) ; Linear Algebra (PeIGS)
    - NWChem, GAMESS-UK, and TurboMole
- Application performance analysis
  - VAMPIR and instrumenting the GA Tools
- Summary

Scaling of Molecular Computations

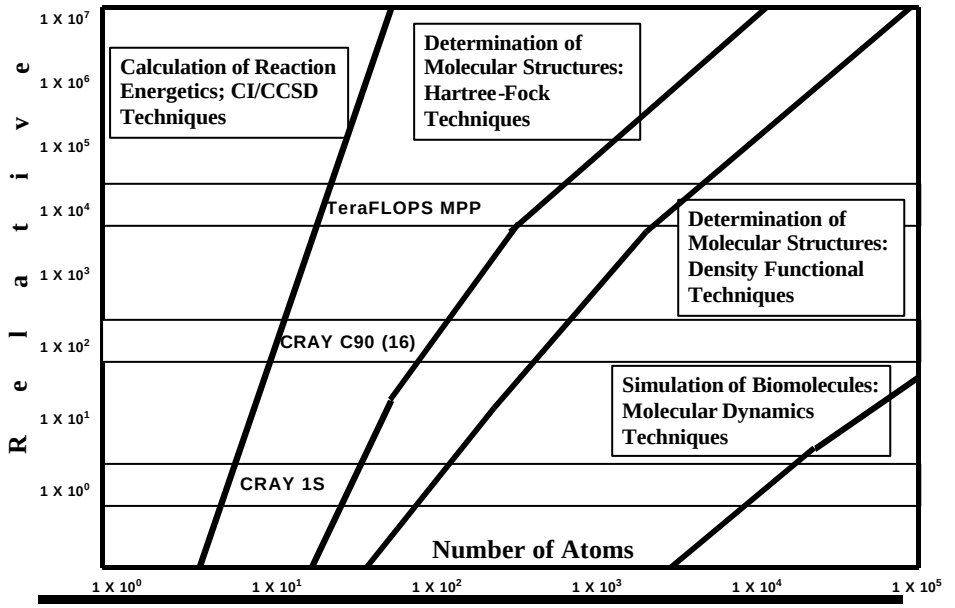

Electronic structure of molecules. I

- What can we predict?
  - In principle everything, in practice ….
  - Very small molecules (4 atoms?) - full chemical reaction dynamics
  - Small molecules (10 atoms?) - spectra, energetics and structure at a level of precision rivaling experiment
  - Larger molecules - lower precision - higher precision for trends
  - Application areas - drug design, petrochemicals, catalysis, waste processing,…
- One-electron (Hartree-Fock) model
  - Tremendous interpretative value
  - Predicts equilibrium structure of many molecules
  - 99% of total molecular energy
  - … but the missing 1% includes much information about chemical binding
- Electron correlation
  - Fluctuations arising from “instantaneous” interactions of electrons
  - Near degeneracy or resonances
  - Describe these effects by either
    - explicit 2-electron functions - not viable
    - including other occupation patterns or configurations in the wavefunction

Electronic structure of molecules .II

- Basis sets - the algebraic approximation
  - Subsequent theory expressed as (mostly dense) matrix operations
  - Gaussian basis - integrals tractable, no loss of precision in practice
  - The root of nearly all our scaling woes but presently indispensable
- Density functional theory
  - Exact energy is a function(al) of the 1-electron density - a 3-D entity
  - But … don't know the functional - educated guesses; nothing systematic
  - Scales from O(N) to $O(N^4)$
- Configuration Interaction
  - A linear expansion mixing important occupations - sparse eigenproblem
  - Slowly convergent - large expansions are necessary; state of art is 10^9
  - Scales from $O(N^6)$ to O(N!)
- Many-body methods - perturbation and coupled-cluster theory
  - A non-linear expansion mixing important orbital occupations - size extensive
  - Scales from $O(N^5)$ to $O(N^7)$

Beowulf Chemistry Sites

- http://www.wulfpack.med.jhmi.edu (Grossfield)
  - CHARMM
- http://www.lobos.nih.gov (Brookes et al., Pentiums)
  - CHARMM, GAMESS
- http://www.hpti.com/clusterweb (Lonergan, 34 XP1000s + Myrinet)
  - CHARMM and GAUSSIAN
- http://www.soton.ac.uk/~chemphys/jessex/beowulf.html
  - MC and MD simulations.
- http://www.ccr.buffalo.edu (Furlani, 64 Ultra 5 CPUs)
  - GAMESS and CRYSTAL 95
- http://www.dhpc.adelaide.edu.au/projects/beowulf/perseus.html
- Burger, Zurich (16 PII Cluster)
  - Turbomole, DMOL, ADF
- http://www.t12.lanl.gov/~mchallacombe (26 CPU PII cluster)
  - variety of QC software
- http://zinc10.chem.ucalgary.ca (94 Alpha EV56 Compaq PW 500 au)
  - ADF and PAW

Commodity Systems (CSx)
Prototype / Evaluation Hardware

Systems	Location	CPUs	Configuration
CS1	Daresbury	32	Pentium III / 450 MHz; fast ethernet (EPSRC)
CS2	Daresbury	64	24 X dual UP2000/EV67-667, QSNet Alpha/LINUX cluster, 8 X dual CS20/EV67-833
CS3	RAL	16	Athlon AMD K7 850MHz; myrinet interconnect
CS4	Sara	32	Athlon AMD K7 1.2 GHz; fast ethernet
CS6	CLiC	528	Pentium III / 800 MHz; fast ethernet (Chemnitzer Cluster)
CS7	Daresbury	64	AMD K7/1000 MP; SCALI SCI interconnect
Protoype Systems			
CS0	Daresbury	10	10 CPUS, Pentium II/266
CS5	Daresbury	16	8 X dual Pentium III/933, SCALI

High-End Systems

- Cray T3E/1200E
  - 816 processor system at Manchester (CSAR service)
  - 600 Mz EV56 Alpha processor with 256 MB memory
- IBM SP/WH2-375 (32 CPU system at DL)
  - 4-way Winterhawk2 SMP “thin nodes” with 2 GB memory
  - 375 MHz Power3-II processors with 8 MB L2 cache
- Compaq AlphaServer SC - 667 (APAC) and 833 MHz CPUs
  - 4-way ES40/667 and /833 SMP nodes with 2 GB memory
  - Alpha 21264a (EV67) CPUs with 8 MB L2 cache
  - Quadrics “fat tree” interconnect (5 usec latency, 150 MB/sec B/W)
- SGI Origin 3800
  - SARA (1000 CPUs) - Numalink with R14k/500 & R12k/400 CPUs
- Cray Supercluster at Eagen
  - Linux Alpha Cluster (96 X API CS20s - dual 833 MHz EV67 CPUs)
  - Myrinet interconnect, Red Hat 6.2

Chemistry Codes

- Performance comparisons between Commodity-based systems and proprietary High-end systems
  - current MPP (CSAR Cray T3E/1200E) and ASCI-style SMP-node platforms (IBM SP / WH2-375, Compaq AlphaServer SC (ES40/6-667, 6-833), SGI Origin 3800 and Prototype Cray Supercluster:
- Molecular Simulation
  - DL_POLY - parallel MD code with many applications
  - CHARMM - macromolecular MD and energy minimisation
- Ab initio Electronic Structure
  - GAMESS-UK, NWChem and Turbomole

Performance Metric (% 32-node high-end system)

Performance Metrics

Attempt 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)

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

$$\frac{[\,T_{32\text{-node}}\ \textit{T3E} \,/\, T_{32\text{-node}}\ \textit{CS1 Pentium III/450 + FE}\,]}{T_{32\text{-node}}\ \textbf{T3E} \,/\, T_{32\text{-node}}\ \textbf{CS6 Pentium III/800 + FE}}$$

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

2. T (32-CPUs SGI Origin 3800) / T (32 CPUs) CS2 Alpha Linux Cluster

$$T_{32\text{-CPU}}\ \textbf{SGI Origin 3800} \,/\, T_{32\text{-CPU}}\ \textbf{CS2 Alpha Linux Cluster}$$

The GAMESS-UK Serial Benchmark

Performance relative to the Compaq Alpha ES40/833

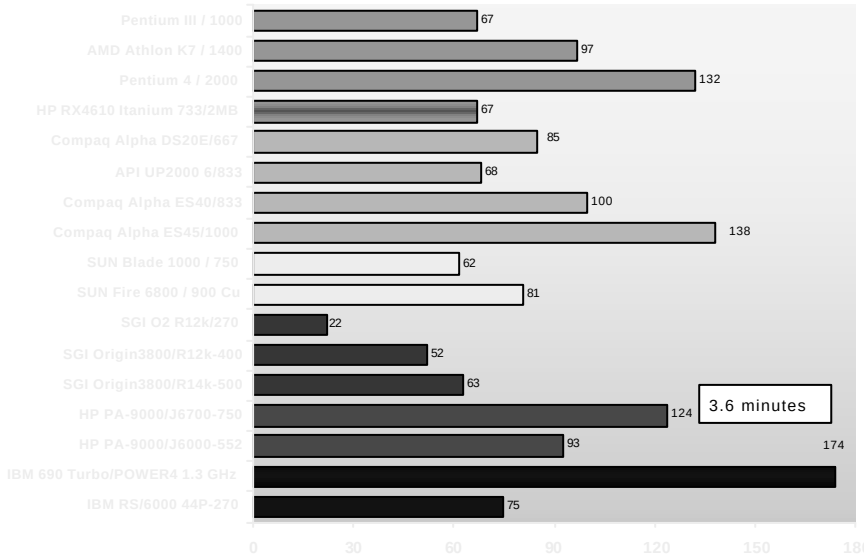

STREAM: Measured Sustainable Memory Bandwidth in HPC (TRIAD)

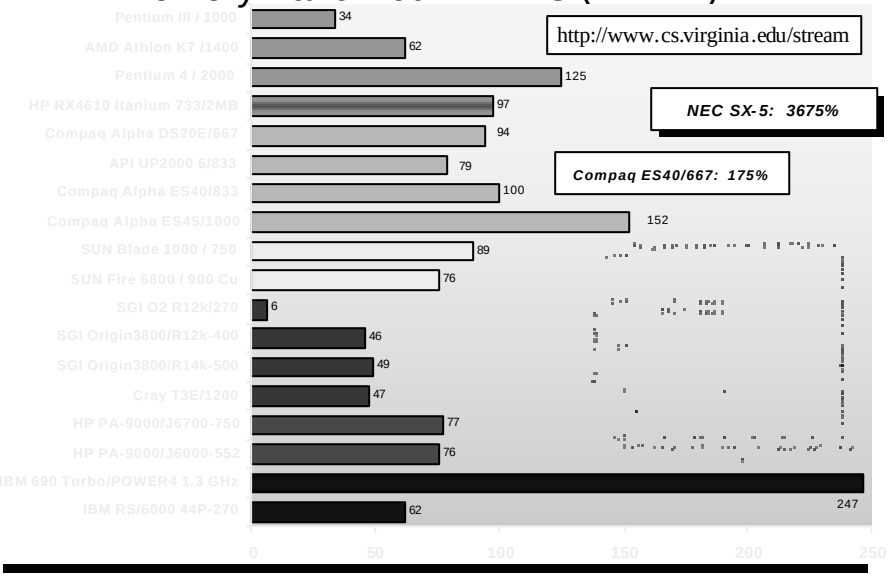

Molecular Simulation

Molecular Dynamics Codes: DL_POLY and CHARMM

DL_POLY: A Parallel Molecular Dynamics Simulation Package

- Developed as CCP5 parallel MD code by W. Smith and T.R. Forester
- UK + International user community
- Adopted by Materials Consortium 1995

Boundary Conditions

- None (e.g. isolated macromolecules)
- Cubic periodic boundaries
- Orthorhombic periodic boundaries
- Parallelpiped periodic boundaries
- Truncated octahedral periodic boundaries
- Rhombic dodecahedral periodic boundaries
- Slabs (i.e. x,y periodic, z nonperiodic)

Target Systems

- Atomic systems & mixtures (Ne, Ar, etc.)
- Ionic melts & crystals (NaCl, KCl etc.)
- Polarisable ionics (ZSM-5, MgO etc.)
- Molecular liquids & solids (CCl_4, Bz etc.)
- Molecular ionics (KNO_3, NH_4Cl, H_2O etc.)
- Synthetic polymers ($[PhCHCH_2]_n$etc.)
- Biopolymers and macromolecules
- Polymer electrolytes, Membranes,
- Aqueous solutions, Metals

MD Algorithms/Ensembles

- Verlet leapfrog, Verlet leapfrog + RD-SHAKE
- Rigid units with FIQA and RD-SHAKE
- Linked rigid units with QSHAKE
- Constant T (Berendsen) with Verlet leapfrog and with RD-SHAKE
- Constant T (Evans) with Verlet leapfrog and with RD-SHAKE
- Constant T (Hoover) with Verlet leapfrog

DL_POLY Parallel Benchmarks (Cray T3E/1200)

4. NaCl; MTS Ewald, 27,000 ions
5. NaK-disilicate glass; 8,640 atoms, Ewald
8. MgO microcrystal: 5,416 atoms

9. Model membrane/Valinomycin (MTS, 18,886)
7. Gramicidin in water (SHAKE, 13,390)
6. K/valinomycin in water (SHAKE, AMBER, 3,838)
1. Metallic Al (19,652 atoms, Sutton Chen)
3. Transferrin in Water (neutral groups + SHAKE, 27,593)
2. Peptide in water (neutral groups + SHAKE, 3993).

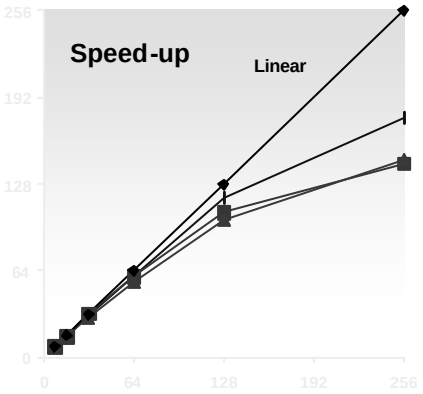

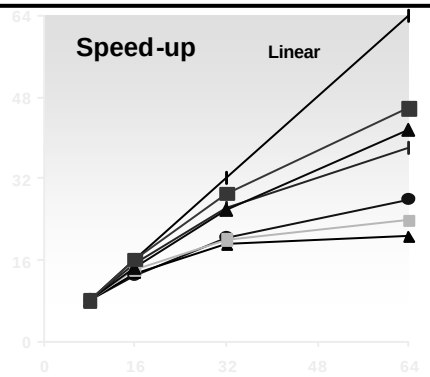

DL_POLY: Cray/T3E, High-end and Commodity-based Systems I.

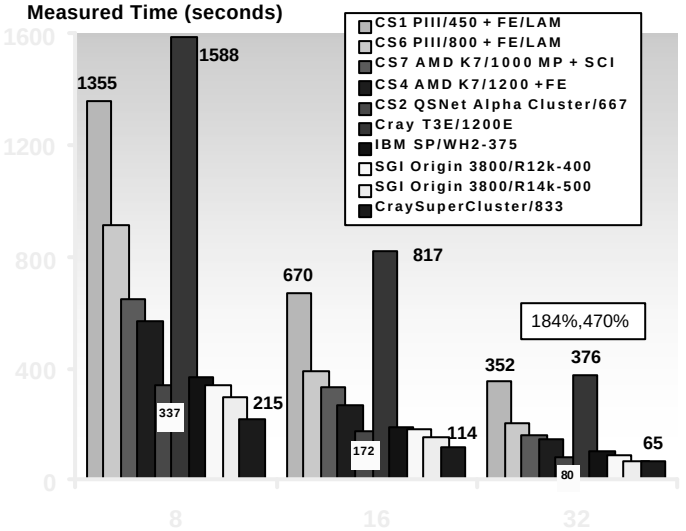

Bench 4.
NaCl; 27,000 ions,
Ewald + MTS,
75 time steps,
Cutoff=24Å

$T3E_{128}$ =94

DL_POLY: Cray/T3E, High-end and Commodity-based Systems II.

Measured Time (seconds)

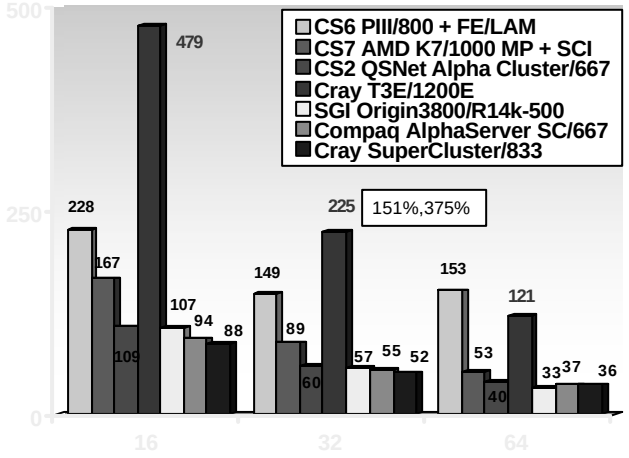

Bench 5.
NaK-disilicate glass; 8,640 atoms, Ewald: 270 time steps

$T3E_{128} = 75$

Number of CPUs

DL_POLY: Cray/T3E, High-end and Commodity-based Systems

$T3E_{128} = 75$

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

Measured Time (seconds)

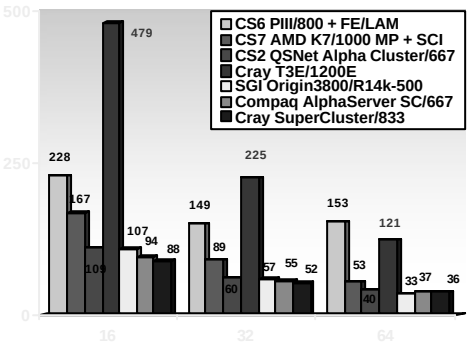

Number of CPUs

Speed-up

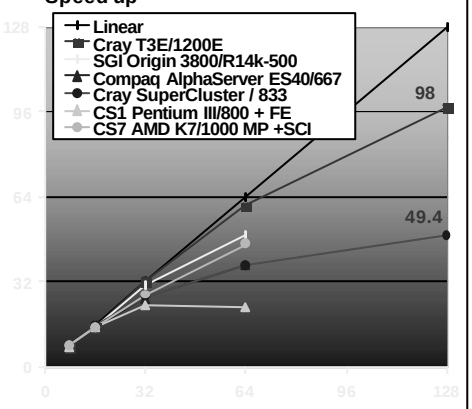

Number of CPUs

DL_POLY: Cray/T3E, High-end and Commodity-based Systems III.

Measured Time (seconds)

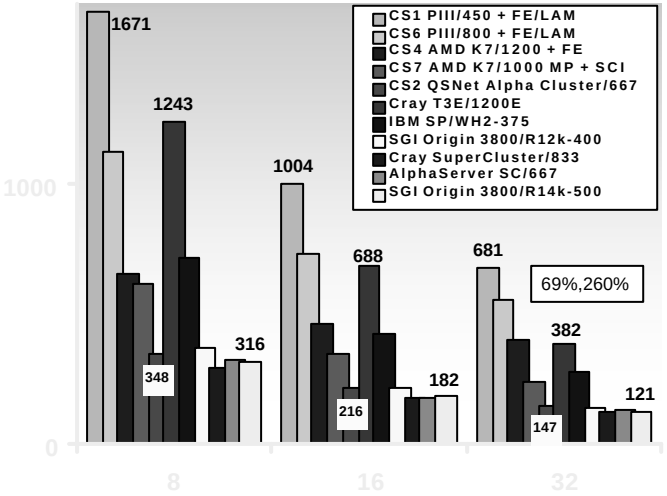

Bench 7.
Gramicidin in water;
rigid bonds and SHAKE,
13,390 atoms,
500 time steps

$T3E_{128}$ =166

Number of CPUs

Vampir

CS-1: 8 node DLPOLY trace output (single timestep)

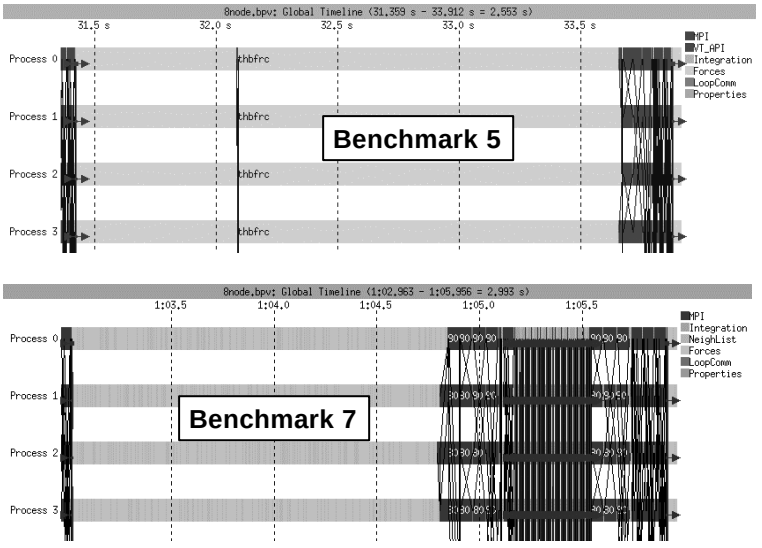

CHARMM

- CHARMM (Chemistry at HARvard Macromolecular Mechanics) is a general purpose molecular mechanics, molecular dynamics and vibrational analysis package for modelling and simulation of the structure and behaviour of macromolecular systems (proteins, nucleic acids, lipids etc.)
- Supports energy minimisation and MD approaches using a classical parameterised force field.
- J. Comp. Chem. 4 (1983) 187-217
- Parallel Benchmark - MD Calculation of Carboxy Myoglobin (MbCO) with 3830 Water Molecules.
- QM/MM model for study of reacting species
  - incorporate the QM energy as part of the system into the force field
  - coupling between GAMESS-UK (QM) and CHARMM.

Parallel CHARMM Benchmark

Benchmark MD Calculation of Carboxy Myoglobin (MbCO) with 3830 Water Molecules: 14026 atoms, 1000 steps (1ps), 12-14 A shift.

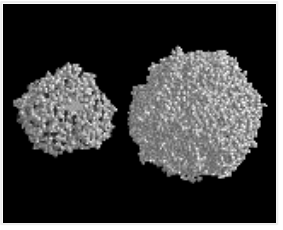

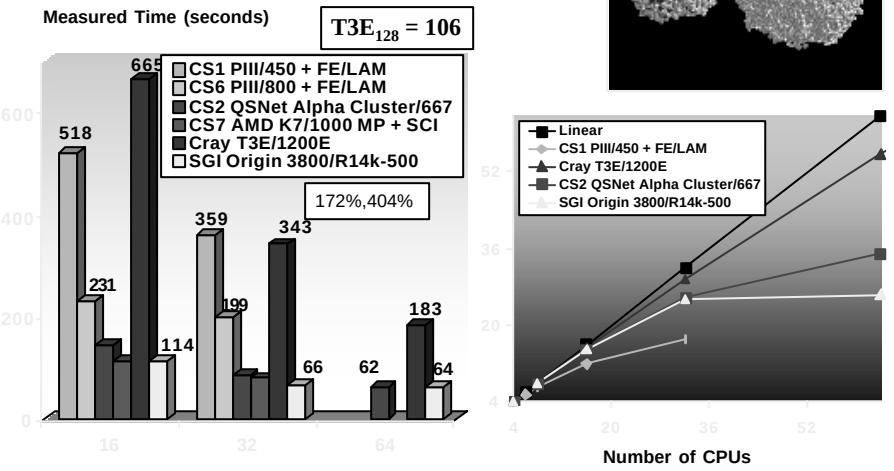

Parallel CHARMM Benchmark:

LAM MPI vs. MPICH

Benchmark MD Calculation of Carboxy Myoglobin (MbCO) with 3830 Water Molecules:

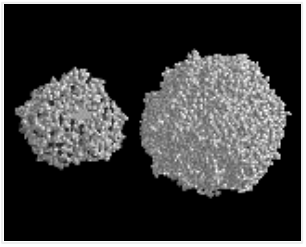

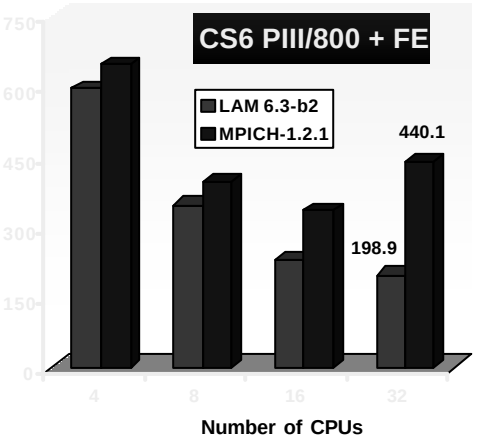

Molecular Electronic Structure

Ab initio Electronic structure Codes:
NWChem, GAMESS-UK and
Turbomole

Distributed Data SCF

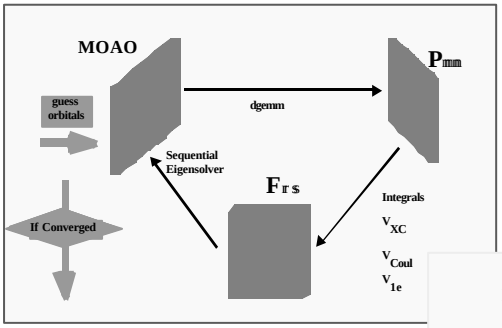

Pictorial representation of the iterative SCF process in (i) a sequential process, and (ii) a distributed data parallel process: **MOAO** represents the molecular orbitals, **P** the density matrix and **F** the Fock or Hamiltonian matrix

Sequential

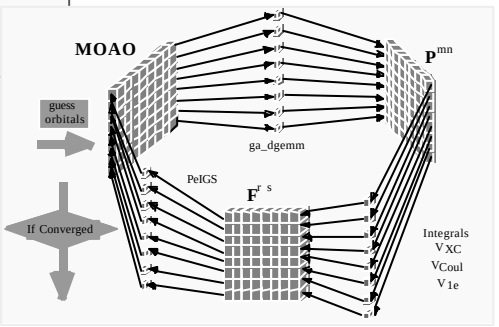

Distributed Data

High-End Computational Chemistry The NWChem Software

- Developed as part of the construction of the Environmental Molecular Sciences Laboratory (EMSL) at PNNL.
- Funded to be used as an integrated component in solving DOE's grand challenge environmental restoration problems
- Designed and developed to be a highly efficient and portable MPP computational chemistry package, providing computational chemistry solutions which are scalable with respect to chemical system size as well as MPP hardware size
- Extensible framework supporting development of new methods in computational chemistry; NWChem Architecture
  - Object-oriented design
    - abstraction, data hiding, handles, APIs
  - Parallel programming model
    - non-uniform memory access, global arrays (GAs)
  - Infrastructure
    - **Global Arrays (GA)** , Parallel I/O, RTDB, MA, **Linear algebra (PeiGS)** ...

High-End Computational Chemistry
The NWChem Software

- Capabilities (Direct, Semi-direct and conventional):
  - RHF, UHF, ROHF using up to 10,000 basis functions; analytic 1st and 2nd derivatives.
  - DFT with a wide variety of local and non-local XC potentials, using up to 10,000 basis functions; analytic 1st and 2nd derivatives.
  - CASSCF; analytic 1st and numerical 2nd derivatives.
  - Semi-direct and RI-based MP2 calculations for RHF and UHF wave functions using up to 3,000 basis functions; analytic 1st derivatives and numerical 2nd derivatives.
  - Coupled cluster, CCSD and CCSD(T) using up to 3,000 basis functions; numerical 1st and 2nd derivatives of the CC energy.
  - Classical molecular dynamics and free energy simulations with the forces obtainable from a variety of sources

Measured Parallel Efficiency for NWChem - DFT on IBM-SP; Wall Times to Solution for SCF Convergence

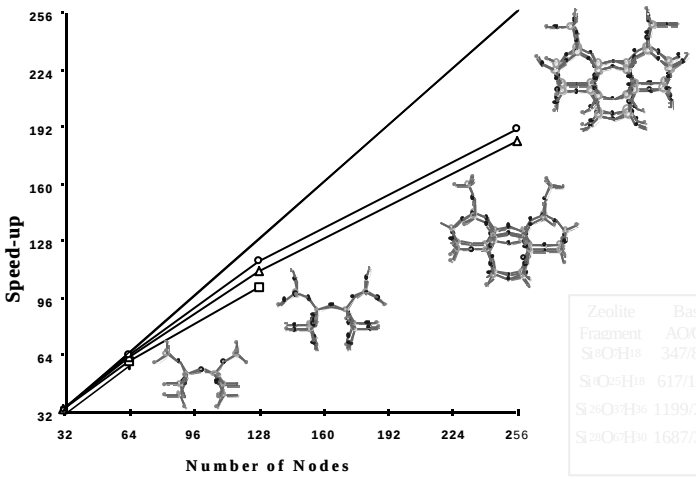

Zeolite Fragment	Basis AO/CD	Number of Nodes	Wall Time to Solution
$Si_8O_7H_{18}$	347/832	64	238s
$Si_{17}O_{25}H_{18}$	617/1444	128	364s
$Si_{26}O_{37}H_{36}$	1199/2818	256	1137s
$Si_{28}O_{67}H_{30}$	1687/3928	256	2766s

Global Arrays

Physically distributed data

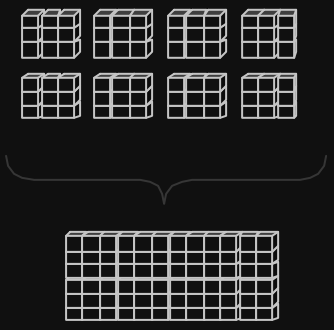

Single, shared data structure

- Shared-memory-like model
  - Fast local access
  - NUMA aware and easy to use
  - MIMD and data-parallel modes
  - Inter-operates with MPI, ...
- BLAS and linear algebra interface
- Ported to major parallel machines
  - IBM, Cray, SGI, clusters,...
- Originated in an HPCC project
- Used by 5 major chemistry codes, financial futures forecasting, astrophysics, computer graphics

Tools developed as part of the NWChem project at PNNL; R.J. Harrison, J. Nieplocha et al.

PeIGS 3.0 Parallel Performance

(Solution of real symmetric generalized and standard eigensystem problems)

Developed as part of the NWChem project at PNNL; R.J. Harrison, J. Nieplocha et al.

Features (not available elsewhere):

- **Inverse iteration using Dhillon-Fann-Parlett's parallel algorithm (fastest uniprocessor performance and good parallel scaling)**
- **Guaranteed orthonormal eigenvectors in the presence of large clusters of degenerate eigenvalues**
- **Packed Storage**
- **Smaller scratch space requirements**

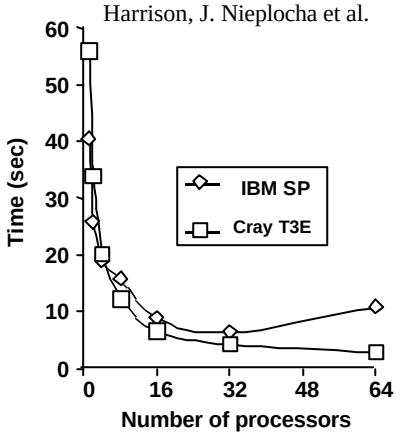

Full eigensolution performed on a matrix generated in a charge density fitting procedure (966 fitting functions for a fluorinated biphenyl).

Parallel Eigensolvers

Real symmetric eigenvalue problems

SGI Origin 3800/R12k-400 ("green")

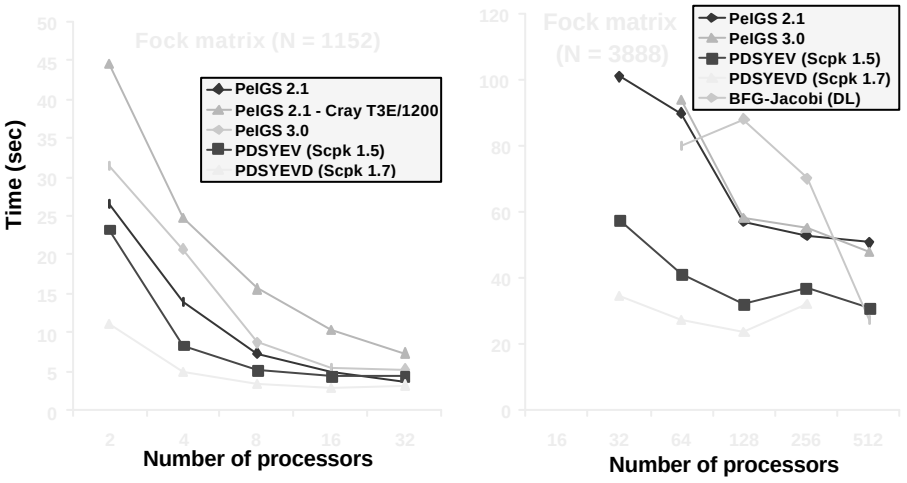

Parallel Eigensolvers

Real symmetric eigenvalue problems

CS7 AMD K7/1000 MP + SCALI

Scalapack (PDSYEV)

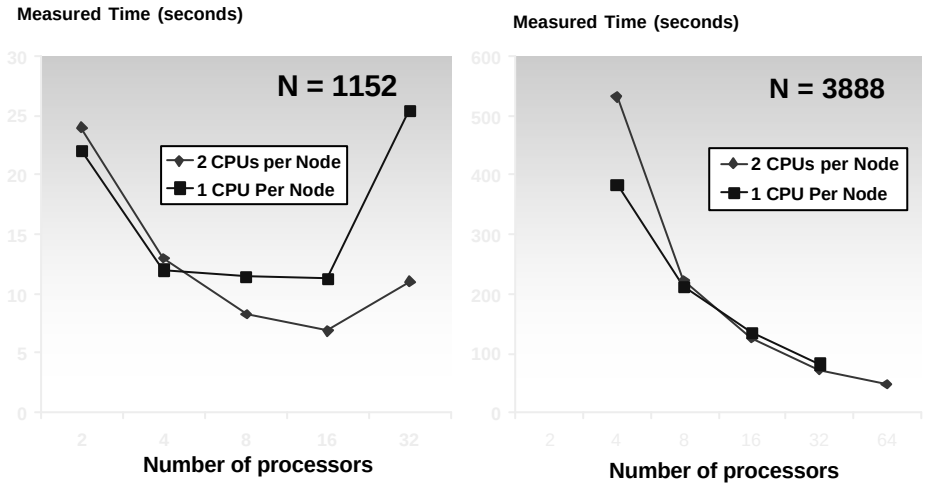

Case Studies - Zeolite Fragments

$Si_8O_7H_{18}$ **347/832**

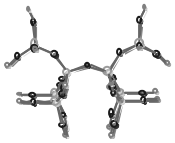

$Si_8O_{25}H_{18}$ **617/1444**

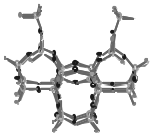

$Si_{26}O_{37}H_{36}$ **1199/2818**

$Si_{28}O_{67}H_{30}$ **1687/3928**

- DFT Calculations with Coulomb Fitting

Basis (Godbout et al.)

DZVP - O, Si
DZVP2 - H

Fitting Basis:

DGAUSS-A1 - O, Si
DGAUSS-A2 - H

- NWChem & GAMESS-UK

Both codes use auxiliary fitting basis for coulomb energy, with 3 centre 2 electron integrals held in core.

DFT Coulomb Fit - NWChem

$Si_8O_7H_{18}$ **347/832**

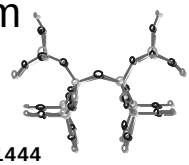

$Si_8O_{25}H_{18}$ **617/1444**

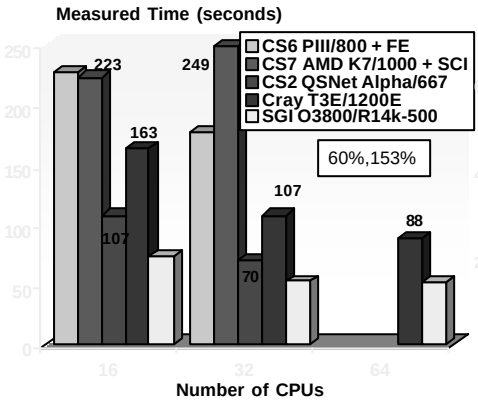

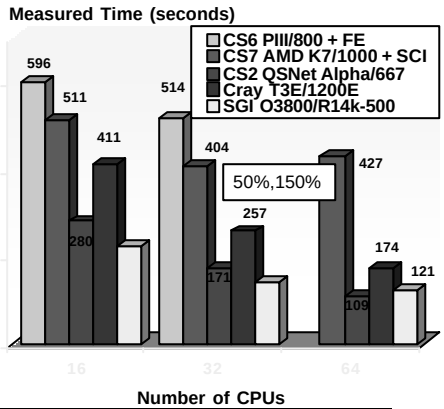

DFT Coulomb Fit - NWChem

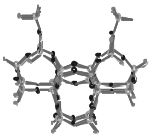

$Si_{26}O_{37}H_{36}$ 1199/2818 $Si_{28}O_{67}H_{30}$ 1687/3928

$T_{IBM\text{-}SP/P2SC\text{-}120}$ (256) = 1137

$T_{IBM\text{-}SP/P2SC\text{-}120}$ (256) = 2766

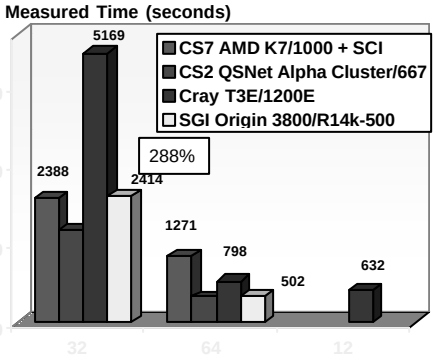

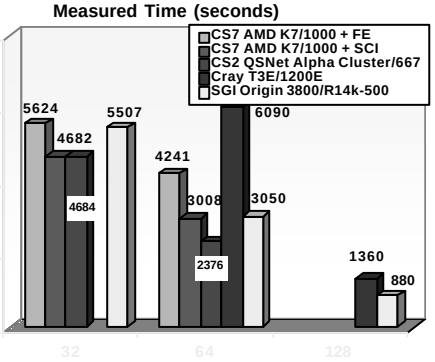

NWChem - DFT (LDA)
Performance on the SGI Origin 3800

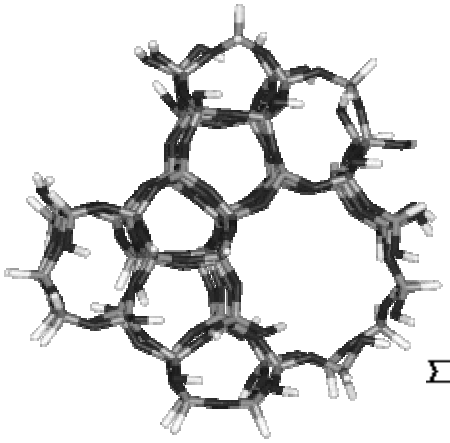

Zeolite ZSM-5

- DZVP Basis (DZV_A2) and Dgauss A1_DFT Fitting basis:

  AO basis: 3554
  CD basis: 12713

- MIPS R14k-500 CPUs (Teras)

Wall time (13 SCF iterations):
64 CPUs = 5,242 seconds
128 CPUs= 3,951 seconds

Est. time on 32 CPUs = 40,000 secs

- 3-centre 2e-integrals = 1.00 X 10^{12}
- Schwarz screening = 5.95 X 10^{10}
- % 3c 2e-ints. In core = 100%

GAMESS-UK

GAMESS-UK is the general purpose ab initio molecular electronic structure program for performing SCF-, MCSCF- and DFT-gradient calculations, together with a variety of techniques for post Hartree Fock calculations.

- The program is derived from the original GAMESS code, obtained from Michel Dupuis in 1981 (then at the NRCC), and has been extensively modified and enhanced over the past decade.
- This work has included contributions from numerous authors†, and has been conducted largely at the CCLRC Daresbury Laboratory, under the auspices of the UK's Collaborative Computational Project No. 1 (CCP1). Other major sources that have assisted in the on-going development and support of the program include various academic funding agencies in the Netherlands, and ICI plc.

Additional information on the code may be found from links at:

http://www.dl.ac.uk/CFS

† *M.F. Guest, J.H. van Lenthe, J. Kendrick, K. Schoffel & P. Sherwood, with contributions from R.D. Amos, R.J. Buenker, H.H. van Dam, M. Dupuis, N.C. Handy, I.H. Hillier, P.J. Knowles, V. Bonacic-Koutecky, W. von Niessen, R.J. Harrison, A.P. Rendell, V.R. Saunders, A.J. Stone and D.Tozer.*

Parallel Implementations of GAMESS-UK

- Extensive use of Global Array (GA) Tools and Parallel Linear Algebra from NWChem Project (EMSL)
- SCF and DFT energies and gradients
  - Replicated data, but …
  - GA Tools for caching of I/O for restart and checkpoint files
  - Storage of 3-centre 2-e integrals in DFT Jfit
  - Linear Algebra (via PeIGs, DIIS/MMOs, Inversion of 2c-2e matrix)
- SCF second derivatives
  - Distribution of <vvoo> and <vovo> integrals via GAs
- MP2 gradients
  - Distribution of <vvoo> and <vovo> integrals via GAs

GAMESS-UK ΔSCF Performance

Cray T3E/1200E, High-end and Commodity-based Systems

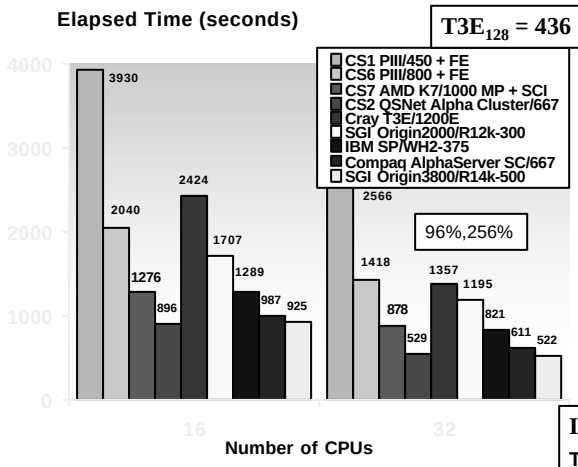

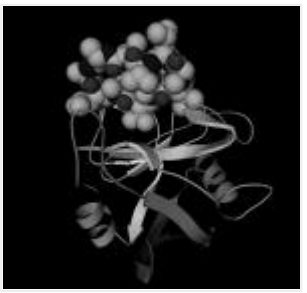

Impact of Serial Linear Algebra:

$T_{IBM-SP}(16) = 2656\ [1289]$

$T_{IBM-SP}(32) = 2184\ [\ 821]$

Cyclosporin:(3-21G Basis, 1000 GTOS)

ΔSCF Performance - Cray T3E/1200E , SGI Origin3800/R14k-500 and Compaq AlphaServer SC/667

Cyclosporin:(3-21G Basis, 1000 GTOS)

Gaussian 98
Serial: L302 - 90 secs;
L401 - 292 secs.
Serial linear algebra

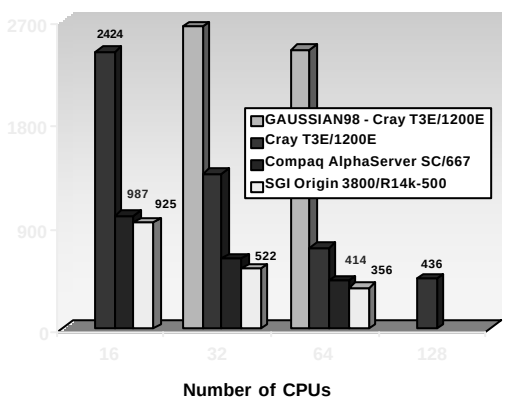

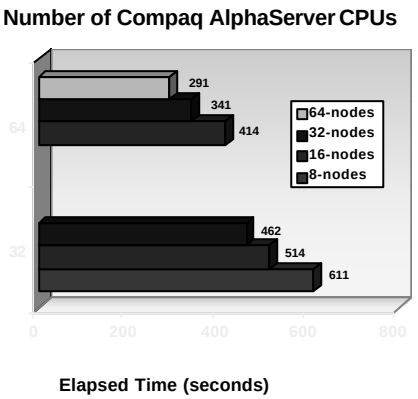

GAMESS-UK. DFT B3LYP Performance

Cray T3E/1200, High-end and Commodity-based Systems

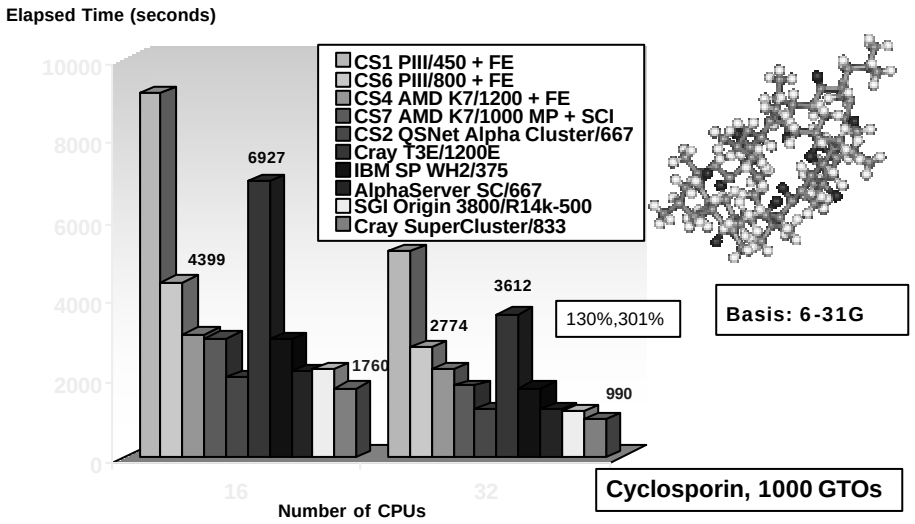

130%,301%

Basis: 6-31G

Cyclosporin, 1000 GTOs

GAMESS-UK. DFT B3LYP Performance

The Cray T3E/1200 and High-end Systems

Cyclosporin, 1000 GTOs

Basis: 6-31G

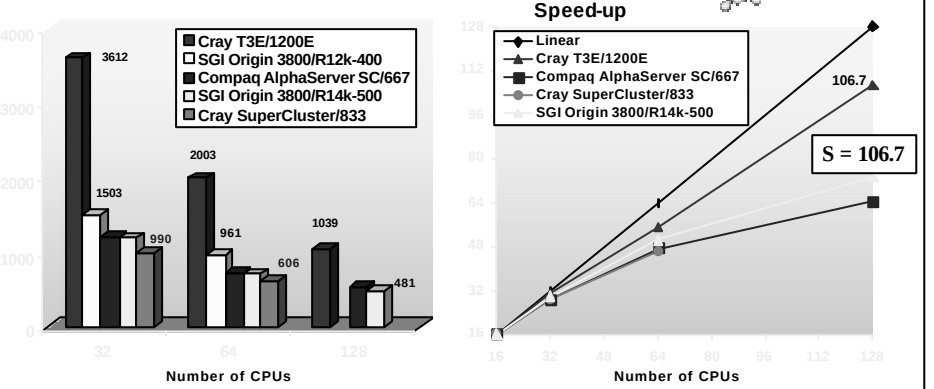

Auxilliary Basis Coulomb Fit (I)

The approach is based on the expansion of the charge density in an auxiliary basis of Gaussian functions

$$r(r) = \sum_{pq} D_{pq} |pq\rangle \approx \sum_{u} \left(\sum_{pq} D_{pq} C_u^{pq} \right) |u\rangle = \sum_{u} d_u |u\rangle$$

As suggested by Dunlap, a variational choice of the fitting coefficients C can be obtained as follows:

$$C^{pq} = V^{-1} b^{pq}$$

Where *V* is the matrix of 2-centre 2-electron repulsion integrals in the charge density basis and *b* are the three centre electron repulsion integrals between the wavefunction basis set and the charge density basis.

Auxilliary Basis Coulomb Fit (ii)

- The number of 3-centre integrals is significantly smaller than the 4-centre integrals used in the conventional coulomb evaluation, but for large molecules additional screening is required.
- We make use of the Schwarz inequality

$$(pq|u) \leq \sqrt{(pq|pq)}\sqrt{(u|u)}$$

- Where p and q are AO basis functions and u are the fitting functions. Since screening is applied on a shell basis, the maximal integrals for each shell quartet are stored.
- Using this screening, and exploiting the aggregate memory of a parallel machine, it is possible to hold a significant fraction of the 3-centre integrals in core.

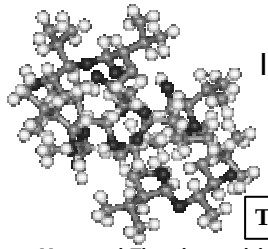

GAMESS-UK: DFT HCTH on Valinomycin.

Impact of Coulomb Fitting: Cray T3E/1200, High-end and Commodity Systems

Basis: DZV_A2 (Dgauss)
A1_DFT Fit: 882/3012

$T_{T3E/1200E}(128) = 2139$

$T_{T3E/1200E}(128) = 995$

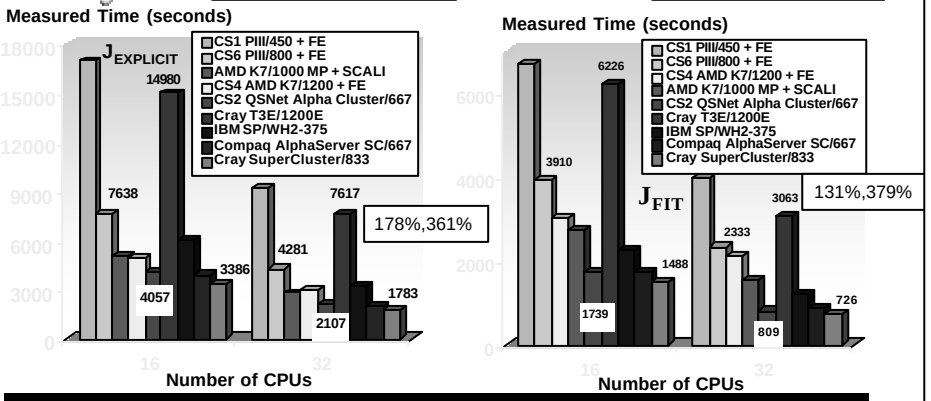

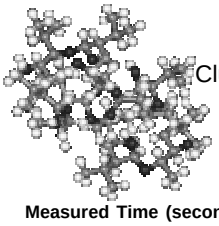

GAMESS-UK: DFT HCTH on Valinomycin.

Impact of Coulomb Fitting: Cray T3E/1200, Cray Super Cluster/833, Compaq AlphaServer SC/667 and SGI Origin R14k/500

Basis: DZV_A2 (Dgauss)
A1_DFT Fit: 882/3012

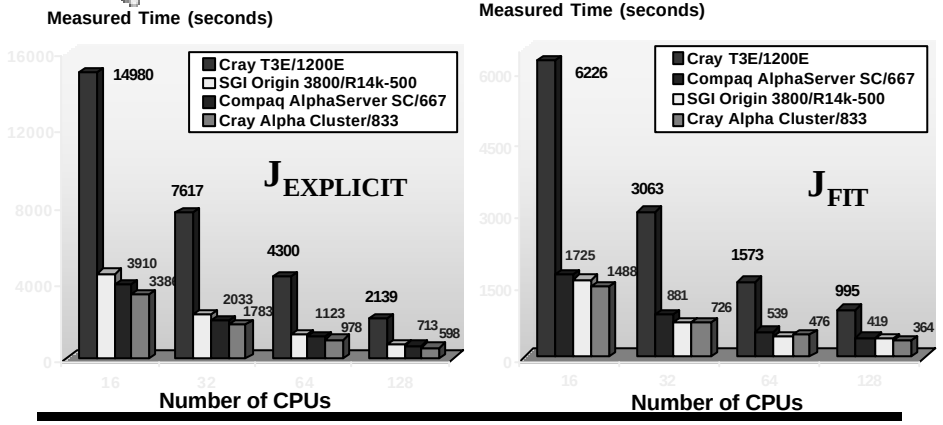

GAMESS-UK: DFT HCTH on Valinomycin. Speedups for both Explicit and Coulomb Fitting.

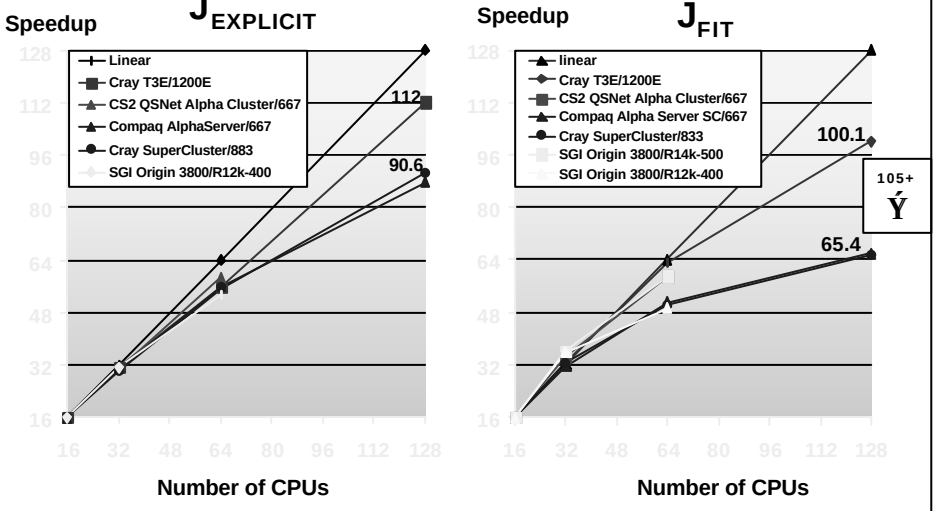

DFT HCTH Performance : Valinomycin

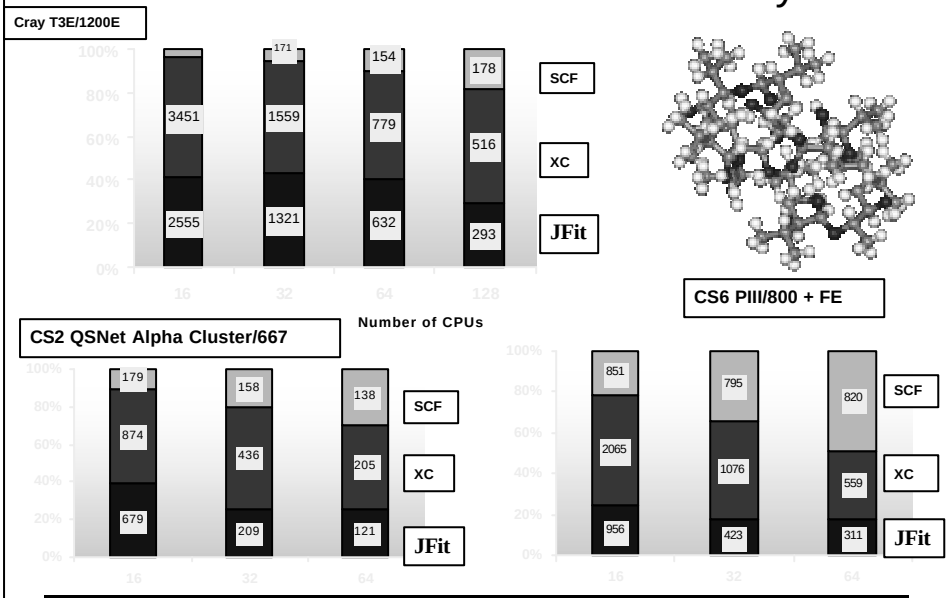

GAMESS-UK: DFT S-VWN

Impact of Coulomb Fitting: Compaq AlphaServer SC /833

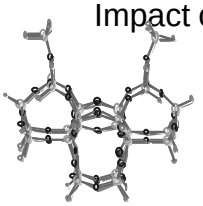

Basis: DZVP, DZVP2 (DGAUSS)
Fit: DGAUSS A1, A2

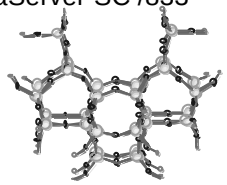

$Si_{26}O_{37}H_{36}$ 1199/2818 $Si_{28}O_{67}H_{30}$ 1687/3928

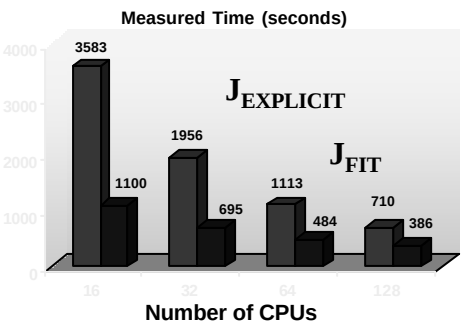

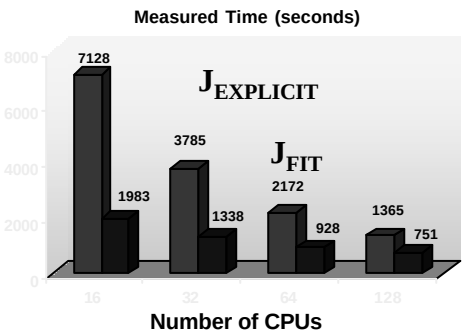

DFT Coulomb Fit - GAMESS-UK

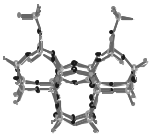

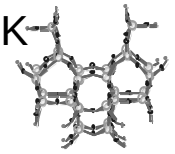

$Si_{26}O_{37}H_{36}$ 1199/2818 $Si_{28}O_{67}H_{30}$ 1687/3928

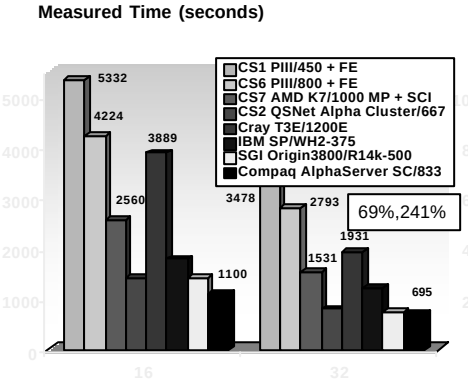

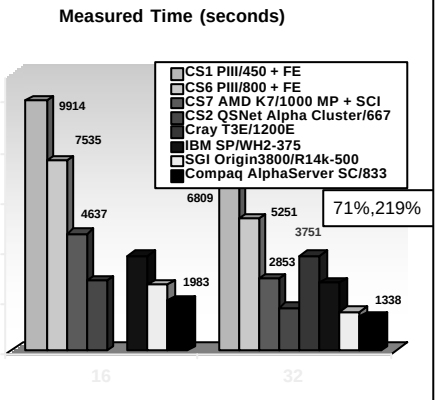

DFT JFit Performance : $Si_{26}O_{37}H_{36}$

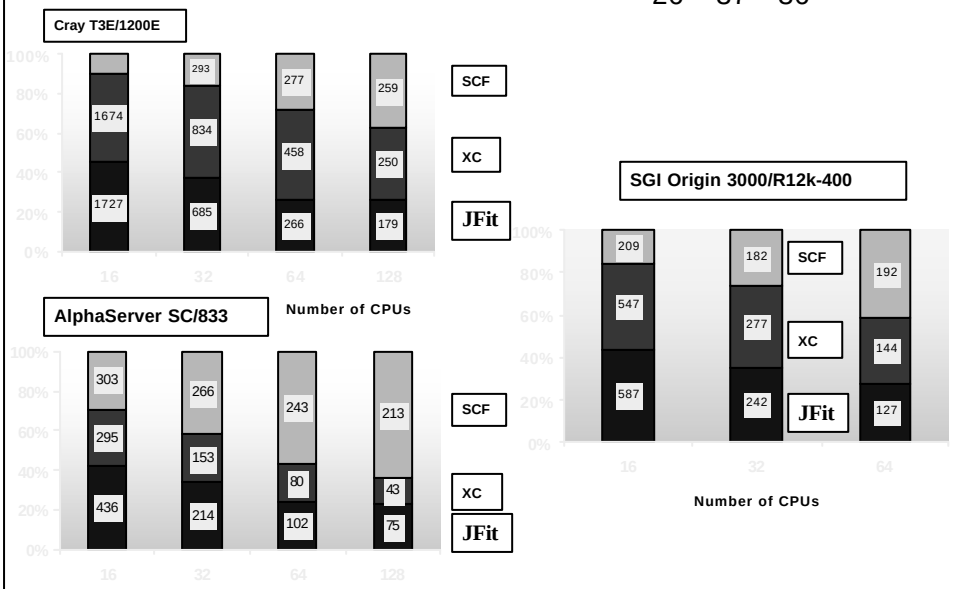

MP2 Gradient Algorithms

Serial

- Conventional
  - integrals written to disk
  - read back, transformed, written out, resorted etc.
  - heavy I/O demands
- Direct/Semi-direct (Frisch, Head-Gordon & Pople, Hasse and Ahlrichs)
  - replace all/some I/O with batched integral recomputation

Parallel

- Poor I/O-to-compute performance of MPPs
  - direct approach
- Current MPPs have large global memories
- Store subset of MO integrals
  - reduce number of integral recomputations
  - increase communication overhead
- Subset includes VOVO, VVOO, VOOO,
  - VVVO-class too large to store
  - compute VVVO-terms in separate step

Performance of MP2 Gradient Module

Cray T3E/1200, High-end and Commodity-based Systems

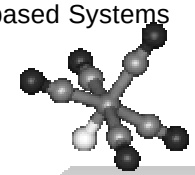

- $Mn(CO)_5H$ - MP2 geometry optimisation
- BASIS: TZVP + f (217 GTOs

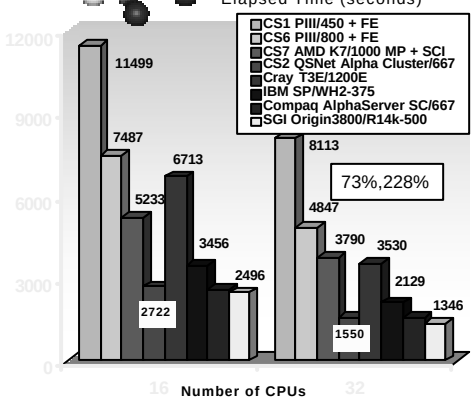

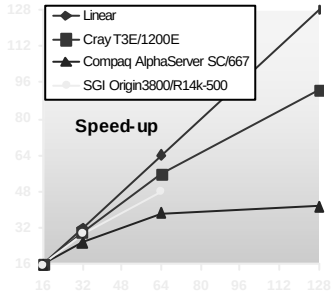

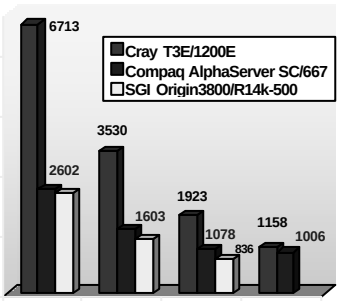

SCF Analytic 2nd Derivatives Performance

Cray T3E/1200, High-end and Commodity-based Systems

$(C_6H_4(CF_3))_2$: Basis 6-31G (196 GTO)

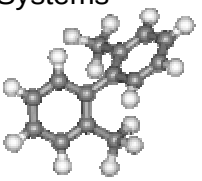

- **Terms from MO 2e-integrals in GA storage (CPHF & pert. Fock matrices); Calculation dominated by CPHF:**

Gaussian98 - L1002 (CPU) - 32 nodes: 1181 secs; 64 nodes: 1058 secs.

GAMESS-UK (total job time); 128 nodes: 499 secs.

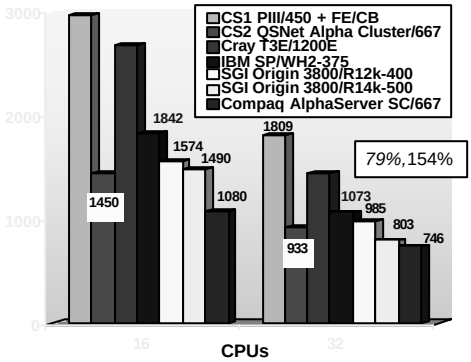

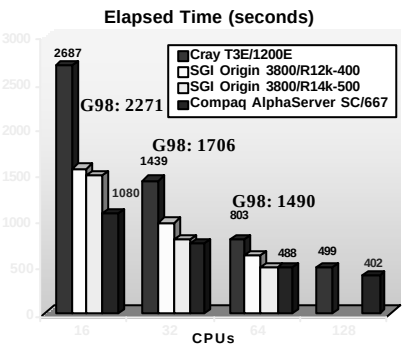

The QM/MM Modelling Approach

- Couple quantum mechanics and molecular mechanics approaches
- QM treatment of the active site
  - reacting centre
  - problem structures (e.g. complex transition metal centre)
  - excited state processes (e.g. spectroscopy)
- Classical MM treatment of environment
  - enzyme structure
  - zeolite framework
  - explicit and/or dielectric solvent models

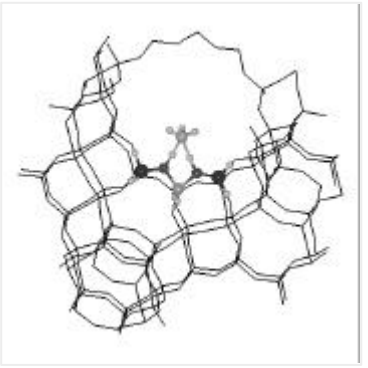

QM/MM Modelling - Challenges

- Methodological validation
  - establish reliability of both QM and MM schemes
  - QM/MM coupling schemes introduce additional artefacts
  - consistency of QM and MM energy expressions
- Computational demands
  - macromolecular systems, with extended conformational space
    - conformational search problems
    - entropic contributions
  - QM component means an expensive energy and gradient evaluation
- Software Complexity
  - range of forcefield types
  - wide variation in QM and MM program design
  - close integration needed for performance (e.g. HPC), but weak coupling simplifies maintenance (e.g. incorporating new versions of QM and MM packages)

Quantum Simulation in Industry (QUASI)

- Software Development
  - Address barriers to uptake of existing QM/MM methodology
    - explore range of QM/MM coupling schemes
    - enhance performance of geometry optimisation for large systems
    - maintain flexible approach, address enzymes, zeolites and metal oxide surfaces
    - adopt modular scheme with interfaces to industry standard codes
- High Performance Computing
  - Scalable MPP implementation
  - QM/MM MD simulation based on semi-empirical *ab-initio* and DFT methods
- Demonstration Applications
  - Value of modelling technology and HPC to industrial problems
  - Beowulf EV6-based solution
- Exploitation
  - Disseminate results through workshop, newsletters etc.

QUASI Partners

CLRC Daresbury Laboratory
P. Sherwood, M.F. Guest, A.H. de Vries
Royal Institition of Great Britain
C.R.A Catlow, A. Sokol
University of Zurich / MPI Mulheim
W. Thiel, S. Billeter, F. Terstegen.
ICI Wilton (UK)
J. Kendrick (CAPS), J. Casci (Catalco)
Norsk Hydro (Porsgrunn, Norway)
K. Schoeffel, O. Swang (SINTEF)
BASF (Ludwigshafen, Germany)
A. Schaefer

TURBOMOLE : RI-DFT Performance

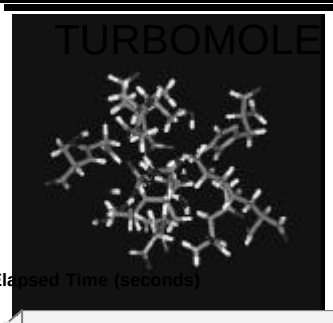

- R. Ahlrichs, M. Bar, M. Haser, H. Horn and Ch. Kolmel, Chem. Phys. Letts., 162 (1989) 165.

Benchmark Example:

- part of the TIM-enzyme + Ligand (197 atoms)
- SVP basis (1490 GTOs)
- BP96 Functional

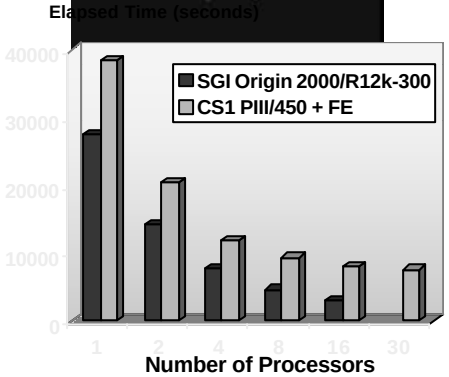

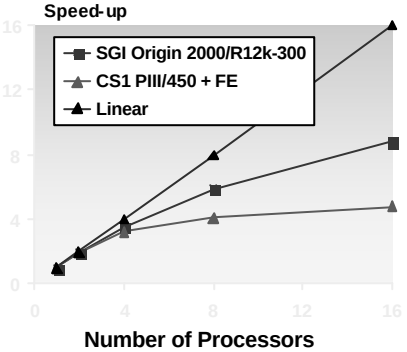

QM/MM Implementations

- Specialised for a classical modelling approach, by integrating QM code into MM package
  - CHARMM + GAMESS(US), MNDO (Harvard & NIH)
  - AMBER + Gaussian (UCSF, Manchester)
  - GULP + TURBOMOLE (Berlin)
  - CHARMM + GAMESS (UK) (Daresbury & NIH)
    - Gaussian blur / double link atom
- Generalised approaches, e.g. ChemShell
  - Modular architecture
    - Tcl Interpreter
      - Optimisation, dynamics, QM/MM coupling
    - 3rd party code interfaces
      - GAMESS-UK, Gaussian, TURBOMOLE, MNDO, DL_POLY, GULP, CHARMM etc
  - Basis for QUASI project "Quantum Simulation in Industry".
    - applications to metal oxide, zeolite, and enzyme systems

GAMESS-UK Version 6.2
IV. QM/MM Interface with CHARMM

- Implemented in collaboration with Bernie Brooks, Eric Billings, (NIH, Bethesda Maryland)
- Functionality:
  - Similar to existing *ab-initio* interfaces; CHARMM side follows coupling to GAMESS(US) (Milan Hodoscek)
  - Support for Gaussian delocalised point charges implemented in GAMESS-UK, based on 2- and 3- centre integral and derivative integral drivers from the CCP1 DFT module, (Phillip Young).
- Availability:
  - CHARMM-capable code incorporated into GAMESS-UK Version 6.2.
  - CHARMM (implemented in c26b2) requires independent licensing from Prof Martin Karplus.

Enzyme Catalysis Applications

- **Enzyme/inhibitor binding energetics for thrombin**
- **Mechanistic studies of enzyme catalysis - triosephosphate isomerase (TIM)**

Demonstration phase	Target Applications
■ **Variation of inhibitor binding enthalpies and free energies with QM region and electrostatic interactions** ■ **Determination of activation energies, variation with QM scheme and QM/MM coupling.** ■ **Comparison of substrate structure with X-ray results**	■ **Influence of active site features on inhibitor binding energies and activation energies.** ■ **Systematic study of free energies of binding for novel inhibitors, inhibitor design** ■ **Understanding the mechanism of TIM action.**

Lead Partner: BASF

QM/MM Applications

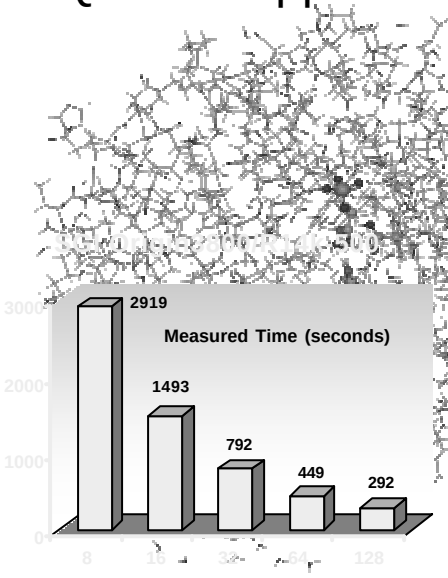

Triosephosphate isomerase (TIM)

- **Central reaction in glycolysis, catalytic interconversion of DHAP to GAP**
- **Demonstration case within QUASI (Partners UZH, and BASF)**

- QM region 35 atoms (DFT BLYP)
  - **include residues with possible proton donor/acceptor roles**
  - **GAMESS-UK, MNDO, TURBOMOLE**
- MM region (4,180 atoms + 2 link)
  - **CHARMM force-field, implemented in CHARMM, DL_POLY**

pallas

Vampir 2.5

Visualization and Analysis of MPI Programs

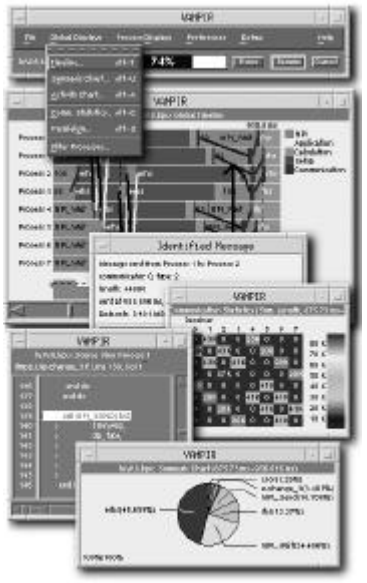

Performance Analysis of GA-based Applications using Vampir

GAMESS-UK on High-end and Commodity class machines

- **extensions to handle GA applications**

Performance analysis of GA-based applications using the Vampir tool

- Tool for Performance Analysis - VAMPIR & VAMPIR Trace
  - VAMPIR - analysis of trace files
  - VAMPIR Trace
    - Trace Library for MPI applications
    - Extensions to handle GA applications
- Case Studies
  - DFT Calculations on Zeolite Fragments (347 - 1687 GTOs) with Coulomb Fitting
  - High-end and Commodity-based Systems
- NWChem and GAMESS-UK
  - Distributed data (NWChem) and Replicated Data (GAMESS-UK)
  - Analysis of GAs and PeIGs

Instrumenting single-sided memory access

- Approach 1: Instrument the puts, gets and data server
  - Advantage: robust and accurate
  - Disadvantage: one does not always have access to the source of the data server
- Approach 2: Instrument the puts and gets only, “cheating” on the source and destination of the messages
  - Advantage: no instrumentation of the data server required
  - Disadvantage: timings of the messages are inaccurate in case of non-blocking operations

Runtime tracing options

- The tracing of activities can be modified at runtime through a configuration file.
- Tracing of messages can not be changed.
- VTTRACEON and VTTRACEOFF should be used sparingly.

```
Logfile-name /home/user/prog.bpv
Symbol nnodes off
Symbol nodeid off
Symbol GA_Nnodes off
Symbol GA_Nodeid off
```

Practical issues

- The vampirtrace library and evaluation licenses can be downloaded from http://www.pallas.com/
- Evaluation licenses are limited to 32 processors

GAMESS-UK / $Si_8O_{25}H_{18}$: 8 CPUs: One DFT Cycle

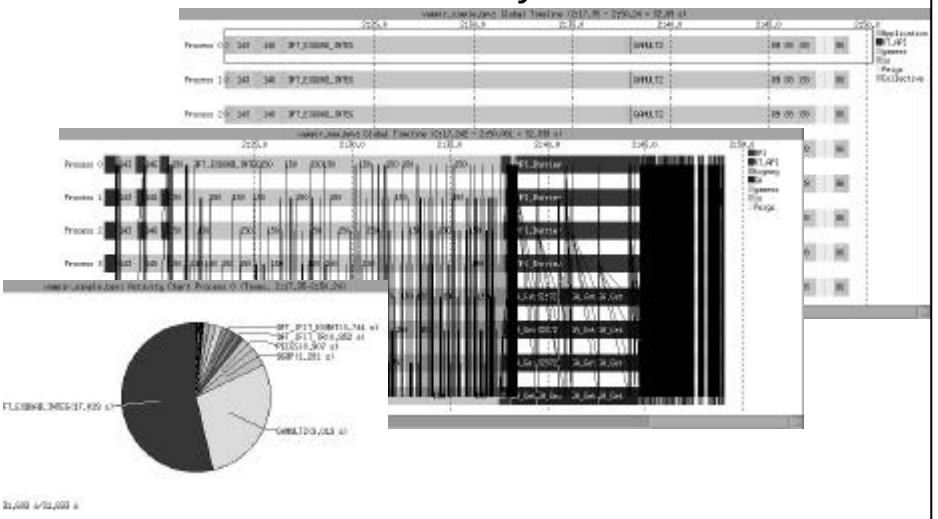

GAMESS-UK / $Si_8O_{25}H_{18}$: 8 CPUs

$Q^{\dagger}HQ$ (GAMULT2) and PEIGS

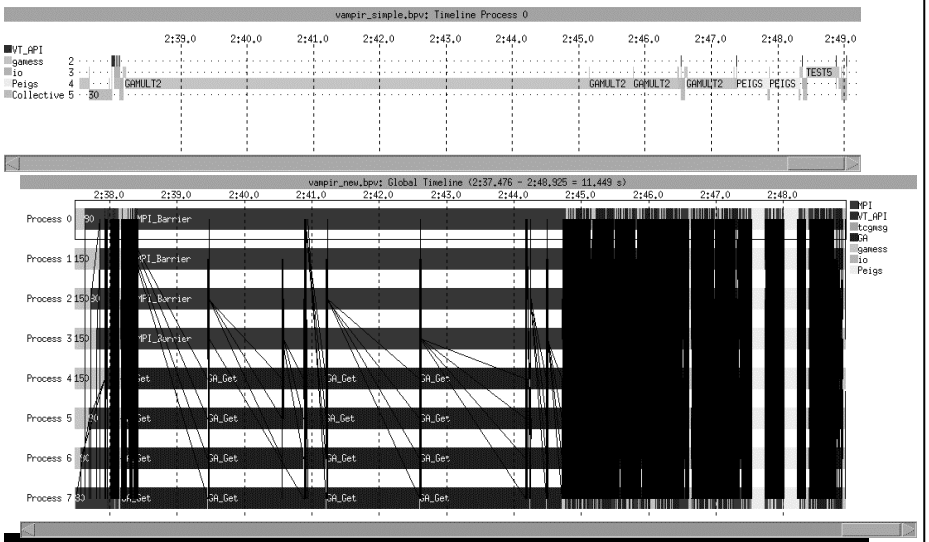

Performance Analysis of GA-based Applications using Vampir - Summary

- Daresbury and PALLAS collaboration
- Tool for Performance Analysis - VAMPIR & VAMPIR Trace
  - Extended to handle GA Applications
- Applied in a number of DFT Calculations on Zeolite Fragments on a variety of high-end and commodity-based platforms
- Instrumentation of both NWChem and GAMESS-UK:
  - Distributed data (NWchem)
  - Replicated Data (GAMESS-UK)
  - Analysis of GAs and PeIGs
- Findings
  - non-intrusive
  - Tracing of substantial runs possible
    - Size of trace files in distributed data applications
  - Use in quantifying scaling problems
    - e.g. GA_MULT2 in GAMESS-UK

Materials Simulation Codes

Plane Wave DFT Codes:
- CASTEP
- VASP
- CPMD

These codes have similar functionality, power and problems. CASTEP is the flagship code of UKCP and hence subsequent discussions will focus on this.

Local Gaussian Basis Set Codes:
- CRYSTAL

This code presents a different set of problems when considering performance on HPC(x).

SIESTA and CONQUEST:
- O(n) scaling codes which will be extremely attractive to users.
- Both are currently development rather than production codes.

for MgO and TiO_2

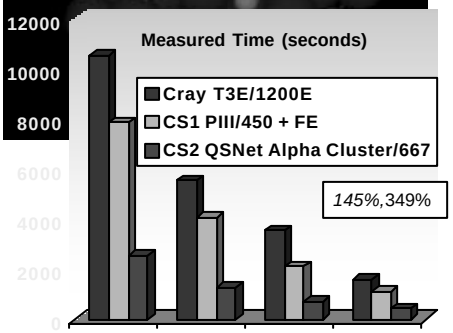

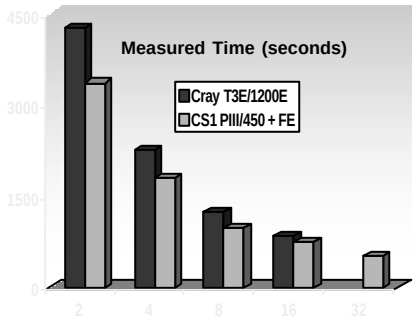

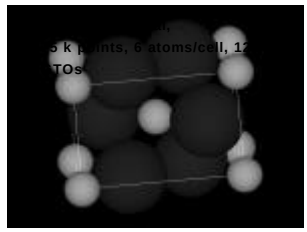

CRYSTAL - 2000

- Distributed Data implementation
- Benchmark:
  - An Acid Centre in Zeolite-Y (Faujasite)
  - Single point energy
  - 145 atoms / cell, No symmetry / 8k-points
  - 2208 basis functions, (6-21G*)

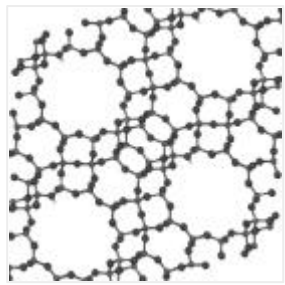

Elapsed Time (seconds)

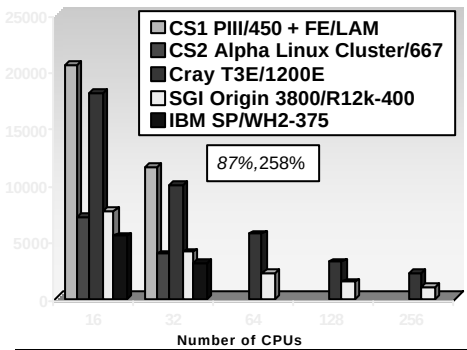

Speed-up

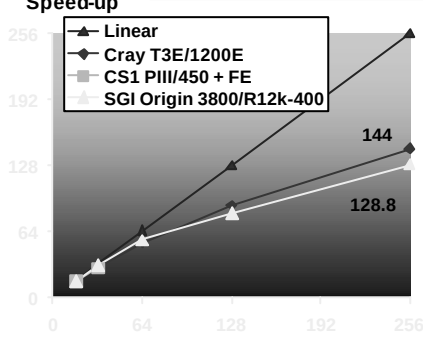

Plane Wave Methods: CASTEP

$$y_j^{\vec{k}}(\vec{r}) = \sum_{\vec{G}\,:\,(\vec{k}+\vec{G})^2 < E_{cut}} C_{j,\vec{G}}^{\vec{k}}\, e^{-i(\vec{k}+\vec{G}).\vec{r}}$$

- Direct minimisation of the total energy (avoiding diagonalisation)
- Pseudopotentials must be used to keep the number of plane waves manageable
- Large number of basis functions N~10^6 (especially for heavy atoms).

The plane wave expansion means that the bulk of the computation comprises large 3D Fast Fourier Transforms (FFTs) between real and momentum space.

- These are distributed across the processors in various ways.
- The actual FFT routines are optimized for the cache size of the processor.

CASTEP - The UK Car-Parrinello Consortium

UK Car-Parrinello Consortium

- The Cambridge Serial Total Energy Package CASTEP (M. Payne et al.) calculates the total energy, forces and stresses in a 3D-periodic system.
- Rev. Mod.Phys. 64 (1992) 1045
- DFT, plane-waves, pseudo-potentials & FFT's

CASTEP 4.2β Key Features:

- Ultrasoft pseudo-potentials with non-linear core corrections
- Range of minimisation methods: Density Mixing, RM-DIIS, Conjugate Gradients band-by-band & all-bands. Full structural relaxation and MD
- LD and GGAs, spin-polarisation

Parallelization of CASTEP

- A Number of parallelization methods are implemented:
  - **k-point:** a processor holds all the wavefunction for a k-point (MPI_ALLTOALLV is NOT required) BUT for large unit cells $N_k \Rightarrow 1$ i.e small CPU count.
  - **G-vector:** a processor holds part of the wavefunction for all k-points (MPI_ALLTOALLV is over ALL CPUs) i.e. biggest systems with 1 K point
  - **mixed kG:** k-points are allocated amongst processors, the wavefunctions sub-allocated amongst processors associated with their particular k-points i.e. MPI_ALLTOALLV is over N_{CPUs} / N_k - intermediate cases.
- On HPC hardware the desired method is either k or kG as this minimizes inter-processor communication, specifically MPI_ALLTOALLV.
- However, on large numbers of processors such distributions will still be problematic. New algorithms will therefore need to be developed to overcome latency problems.

CPU Optimizations: Efficiency on Commodity-based Systems

- CASTEP on the Cray T3E and SGI Origin 3800 systems use FFT code fully optimized for the processor L1-cache.
- Extended to other cache-based processors. FFT operations are performed on chunks of data that fit in L1-cache and run at maximum speed. i.e. FFT exploited efficiently on cache-based CPUs.
- Example:

  64 processor (32 dual 1GHz AMD K7) system, SCALI interconnect

TiN: A TiN 32 atom slab, 8 k-points, single point energy calculation with Mulliken analysis,

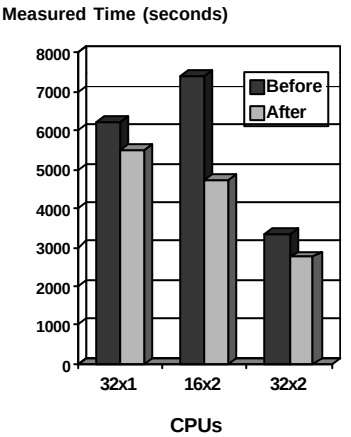

CASTEP 4.2 - Parallel Benchmark

Measured Time (seconds)

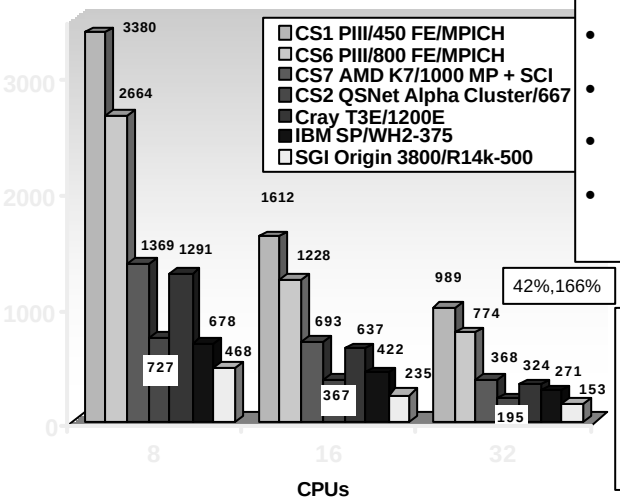

Chabazite

- Acid sites in a zeolite. (Si_{11} O_{24} Al H)
- Vanderbilt ulatrasoft pseudo-potential
- Pulay density mixing minimiser scheme
- single k point total energy, 96 bands
- 15045 plane waves on 3D FFT grid size = 54x54x54; convergence in 17 SCF cycles

	Time (comms)
IBM SP/WH2-375	157
Cray T3E/1200E	90
CS1 PIII/450+FE	660
CS6 PIII/800+FE	600
CS7 AMD K7/1000 + SCI	242
CS2 QSNet Alpha	111
SGI Origin 3800/R14k	71

CPMD - Car-Parrinello Molecular Dynamics

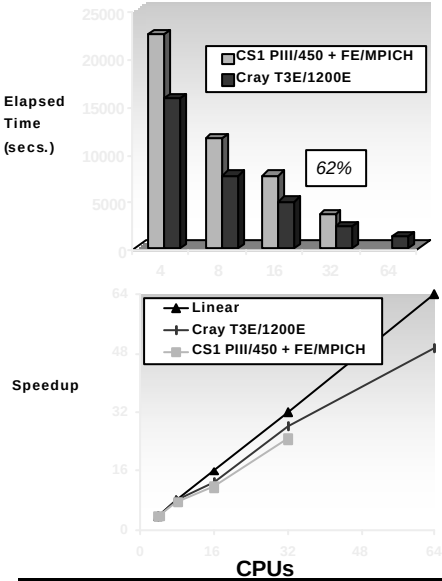

CPMD

- Version 3.3: Hutter, Alavi, Deutsh, Bernasconi, St. Goedecker, Marx, Tuckerman and Parrinello (1995-1999)
- DFT, plane-waves, pseudo-potentials and FFT's

Benchmark Example: Liquid Water

- Physical Specifications:
  32 molecules, Simple cubic periodic box of length 9.86 A, Temperature 300K
- MD parameters;
  Time step 7 au = 0.169 fs; Length test run 200 steps = 34 fs
- Electronic Structure;
  BLYP functional, Trouillier Martins pseudopotential, Recriprocal space cutoff 70 Ry = 952 eV

Sprik and Vuilleumier (Cambridge)

CPMD on High-end Computers

- **Performances on NH2 nodes**
  - 256 CPUs and 8 MPI tasks/node including IO (RD30WFN / WR30WFN)
  - 390 Mflops per CPU, 99.8 Gflops
  - 252*252*252 - 8 MPI tpn
  - Use of ESSL routines instead of Lapack; Some routines were "OpenMPed"
- **Mixed mode** MPI and OMP (IBM)
- CPMD is dominated by ESSL SMP routines (ROTATE and OVLAP).
- 4 MPI tpn and 2 SMP tpn is 1.36 faster than flat MPI on a given MESH
- **Power4 estimates:** The key is FFTCOM (MPall2all). Average speedup of 2.0 on a R-H LPARd system

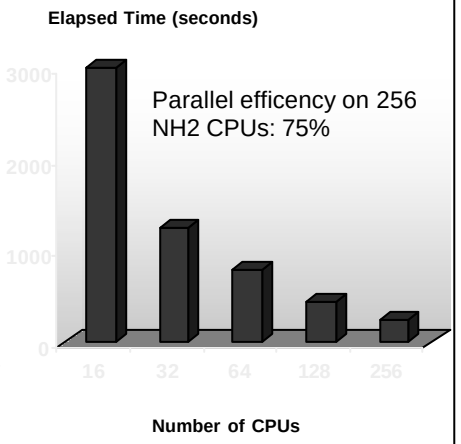

ANGUS: Combustion modelling (regular grid)
The Cray T3E/1200, IBM SP/WH2 and Beowulf Systems

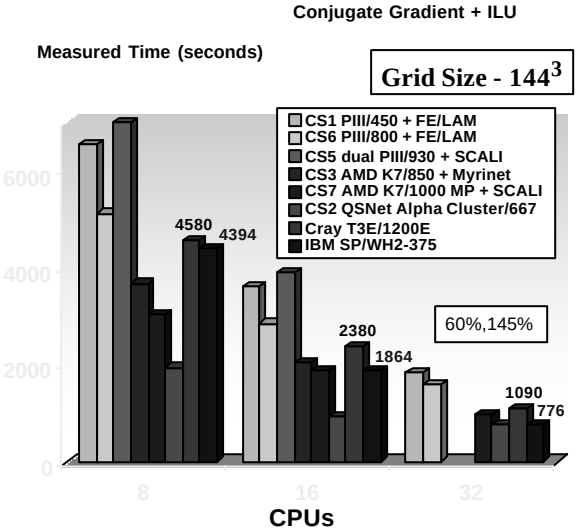

Direct numerical simulations (DNS) of turbulent pre-mixed combustion solving the augmented Navier-Stokes equations for fluid flow.

Discretisation of equations is performed using standard 2nd order central differences on a 3D-grid.

Pressure solver utilises either a conjugate gradient method with modified incomplete LU pre-conditioner or a multi-grid solver (both make extensive use of Level 1 BLAS) or fast Fourier transform.

ANGUS: Combustion modelling (regular grid)

Memory Bandwidth Effects: The IBM SP and Alpha Linux System

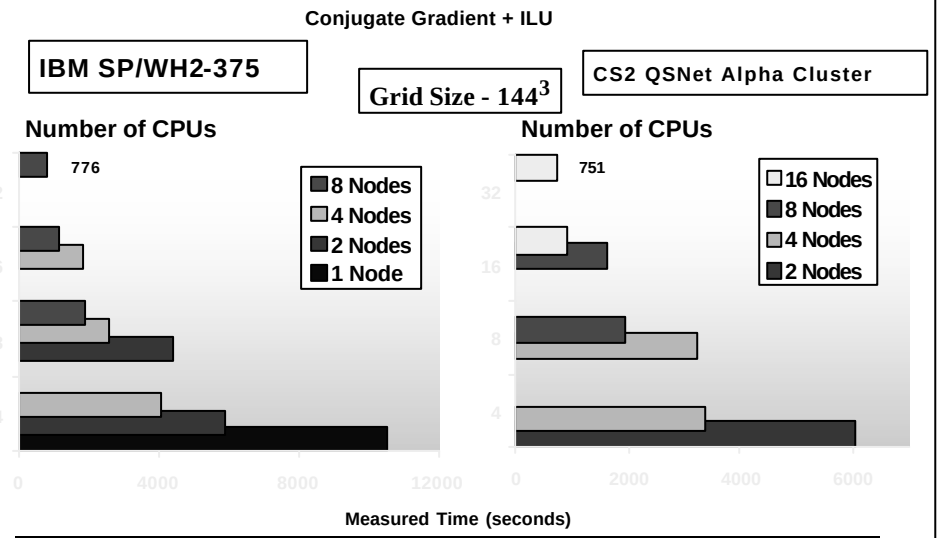

FLITE3D: An Industrial Aerospace Code

F18 test, 3444350 elements

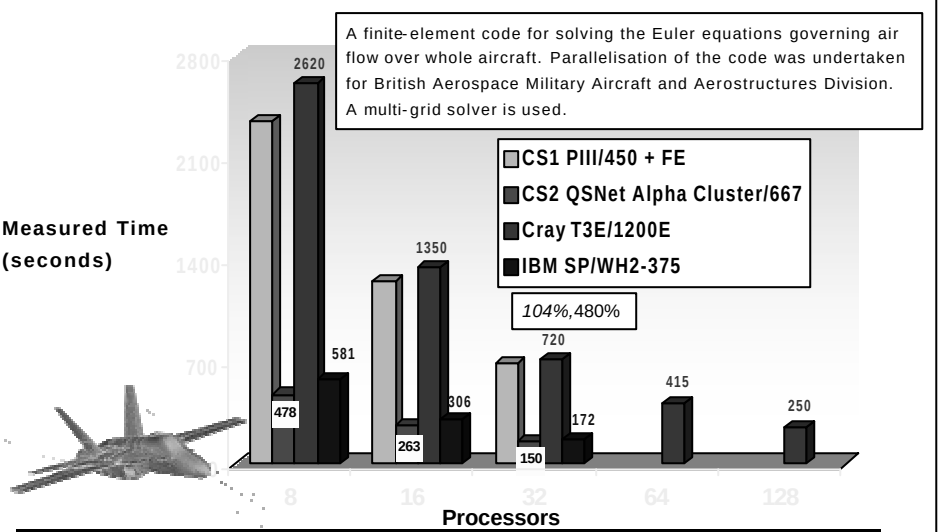

Beowulf Comparisons with the T3E & O3800/R14k-500

CSx - Pentium III + FE
% of 32-node Cray T3E/1200E

GAMESS-UK	CS1	CS6
SCF	53-69%	96%
DFT	65-85%	130-178%
DFT (Jfit)	44-77%	65-131%
DFT Gradient	90%	130%
MP2 Gradient	44%	73%
SCF Forces	80%	127%
NWChem (DFT Jfit)		50-60%
REALC	67%	
CRYSTAL	145%	
DL_POLY		
Ewald-based	95-107%	151-184%
bond constraints	34-56%	69%
CHARMM	96%	172%
CASTEP	33%	42%
CPMD	62%	
ANGUS	60%	68%
FLITE3D	104%	

CS2 - QSNet Alpha Linux Cluster
% of 32-node Cray T3E and O3800/R14k-500

GAMESS-UK		
SCF	256%	99%
DFT †	301-361%	99%
DFT (Jfit)	219-379%	89-100%
DFT Gradient †	289%	89%
MP2 Gradient	228%	87%
SCF Forces	154%	86%
NWChem (DFT Jfit) †	150-288%	74-135%
CRYSTAL †	349%	
DL_POLY		
Ewald-based †	363-470%	95%
bond constraints	143-260%	82%
CHARMM †	404%	78%
CASTEP	166%	78%
ANGUS	145%	
FLITE3D †	480%	

Summary

- Computational Chemistry - Background
- Commodity-based and High-end Systems
  - Prototype Commodity Systems; CS1 - CS7
  - High-end systems from Cray, SGI, IBM and Compaq
  - Performance Metrics
- Application performance
  - Electronic Structure
    - Distributed data: GAs and PeIGS
    - NWChem, GAMESS-UK, and TurboMole
  - Molecular Simulation
    - DL_POLY and CHARMM
- Application performance analysis
  - VAMPIR and instrumenting the GA Tools
  - Linux Alpha Cluster delivers between 150-400% of T3E/1200E, 78-100% of SGI Origin 3800/R14k-500

CS2 - QSNet Alpha Linux Cluster
% of 32 CPU O3800/R14k-500

GAMESS-UK	
SCF	99%
DFT	99%
DFT (Jfit)	89-100%
DFT Gradient	89%
MP2 Gradient	87%
SCF Forces	86%
NWChem	74-135%
DL_POLY	
Ewald-based	95%
bond constraints	82%
CHARMM	78%