

Scaling Computational Chemistry to Peta-FLOP Computers

Robert J. Harrison

Pacific Northwest National Laboratory

This work is funded by a grand-challenge application project supported by the U.S. Department of Energy, the division of Mathematics, Information and Computational Sciences, Office of Computational and Technology Research under contract DE-AC06-76RLO 1830 with Battelle Memorial Institute (Pacific Northwest National Laboratory). This research was performed in part using the Molecular Science Computing Facility in the Environmental Molecular Sciences Laboratory at the Pacific Northwest National Laboratory (PNNL), and also used resources of the National Energy Scientific Computing Center which is supported by the Office of Energy Research of the U.S. Department of Energy under contract DE-AC03-76SF0098.

“Slant” - Greg Bear, 1997


- **The date is 2052 or later**

- **From near the end of chapter 1**

“Giffey believes he will face an unauthorized artificial intelligence - Omphalos’s own advanced petaflop INDA, perhaps even a thinker.”

P.s. What’s an INDA? He never defines the term in this book.

Impact of algorithmic scaling



- **1000-fold speedup?**
 - Only for same calc.
- **Does not imply 1000-x larger or more accurate calculation**

A Renaissance Period in Q-Chem

■ Traditional electronic structure

- Dense matrices and classical linear algebra and BLAS
- Coulomb and exchange interaction via 2-e integrals, $O(N^4)$
- High-precision calculations $O(N^7)$ or worse

■ We will not be running these algorithms 10 years from now

- Fast Coulomb and exchange evaluation $O(N)$ - $O(N^3)$
- More varied choice of representations, improved approximate methods
- Factorization and low-rank representations
- Novel theory

■ When we get petaflop computers comp. chem. will be at least 1,000,000 times faster than on the teraflop computers of today

Is this exciting? Yes!

■ Solution of longstanding problems

- Detailed chemical reaction dynamics in solution
- Molecular level environmental chemistry
- Ab initio design of catalysts
- Ab initio molecular dynamics

■ New “Grand Challenges”

■ Computational chemistry as a design tool

- Interactive, quantitative response to posed questions
- Will the real impact of petaflop computing be that you can have all the tflop computing you need without waiting? On your desk?

Overview

- **Computational chemistry past and present**
 - Objectives and methodology
 - Algorithms - scalar, vector, and parallel
 - Burning issues - scaling w.r.t. system size, bottlenecks
- **The near future**
 - Approaching linear scaling - reality and impact on memory locality
 - Fully numerical solution - for effective 1-particle models
- **Impact and use of petaflops computers**
 - Example applications
 - Machine requirements; architectural flexibility

Electronic structure of molecules - I

■ What can we predict?

- In principle everything! In practice
- For very small molecules (4 atoms?) - (nearly) full chemical reaction dynamics
- For small molecules (10 atoms?) - spectra, energetics and structure at a level of precision rivalling experiment
- For larger molecules - lower precision - higher precision obtained from trends
- Application areas - drug design, petrochemicals, catalysis, waste processing,...

■ The electronic Schrödinger equation

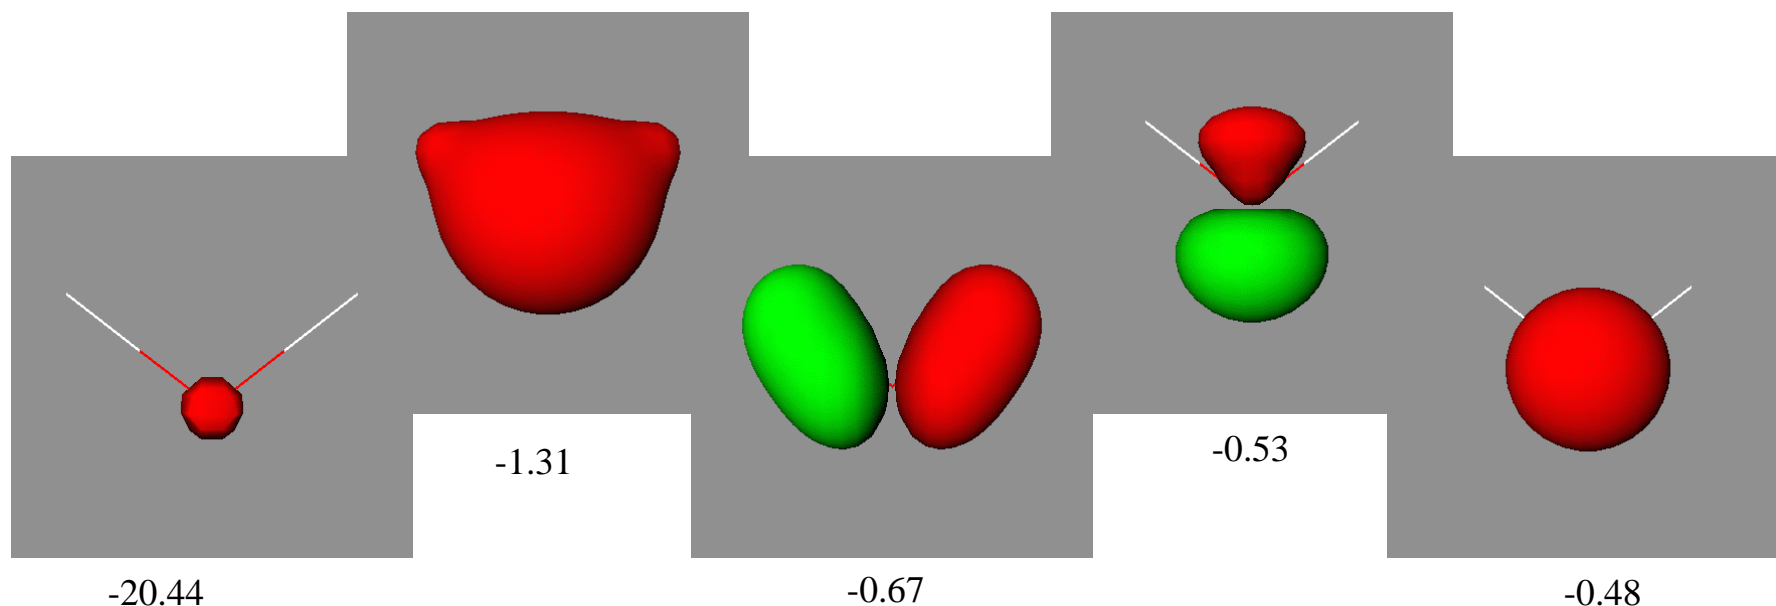
- A linear eigen-equation in the coordinates of the N electrons

$$\hat{H} \psi = E \psi, \quad \hat{H} = \sum_i \frac{1}{2} \mathbf{p}_i^2 - \sum_i \frac{Z_I}{R_{iI}}$$

Electronic structure of molecules - II

■ Atomic and molecular orbitals

- ▶ An effective one-particle model. Each electron moves in the average field of all other electrons (self-consistent field, Hartree-Fock, ...)
- ▶ Provides the structure of the periodic table and the chemical bond
- ▶ E.g., molecular orbitals for water, H_2O



Electronic structure of molecules - III

■ Known properties of the wavefunction

- Cusps at the nuclei and coalescence of two or more electrons
- Long-range exponential decay
- Shell-structure within the atom
- Molecules consist of (strongly-) perturbed atoms

■ Basis sets - the algebraic approximation

- Incorporate known behaviour - very compact representation of orbitals (3-D)
- Subsequent theory expressed as (mostly dense) matrix operations
- Exponential (Slater) bases - integrals too hard
- Gaussian basis - integrals tractable, no loss of precision in practice
- The root of nearly all our scaling woes but presently indispensable

$$\left| \underline{r} - \underline{A} \right|^n Y_{lm}(\underline{r} - \underline{A}) \longrightarrow \sum_k c_k e^{-\zeta_k (\underline{r} - \underline{A})^2}$$

Electronic structure of molecules - IV

■ 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 - V

■ Density functional theory

- Exact energy is a function(al) only of the 1-electron density - a 3-D entity
- But ... don't know the functional - good, educated guesses; nothing systematic
- Scales from $O(N)$ to $O(N^4)$

■ Configuration Interaction

- A linear expansion mixing important orbital occupations - sparse eigen problem
- Slowly convergent - large expansions are necessary; state of art is 10^9 parameters
- 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

A Typical Current Application

- **What's going to happen to the uranium leaking from the tank?**
 - Input into subsurface transport model are chemical speciation and binding data

- **Interpret the laser line-narrowed fluorescence spectrum of uranium oxide in solution absorbed onto a mineral surface**
 - Predict trends in vibrational frequencies as a function binding site
 - Form a molecular model - $(\text{UO}_2^{2+})(\text{H}_2\text{O})_n$ on a fragment of clay surface (Si,O,Al)
 - Understand the hydrated species with a variety of ligands - both electronic states
 - optimize structures (many possibilities) and compute vibrational frequencies
 - selectively replace waters with other ligands, and examine solvation effects



- compare with experiment, or more rigorous theories

- Validate the model of the mineral surface

Where we need to be

■ Eliminate empirical parameters and approximations

- Currently necessary for efficiency or tractability
- True predictive capability rather than interpolative
- (Semi-)empirical models still valuable

■ Full molecular models

- Reduce the distance between model and reality
 - E.g., explicit description of more solvent molecules
- More detailed computations - E.g., accurate electronic spectra
 - Feasible now only for few atom, isolated molecules
 - Eliminate gap between observation and theory

■ Reduce risk

- We, in principle, can get the right answer for the right reasons

Computational Chemistry Now - CPU

■ Local density functional theory

- Effective $O(N-N^3)$ scaling attainable through fast-Poisson solvers and careful numerical quadrature - $O(N^3)$ is due to dense linear algebra with small prefactor
- Up to 2000 basis functions routine - several hours on high-end workstation.

■ Hartree-Fock and DFT with exact exchange

- $O(N^2-N^3)$ scaling attainable but with much larger prefactors (must compute 2-e integrals for exact exchange)
- Most algorithms for fast-coulomb/exchange are not applicable for high precision
 - 2000 functions, low-accuracy (6-31g*) - a day on a high-end workstation
 - 2000 functions, high-accuracy (aug-cc-pVTZ) - several days on 32 nodes

■ Correlated methods (CCSD(T))

- Formal $O(N^7)$ scaling only avoided with undesirable approximations
- Up to 500 functions feasible - several to many hours on 128 nodes

Comp. Chem. Now - Memory

■ Hartree-Fock and DFT with exact exchange

- Several $O(N^2)$ dense matrices - possible to distribute with some effort
- Recompute or store integrals - trend in disk/cpu speed favors recomputation
- Good memory locality in integral evaluation - 4-32 MB cache? Vectorizable.
- Poorer memory locality in Fock-matrix formation - blocking helps

■ CCSD(T)

- Dominated by dense matrix multiplications of $O(N)$ ($N =$ up to circa 500)
- Executes near peak processor performance
- Requires several $O(N^2)$ matrices per processor
- Requires $O(N^4)$ aggregate memory (or page $O(N^3)$ chunks from disk)



■ NUMA

- ~~Very successful application in computational chemistry - use blocking~~

Robert J. Harrison, 2/16/99 to improve memory locality at all levels of the memory hierarchy *High Performance Computational Chemistry*

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
- Supported by DOE 2000, ACTS
- Used by 5 major chemistry codes, financial futures forecasting, astrophysics, computer graphics

Comp. Chem. The Near Future - I

■ Hartree-Fock and DFT with exact exchange

- Fast coulomb/exchange algorithms being designed for high-quality basis sets
 - Improves scaling only slightly - still $O(N-N^3)$
 - Reduce prefactors dramatically
- Increased memory requirement ((non-uniform) FFT grids, FMM oct-tree, ...)
- No more two-electron integrals

■ Perturbation theory (MP2 - simplest many-body model)

- $O(N)$ algorithms coming - still not applicable to highest precision
- Improve scaling from $O(N^5)$ to $O(N^3)$? Much larger prefactor.

■ Coupled-cluster methods

- Rigorous reduction from $O(N^7)$ to $O(N^6)$ but larger prefactor

- ~~Permits localization - improved sparsity - $O(N)$ theoretically obtainable,~~

Comp. Chem. The Near Future - II

■ Low-medium precision one-electron methods

- 10,000 functions routine (high-end workstation several hours/days)
- Significant sparsity in matrices - must exploit to avoid $O(N^3)$ linear algebra

■ High precision one-electron methods

- 10,000 functions feasible (good basis, exact exchange, few hours on TFLOP?)
- Modest sparsity in the matrices - dense linear algebra (sparse for low-precision)

■ Low-precision correlated methods - MP2

- Viable $O(N^3)$ or better methods - 2,000 functions routine (high-end workstation)
- 10,000 functions feasible (high-quality basis, few hours on TFLOP)

■ ~~High-precision correlated methods - CCSD(T)~~

Robert J. Harrison, 2/16/99

High Performance Computational Chemistry

- Viable $O(N^5 - N^6)$ methods, reduced prefactor - 1500 functions, days on

Ratio of memory to machine speed

■ Crude scaling-law models for the algorithmic requirements

- $m(N) = b \cdot N^l =$ memory required to run calculation
- $t(N) = a \cdot N^k =$ execution time
- $F =$ machine speed in FLOP/s

■ A single 10 day ($T = 10^6$ s) calculation using CCSD(T)

- $N = (FT/a)^{1/k}$
- Current algorithm on $F=10^{12}$ --> $N = 1,000, M = 0.8\text{TB}, M/F = 0.8$
- Current algorithm on $F=10^{15}$ --> $N = 2,683, M = 40\text{TB}, M/F = 0.04$
- Gussed future algorithm on $F=10^{12}$ --> $N = 3,162, M=2.5\text{TB}, M/F=2.5$
- Gussed future algorithm on $F=10^{15}$ --> $N = 17,783, M=450\text{TB}, M/F=0.45 \quad ??$

■ Speedup from architecture vs. algorithm, $N=5000$ on PFLOP

- Current algorithm --> 2.5 years
- Gussed future algorithm --> 1.8 hours

➤ 1,000 from TFLOP-PFLOP but 10,000 potentially from algorithm

SSI - first H/W draft Q-Chem. ideals

- Plan from '98 was initial machine would be 5-TFLOP
- Current strategy unclear (multiple smaller machines?)

Local processor FLOPS	2 GFLOPS (assumed as a likely target)
Scaling (FLOPS)	N^{**7} , $N < 1500$ basis functions
Scaling (local data)	$10N^{**2}$ words = 200 Mbytes
Scaling (global data)	$N^{**4}/100$ words = 320 Gbytes
Local memory (per proc)	512 Mbytes
Local cache	more the better up to at least 32 Mbytes
Local disk (per proc)	don't care as long as bandwidth and aggregate disk is met
IO bandwidth (per proc)	2000 MFLOPS/500 Mwords/sec = 30MBytes/Sec
Interproc bandwidth	2000 MFLOPS/30 Mwords/sec = 500 Mbytes/sec
Interproc latency	$2 * 10^{**4} / 2000$ MFLOPS = 10 microseconds
Aggregate memory	1 Tbyte
Aggregate disk	N^{**4} words = 100 Tbytes
Aggregate cross-section bandwidth	at least $nproc^{**(2/3)} * 500$ Mbytes/sec

- Feedback suggests much more detail w.r.t. memory needed

Memory locality and performance

- **CCSD(T) currently has a large ratio - compute/data**
 - Can in principle block to tolerate any nearly any memory hierarchy
 - To avoid local memory explosion, use matrix mult. of $O(N)$ dense matrices
 - As long as latency and B/W are sufficient to get two square matrices, multiply them and put the result back, conventional algorithm runs at peak speed
 - E.g., IBM SP, 120 MHz P2SC requires $N \geq 600$ for 90% peak speed
- **This ratio will diminish**
 - And the calculation becomes more irregular - load balance, data distribution, ...
- **For nearly all models more time spent in “fast(er)” methods**
 - FMM and modifications thereof - but with high precision ($L=25$ is used)
 - Dispersion and exact exchange interactions
 - 2-D quadrature

Massive parallelism

- **Demonstrated scaling to 1,000+ processors**
 - Expect same approach to scale to circa 10,000 processors with similar system parameters (c.f., Agerwala and projected IBM parameters)
 - Coarse-grain parallelism only
 - Increased emphasis on load-balance, Amdahls law, ...
- **Fine-grain parallelism not yet exploited, because**
 - Used to do it a decade ago on CRAYs, Alliants, ...
 - Compilers were not very helpful - not much more so now - hard problem
 - No need to since coarse grain was successful
- **Parallelism of $O(10^8)$ is readily obtained**
 - How fine can we go?
 - The lower the latency the easier is my job and the compilers job

Out of the mainstream

■ Quantum Monte Carlo

- Directly solves the Schrödinger equation
- Still has non-linear scaling w.r.t. system size and cannot yet answer the same breadth of questions
- Compact kernels with algorithms amenable to many novel massively parallel architectures

■ Direct multiresolution solution of LDA (and other) equations

- Solve the Kohn-Sham equations but implement all steps as fast transforms
- Avoid operations in the orbital basis
- 1-D formulation (Beyelkin)

■ Explicitly correlated wavefunctions; numerical pair equations

■ Improved density functionals

- Reached limit of current incremental refinement approach

Other Issues

■ The memory hierarchy

- Do we really have to worry about each level in the hierarchy?
 - (cache) + (local memory) + (non-local memory) , or do we need to think
 - (L1) + (L2) + (L3 SMP cache) + (local memory) + (near non-local) + ...
 - We would thrash TLBs and other directories in some proposed machines

■ External storage

- Presently we do many out-of-core solutions (e.g., scalable IO on IBM-SP)
 - 1 GFLOP cpu is approx. balanced with 30 MB/s I/O - affordable
- CPU/disk speed - Cannot afford the bandwidth? Saved by cheap memory?
- Replace I/O with more computation?

■ The size and complexity of codes

- All significant packages are 500,000+ lines of code, and always will be
- ~~Source of inertia in adapting to sudden architecture change - investment~~

Conclusions

- **Petaflops computing + decade of theory/algorithm development**
 - **Many problems will need this power**
 - **High-consequence problems (c.f., White) - catalysis, combustion, environmental chemistry, ...**
 - **... many do not**
 - **Deskside/department TFLOP computing will satisfy many applications**
 - **How big will be the pool of supercomputer users?**
 - **Basic research is necessary**
 - **Theoretical and computational chemistry, computer science, applied math**
 - **Cannot be isolated from each other - lesson from HPCC - need a focus**
 - **Shortage of skilled people**
 - **Timelines and expectations must take this into account**
-