DFT Made Faster and More Accurate

Holger Dachsel, Prakashan Korambath,
Shawn Brown, Yihan Shao, Laszlo Fusti-Molnar,
Emil Proynov, Zhengting Gan, Nick Russ,
Chumin Chang, Jing Kong

Q-Chem Inc.
Pittsburgh, PA

Science at Q-Chem Inc.

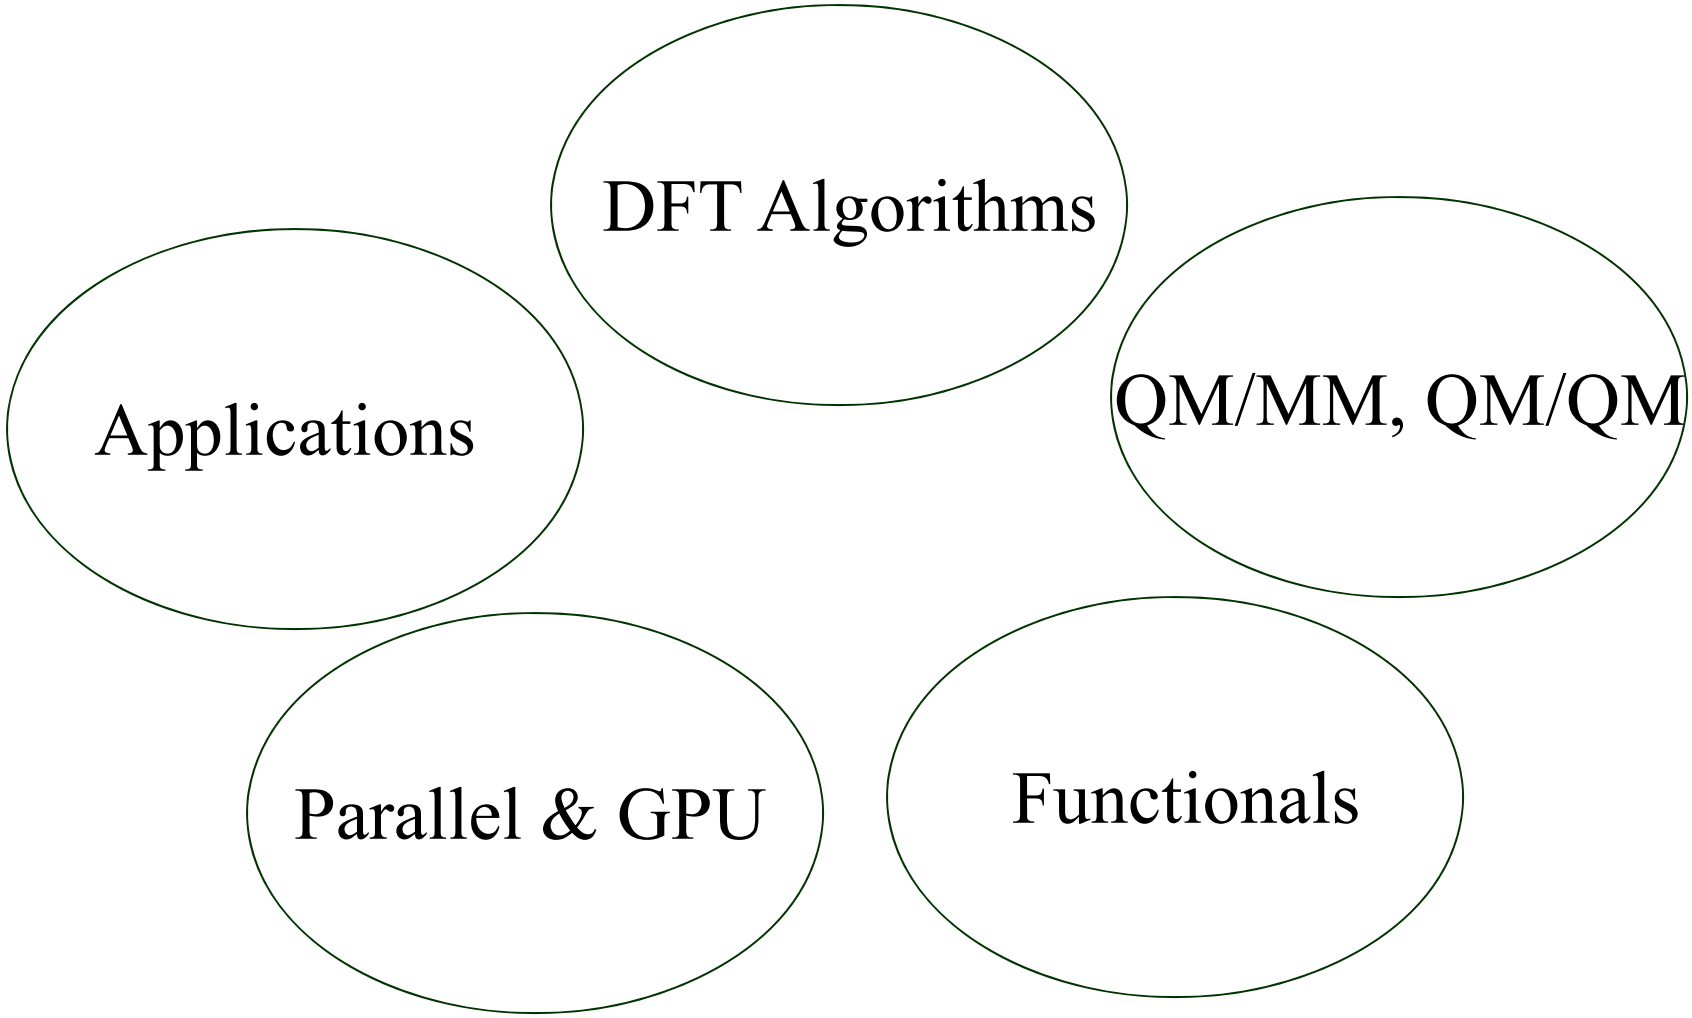

Research at Q-Chem Inc.

Funded by Small Business Research Grants from NIH

Q-CHEM™
A Quantum Leap Into the Future of Chemistry

People

Features Contributed Before

- 1.1 *Basis projection* (Jing)
- 1.2 DFT analytic frequencies (Jing)
- 1.2 *Parallel DFT energy and gradient* (Furlani, Dachsel)
- 2.0 *Parallel DFT frequencies with distributed memory* (Dachsel, Korambath, Brown, Furlani, Martin)
- 2.1 *Incremental DFT* (Brown)
- 3.0 *Fourier Transform Coulomb* (Fusti-Molnar)
- 3.2 *YinYang-Atom, a QM/MM interface* (Shao)

Features Contributed Before

- 4.0 *mrXC - multiresolution eXchange Correlation* (Brown, Fusti-Molnar, Russ, Chang)
- 4.0 *DFT dispersion with XDM* (Gan, Proynov)
- 4.0 *GPU application to DFT* (Gan, Shao)

Multiresolution exchange-correlation (mrXC)

Shawn Brown, Laszlo Fusti-Molnar, Nick Russ,
Chunming Chang

J. Chem. Phys., **124**, 094109 (2006)

Chem. Phys. Lett., **418**, 490 (2006)

mrXC: Motivation

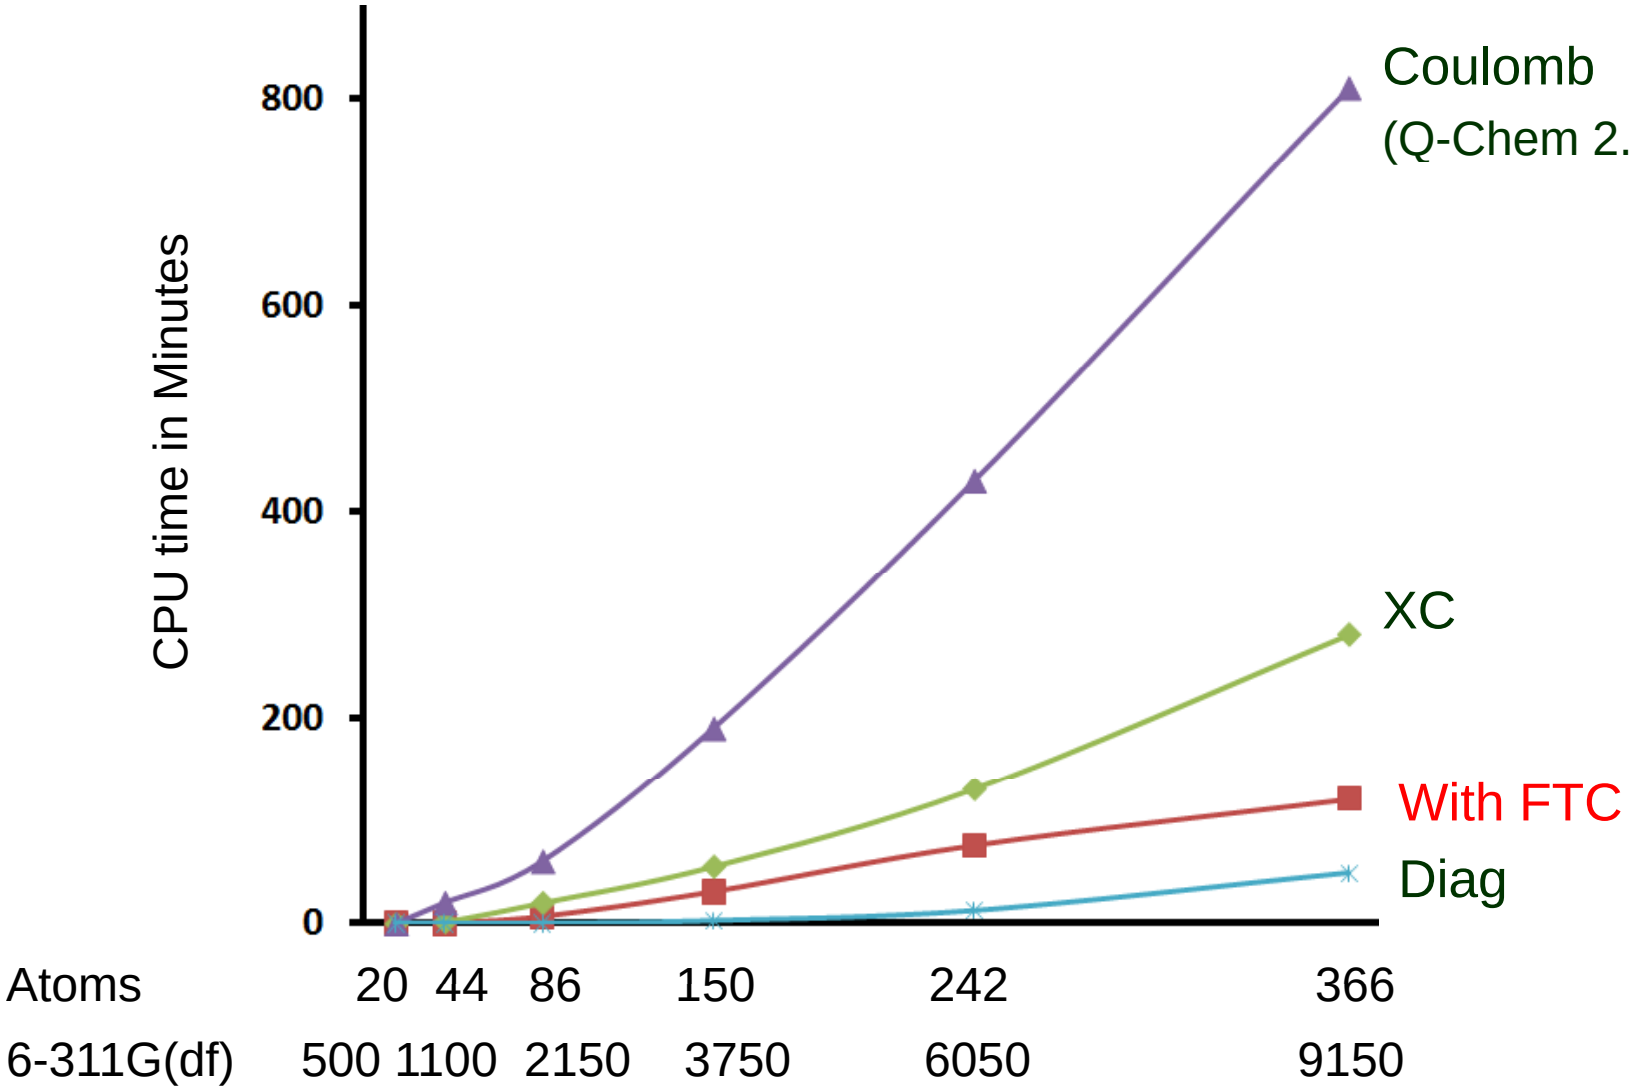

mrXC: the Grids Problem

Even spaced grid, e.g. CPMD

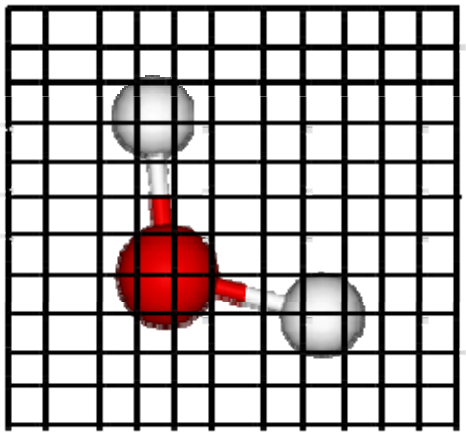

- Efficient for smooth, diffuse functions
- Cannot handle compact, core functions

mrXC: the Grid Problem

Atom-centered grid

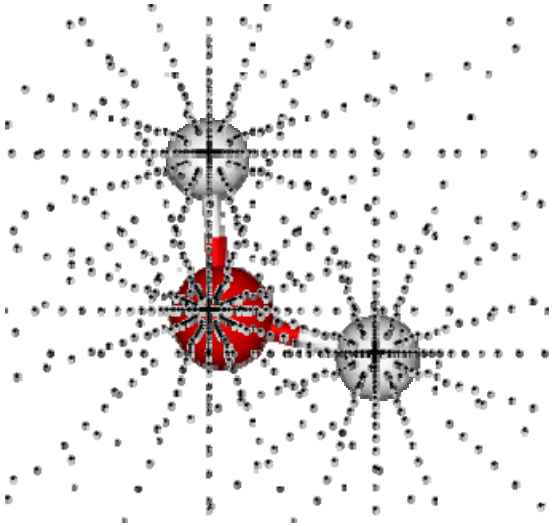

- Can handle both smooth and compact functions
- Not as efficient as the even-spaced grid for smooth functions

mrXC: the Grid Problem

Can we use the two grids in one space?

Not really

$$f(\rho) \neq f(\rho^S) + f(\rho^C)$$

But one can still try…

GAPW: divide the space into regions

$$f(\rho) \approx f(\rho^S) + f(\rho^C) - f(\tilde{\rho})$$

Error: 50μH/atom or more

mrXC: Our Solution

- Keep the integration formally on ACG grid
- Smooth part is actually done on cubic grid

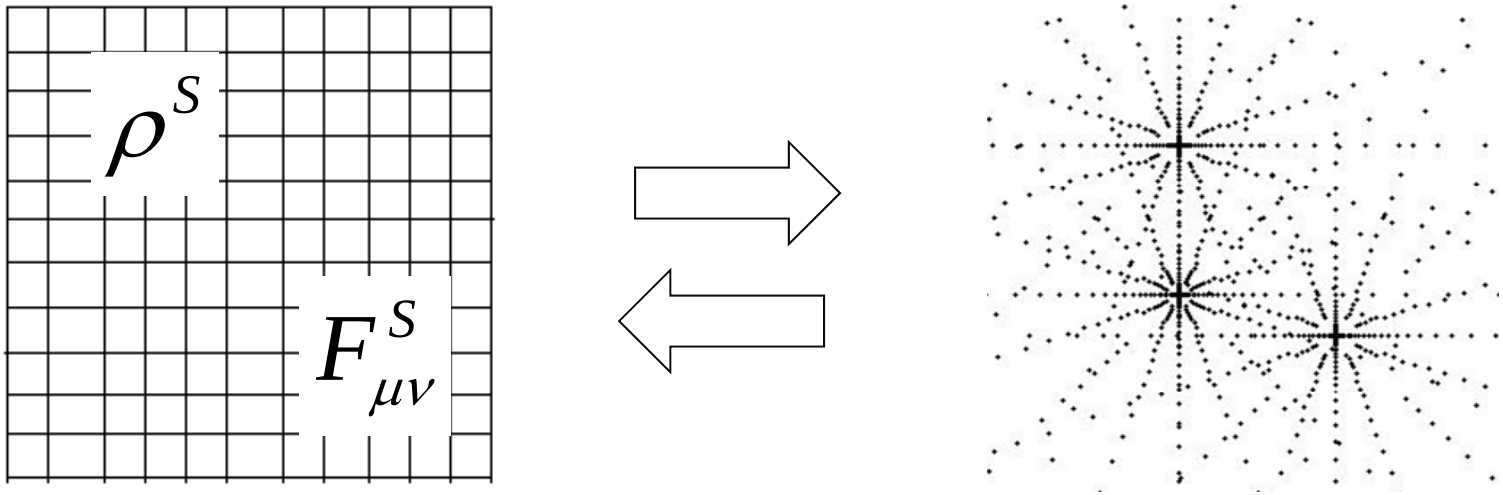

mrXC: Our Solution

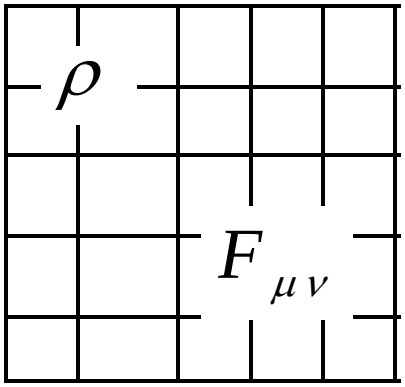

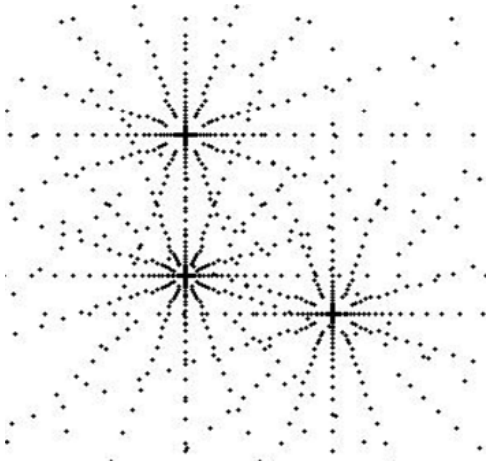

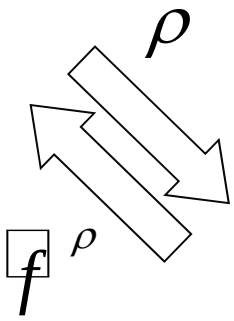

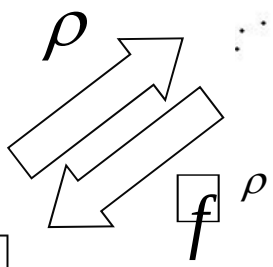

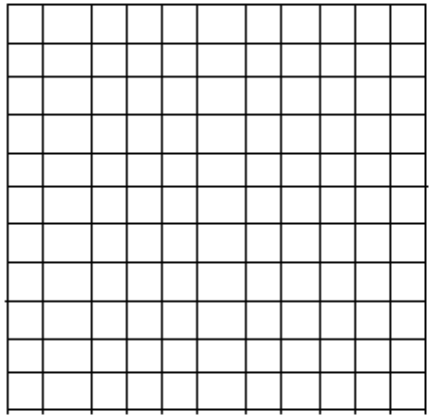

mrXC: Between Resolutions

Added benefit:

Smooth part free with FTC for Coulomb!

(at least for closed-shell)

mrXC: Our Solution

Details:

- grid densities: 4 and 6 in $bohr^{-1}$
- Local Interpolation scheme: Divided Difference
- Interpolation order: 6

mrXC: Example

Taxol: a natural cancer drug

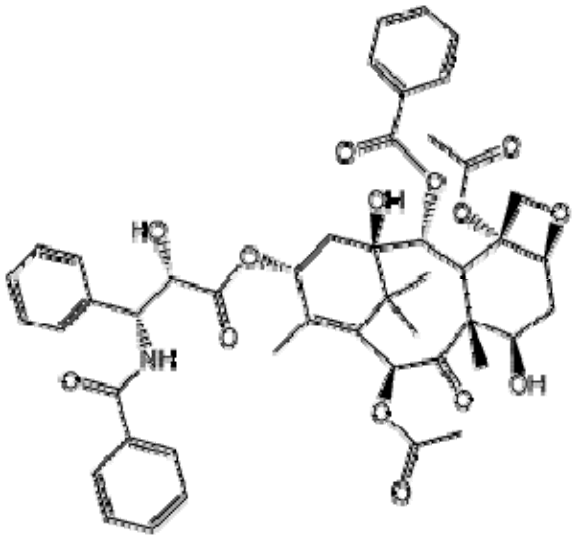

mrXC: Example

Basis sets	# of basis functions	Errors (μH/atom)	Speed-up (Time in sec.)	Speed-up with FTC[a]
6-31G(df,pd)	1925	0.03	3.9 (180)	5.0
cc-pvTZ	2574	0.1	5.8 (367)	9.6

Adding Dispersion to DFT

Zhengting Gan, Emil Proynov, Marek Freindorf, Thomas Furlani

Phys. Rev. A, **79**, 042510 (2009)

DFT: Accuracy Problem

Errors of calculated enthalpies of formation

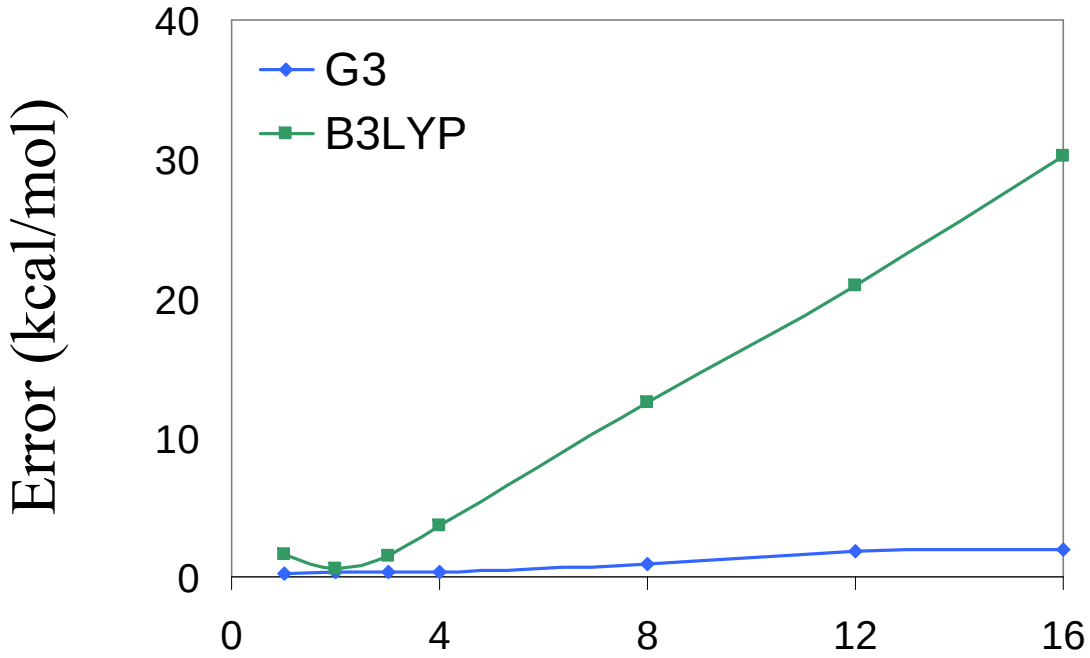

6-311++G(3df,2p)

DFT: Accuracy Problem

MP2/6-31+G*

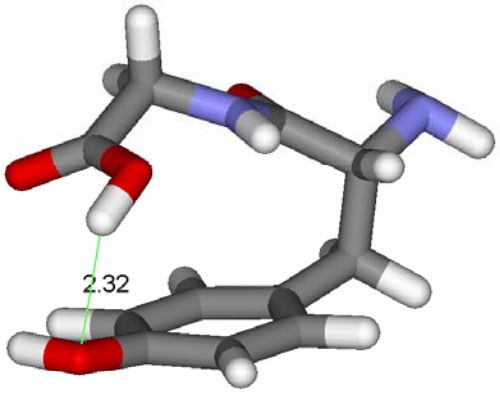

3 kcal/mol

B3LYP/6-31+G*

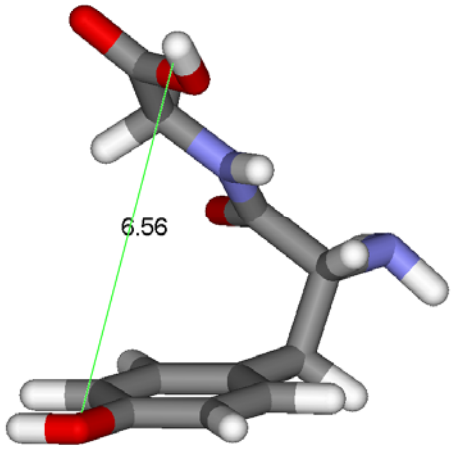

8 kcal/mol

Dispersion: XDM

Exchange Dipole Moment (XDM) Model

Johnson and Becke, JCP, **124**, 174104 (2006)

Exchange-hole

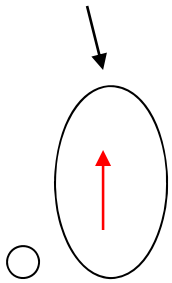

Dispersion attraction

Dispersion: XDM

$$C_{6,ij} = \frac{\alpha_i \alpha_j \left\langle d_X^2 \right\rangle_i \left\langle d_X^2 \right\rangle_j}{\alpha_i \left\langle d_X^2 \right\rangle_j + \alpha_j \left\langle d_X^2 \right\rangle_i}$$

- Electronic, NOT classical
- NO empirical parameter for C_6

Dispersion: XDM

- Critical quantity: the dipole of the exchange hole:

$$d_{X\sigma}(r) = -\int r' h_{X\sigma}(r, r') dr' - r$$

- We used Becke-Russel exchange hole (BR89)

Dispersion: XDM

Our implementation

- Analytical exchange-hole function
- Full SCF
- Analytical gradients
- Efficient: Little extra cost

Dispersion: Example

Errors of calculated enthalpies of formation

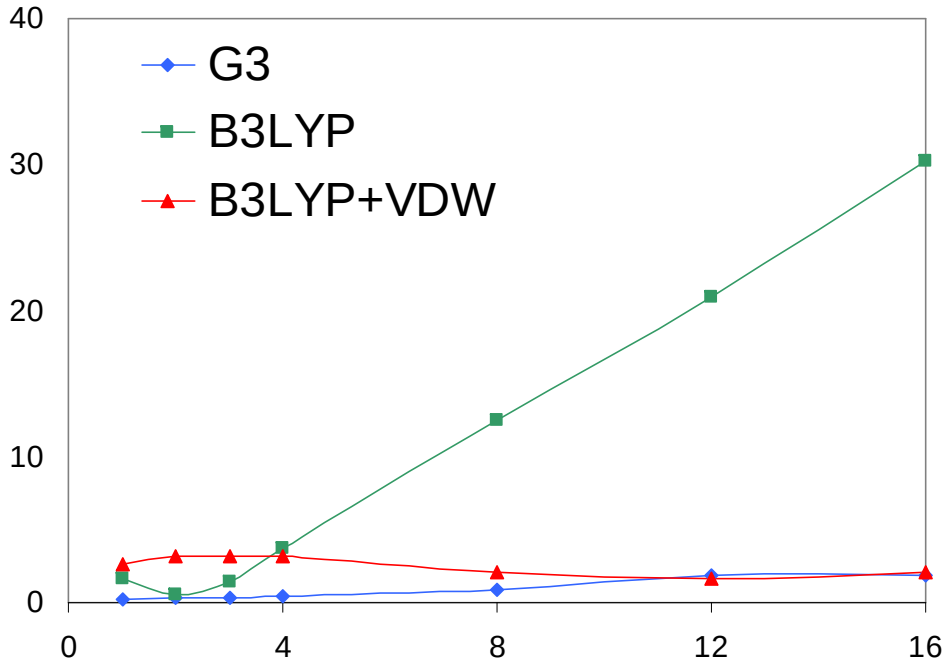

N-alkanes

6-311++G(3df,2p)

Dispersion: Example

B3LYP+XDM/6-31+G*

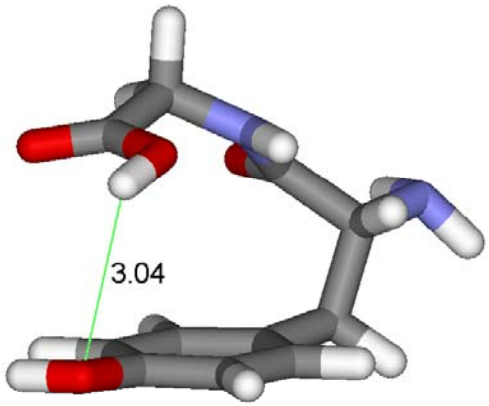

B3LYP/6-31+G*

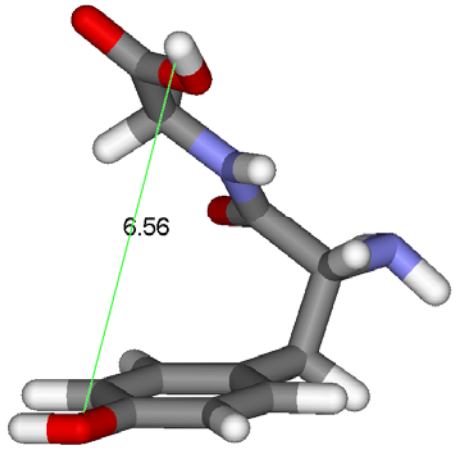

1.5 kcal/mol

8 kcal/mol

Dispersion: Example

Pyran as the model for the glycosidic linkage

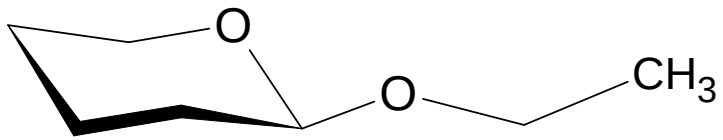

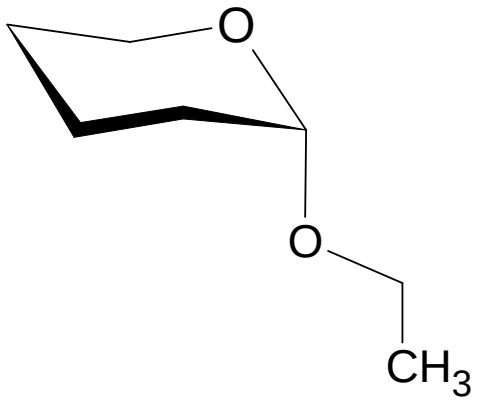

Dispersion: Example

ccpVTZ Basis	E_{eq}-E_{ax} (kcal/mol)
B3LYP	0.63 (-0.12)
CCSD(T) *Very expensive!*	1.47 (0.49)
B3LYP+vdW	1.07 (0.32)
Experiment	(0.47 ± 0.3)

Dispersion: Example

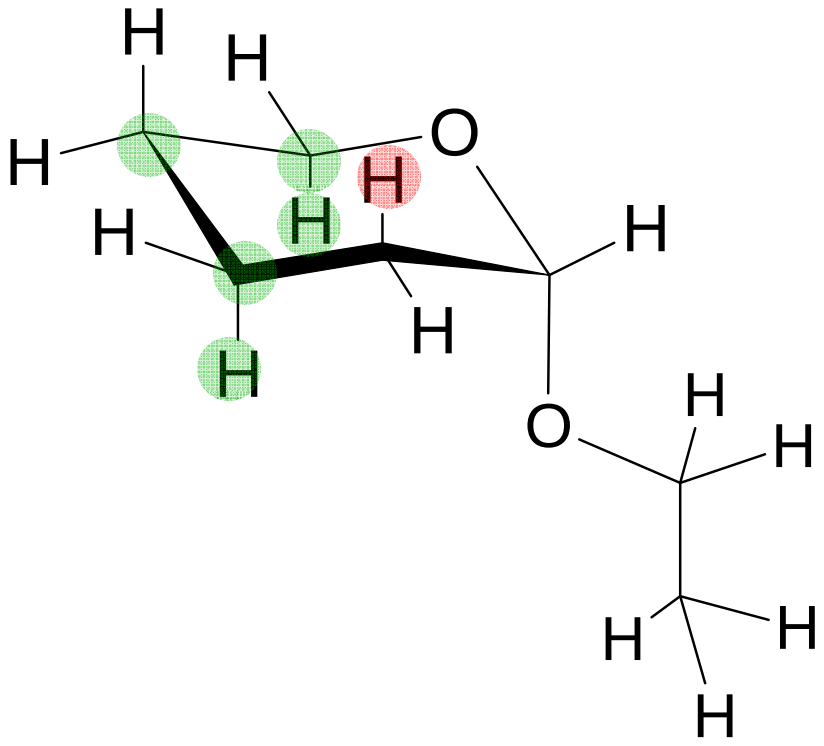

Accelerating DFT with GPU

Zhengting Gan, Yihan Shao
Roberto Olivares-Amaya, Alan Aspuru-Guzik (Harvard)

GPU and Hybrid Computing

- Each graphic card is a micro parallel computer
- High computing density: Teraflops in a desktop
- 64-bit versions are on the market
- BLAS interface available (CUBLAS)
- *Difficulty*: Slow memory, small register

GPU Application to DFT

Standard XC numerical integration uses BLAS1

$$\rho_i = \sum_{\mu\nu} P_{\mu\nu}\phi_{\mu,i}\phi_{\nu,i} \qquad F_{\mu\nu} = \sum_i w_i f_i^{\rho}\phi_{\mu,i}\phi_{\nu,i}$$

- Screen the basis functions
- Screen the basis function pairs
  - *Problem: Use BLAS1 calls only*
- Calculate the basis function values
  - *Problem: too many intermediates (> 50)*

GPU Application to DFT

Solution 1: Get rid of screening of basis pairs

- Trade flops with cache friendliness

GPU Application to DFT

Solution 2: For the calculation of basis functions

- Save low angular components in registers
- Calculate high angular ones in the fly

GPU Application to DFT

How GPU is combined with multicore and multinodes?

- Each GPU or core is run by a separate thread
- Threads from different nodes are put in one resource pool
- Dynamic load balancing: Job tasks are distributed according to the availabilities of the threads

Hybrid Computing

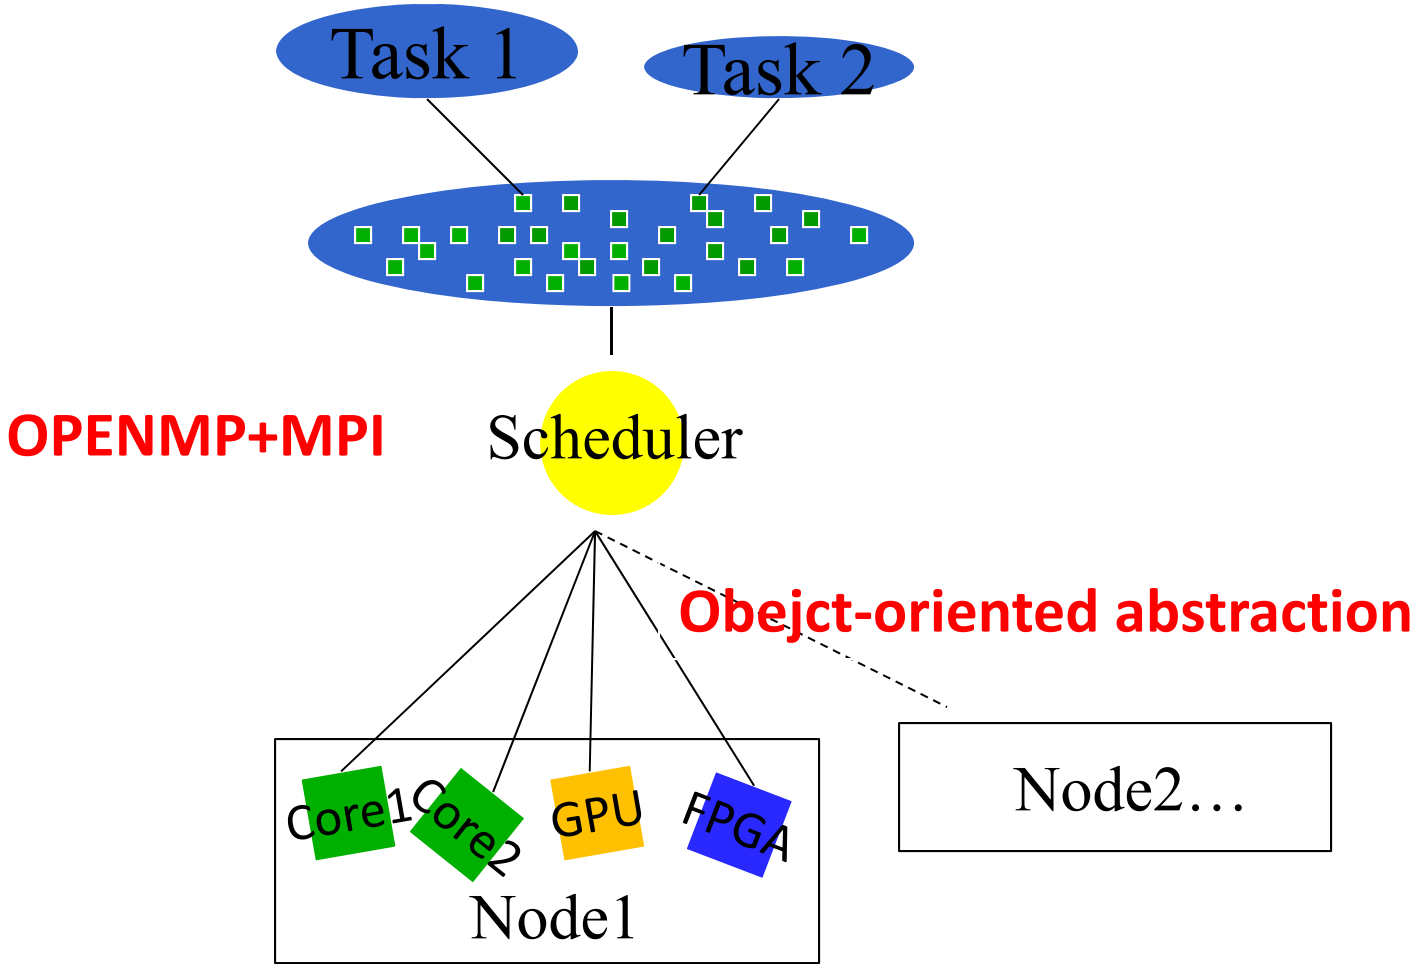

TOOM: Task-based Object-oriented with Openmp and Mpi

GPU Application to DFT

GPU used

- Tesla S1070
- 960 cores @ 1.45GHz
- Peak performance:
  - 312GFlops in double precision
  - 4TFlops in single precsion

CPU used

- Xeon 2.8GHz
- Peak performance: 44.8GGlops

Hybrid Computing

Benchmark calculations using hybrid computing

Taxol(C47H51NO14)
aug-cc-pVTZ, 4025 basis
Lebedev grid (75,302)

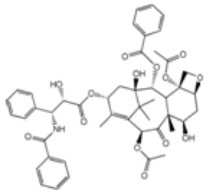

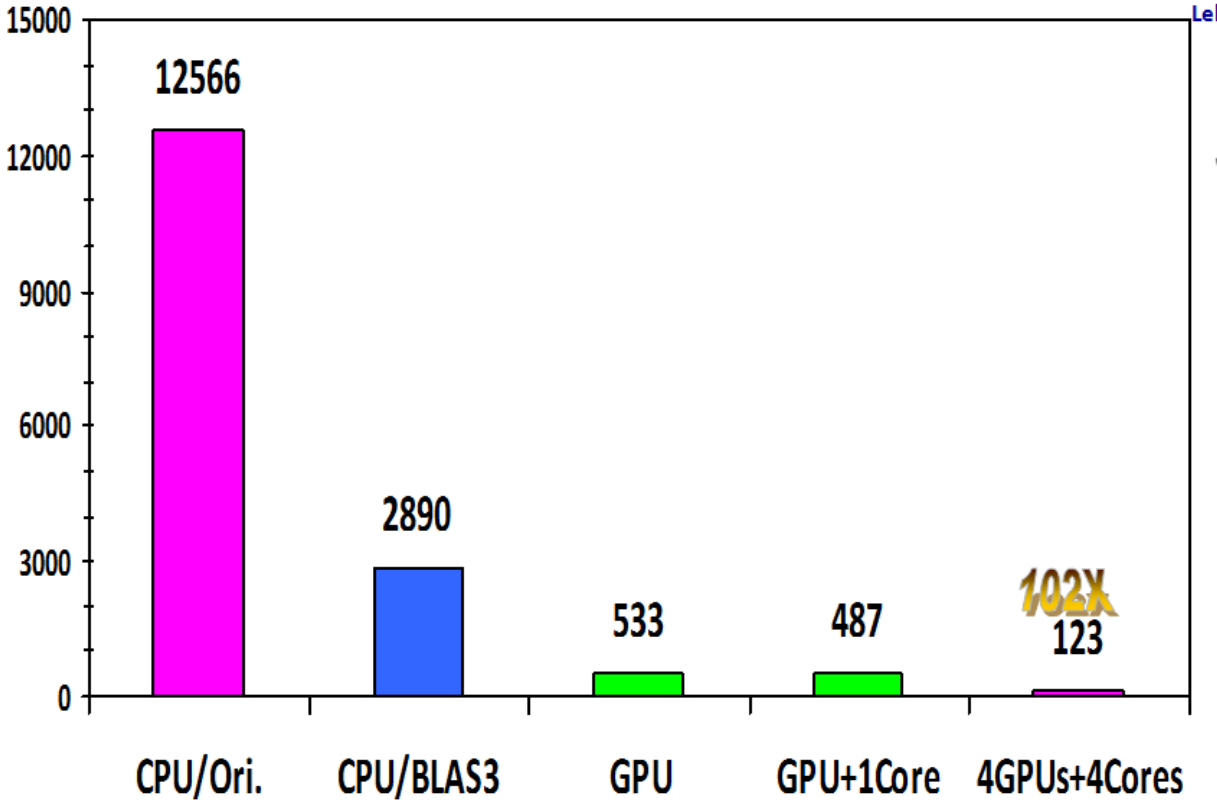

Calculation run on Dual socket Xeon E5462 2.8GHz with S1070 GPU

Some Business

www.q-chem.com

People

Current:

Yihan Shao
Emil Proynov
Zhengting Gan
Nick Russ
Chunmin Chang
ZitaAnn Berry
Hilary Pople
Jaime Martell

Past:

Shawn T. Brown @ PSC
Holger Dachsel @ Julish
Fusti-Molnar @ Florida
Prakashan Korambath @ UCLA
Weimin Zhang @ ScottTrade

Collaborators:

Tom Furlani and Marek Freindorf
at Univ. Buffalo