

New frontiers in quantum chemistry using supercomputers

Jeff Hammond

Leadership Computing Facility
Argonne National Laboratory

4 May 2011



Atomistic simulation in chemistry

- 1 classical molecular dynamics (MD) with empirical potentials
- 2 ab initio molecular dynamics based upon density-function theory (DFT)
- 3 quantum chemistry with wavefunctions
e.g. perturbation theory (PT), Coupled-Cluster (CC) or Quantum Monte Carlo (QMC).

Classical molecular dynamics

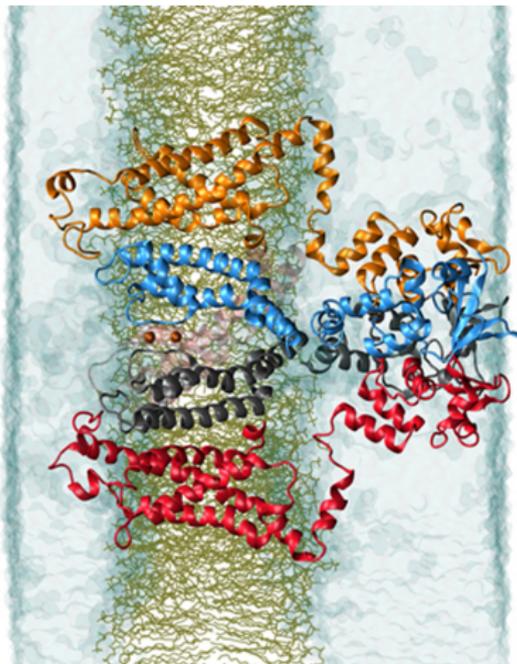


Image courtesy of Benoît Roux via ALCF.

- Solves Newton's equations of motion with empirical terms and classical electrostatics.
- Size: 100K-10M atoms
- Time: 1-10 ns/day
- Scaling: $\sim N_{atoms}$

Data from K. Schulten, et al. "Biomolecular modeling in the era of petascale computing." In D. Bader, ed., *Petascale Computing: Algorithms and Applications*.

Car-Parrinello molecular dynamics

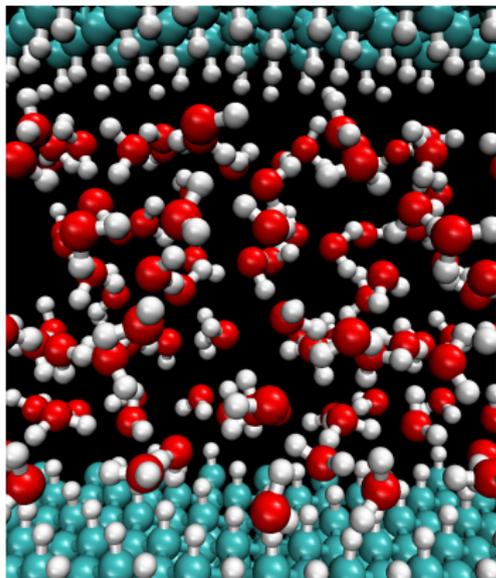
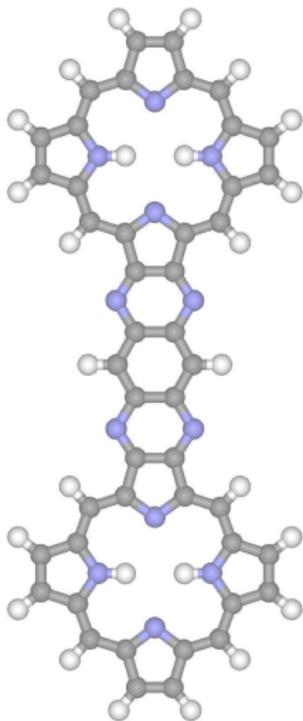


Image courtesy of Giulia Galli via ALCF.

- Forces obtained from solving an approximate single-particle Schrödinger equation; time-propagation via Lagrangian approach.
- Size: 100-1000 atoms
- Time: 0.01-1 ps/day
- Scaling: $\sim N_{el}^x$ ($x=1-3$)

F. Gygi, *IBM J. Res. Dev.* **52**, 137 (2008); E. J. Bylaska et al. *J. Phys.: Conf. Ser.* **180**, 012028 (2009).

Wavefunction theory



- MP2 is second-order PT and is accurate via magical cancellation of error.
- CC is infinite-order solution to many-body Schrödinger equation truncated via clusters.
- QMC is Monte Carlo integration applied to the Schrödinger equation.
- Size: 10-100 atoms, maybe 100-1000 atoms with MP2.
- Time: N/A
- Scaling: $\sim N_{bf}^x$ ($x=2-7$)

Image courtesy of Karol Kowalski and Niri Govind.

Quantum chemistry — standard model

- 1 Separate molecule(s) from environment (closed to both matter and energy)
- 2 Ignore relativity, QED, spin-orbit coupling
- 3 Separate electronic and nuclear degrees of freedom

→ non-relativistic electronic Schrödinger equation in a vacuum at zero temperature.

$$\begin{aligned}\hat{H} &= \hat{T}_{el} + \hat{V}_{el-nuc} + \hat{V}_{el-el} \\ \hat{H} &= -\frac{1}{2} \sum_{i=1}^M \nabla_i^2 + \sum_{n=1}^N \sum_{i=1}^M \frac{Z_n}{R_{ni}} + \sum_{i<j}^M \frac{1}{r_{ij}}\end{aligned}$$

$$\Psi(\mathbf{x}_1, \dots, \mathbf{x}_n, \mathbf{x}_{n+1}, \dots, \mathbf{x}_N) = -\Psi(\mathbf{x}_1, \dots, \mathbf{x}_{n+1}, \mathbf{x}_n, \dots, \mathbf{x}_N)$$

The electron coordinates include both space (r) and spin (σ). We will integrate-out spin wherever possible.

Quantum chemistry — standard model

Wavefunction antisymmetry is enforced by expanding in determinants, which we now capture in second quantization.

- 1 project physical operators (e.g. Coulomb) into one-electron basis — usually atom-center Gaussians
- 2 generate mean-field reference and expand many-body wavefunction in terms of excitations out of that reference

→ Full configuration-interaction (FCI) ansatz.

- 1 truncate exponentially-growing FCI ansatz (CI=linear generator, CC=exponential generator)
- 2 solve CC (or CI) iteratively
- 3 add more correlation via perturbation theory

→ CCSD(T), as one example.

Correct for missing physics using perturbation theory (a posteriori error correction) or mixed (e.g. QM/MM) formalism:

- 1 relativistic corrections
- 2 non-adiabatic corrections
- 3 solvent corrections
- 4 open BC corrections (less common)

Coupled-cluster theory

$$|\Psi_{CC}\rangle = \exp(T)|\Psi_{HF}\rangle$$

$$T = T_1 + T_2 + \dots + T_n \quad (n \ll N)$$

$$T_1 = \sum_{ia} t_i^a \hat{a}_a^\dagger \hat{a}_i$$

$$T_2 = \sum_{ijab} t_{ij}^{ab} \hat{a}_a^\dagger \hat{a}_b^\dagger \hat{a}_j \hat{a}_i$$

$$\begin{aligned} |\Psi_{CCD}\rangle &= \exp(T_2)|\Psi_{HF}\rangle \\ &= (1 + T_2 + T_2^2)|\Psi_{HF}\rangle \end{aligned}$$

$$\begin{aligned} |\Psi_{CCSD}\rangle &= \exp(T_1 + T_2)|\Psi_{HF}\rangle \\ &= (1 + T_1 + \dots + T_1^4 + T_2 + T_2^2 + T_1 T_2 + T_1^2 T_2)|\Psi_{HF}\rangle \end{aligned}$$

Coupled cluster (CCD) implementation

$$R_{ij}^{ab} = V_{ij}^{ab} + P(ia, jb) \left[T_{ij}^{ae} I_e^b - T_{im}^{ab} I_j^m + \frac{1}{2} V_{ef}^{ab} T_{ij}^{ef} + \frac{1}{2} T_{mn}^{ab} I_{ij}^{mn} - T_{mj}^{ae} I_{ie}^{mb} - I_{ie}^{ma} T_{mj}^{eb} + (2T_{mi}^{ea} - T_{im}^{ea}) I_{ej}^{mb} \right]$$

$$I_b^a = (-2V_{eb}^{mn} + V_{be}^{mn}) T_{mn}^{ea}$$

$$I_j^i = (2V_{ef}^{mi} - V_{ef}^{im}) T_{mj}^{ef}$$

$$I_{kl}^{ij} = V_{kl}^{ij} + V_{ef}^{ij} T_{kl}^{ef}$$

$$I_{jb}^{ia} = V_{jb}^{ia} - \frac{1}{2} V_{eb}^{im} T_{jm}^{ea}$$

$$I_{bj}^{ia} = V_{bj}^{ia} + V_{be}^{im} (T_{mj}^{ea} - \frac{1}{2} T_{mj}^{ae}) - \frac{1}{2} V_{be}^{mi} T_{mj}^{ae}$$

Tensor contractions currently implemented as DGEMM plus PERMUTE to align indices (alternatives in-progress).

Motivation for HPC

- Electronic excited-states and electric-field perturbations push the limits of conventional approximations in DFT and are outside the scope of classical methods.
- Interesting chemical processes in biology and material science require model systems too large for a conventional computational resources.
- Answering many chemical questions requires large data sets which cannot be obtained in a reasonable amount time if done sequentially.

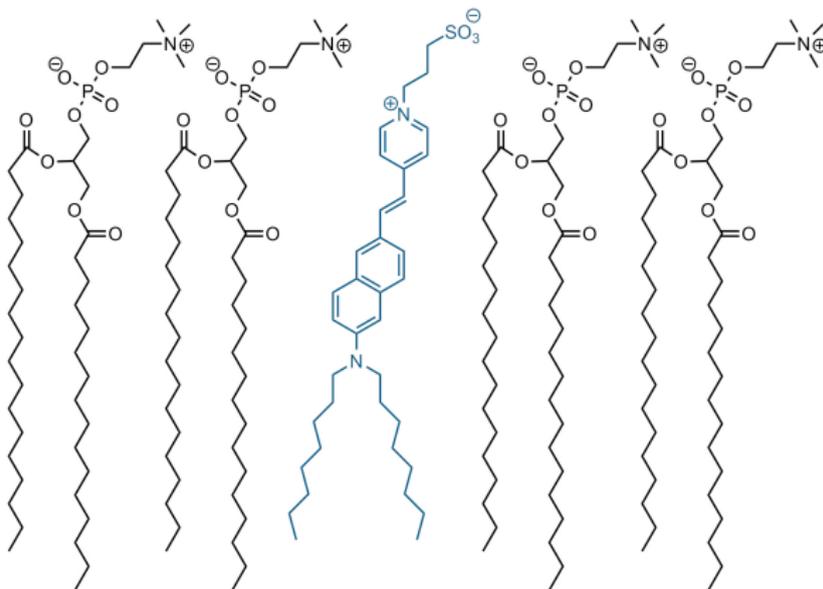
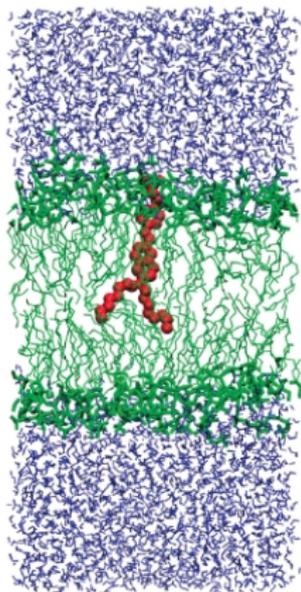
Electronic excited-states in biology

Joint work with Benoît Roux (UC/ANL) and Karol Kowalski (PNNL).

Optical potentiometric probes have become important tools in electrophysiology. These organic molecules display spectroscopic responses to membrane potential and have been used for the study and characterization of model membranes, nerve and muscle tissues, organelles, microorganisms, and red blood cells. They can often be used in place of conventional microelectrodes and lend themselves to many systems not accessible to microelectrodes.

E. Fluhler, V. G. Burnham, L. M. Loew, *Biochemistry* **24** , 5749 (1985). "Spectra, membrane binding, and potentiometric responses of new charge shift probes."

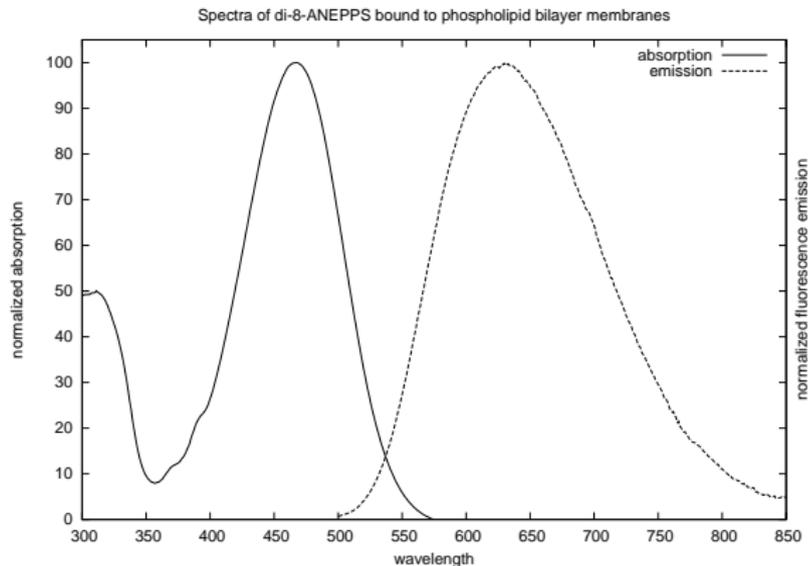
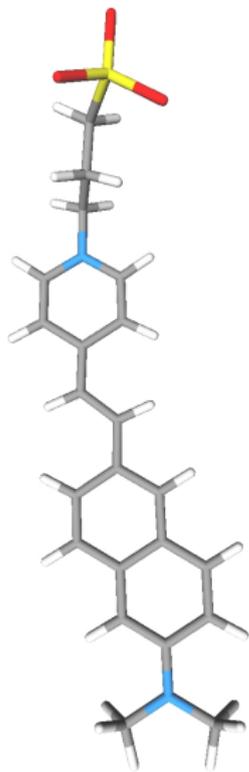
Membrane configuration of di-8-ANEPPS



C. F. Rusu, H. Lanig, O. G. Othersen, C. Kryschik, and T. Clark, *J. Phys. Chem. B* **112**, 2445 (2008).

ANEPPS model structure

peaks = 2.655 eV (3.987 eV) and 1.965 eV



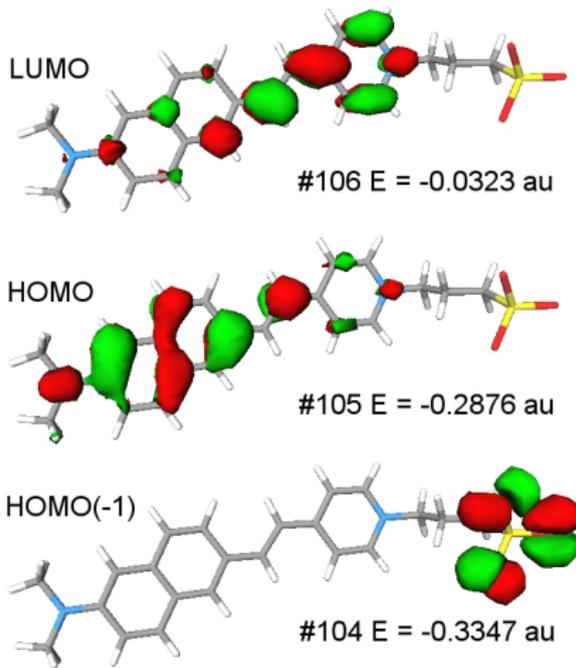
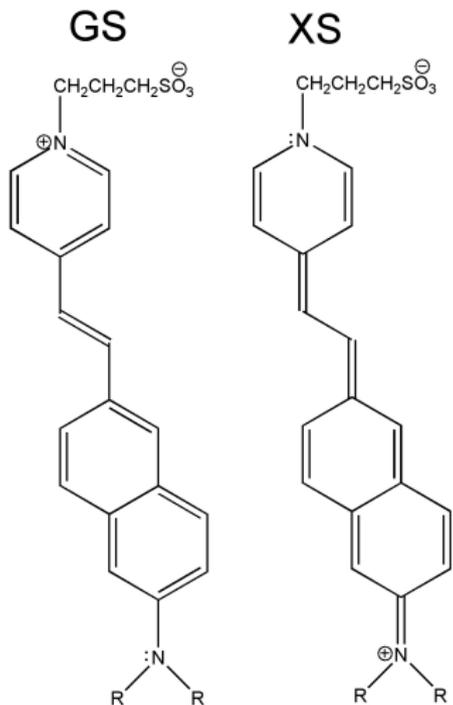
<http://probes.invitrogen.com/media/spectra/data/3167lip.txt>

Computing the spectrum I

Root	au	eV	nm	OS
B3LYP/cc-pVDZ	0.002	0.06	19758.1	0.001
B3LYP/aug-cc-pVDZ	0.016	0.44	2799.2	0.000
PBE0/aug-cc-pVDZ	0.026	0.70	1773.5	0.000
BH&H/aug-cc-pVDZ	0.091	2.47	501.7	0.000
TDHF/aug-cc-pVDZ	0.124	3.38	366.7	1.731
CIS/aug-cc-pVDZ	0.132	3.59	345.8	1.949

The basis set dependence is an illusion.

What are the electrons doing?



Computing the spectrum II

Basis	τ	eV	au	nm
CC2/6-31G*	1.5	1.815	0.067	683.3
CCSD/6-31G*	0.5	3.629	0.133	341.6
CCSD/6-31G*	0.7	3.231	0.119	383.8
CCSD/6-31G*	1.5	2.984	0.110	415.5
CCSD/6-31G*	2.0	2.962	0.109	418.6
CCSD/6-31G*	∞	2.968	0.109	417.7
CCSD/cc-pVDZ	∞	2.945	0.108	421.0
Experiment		2.655	0.098	467

For $\tau > 1.5$, X_1 is dominated by one amplitude.

$\|X_2\|$ is small throughout.

Accurate many-body methods

Method	Basis	τ	eV	au	nm
EOM-CCSD	6-31G*	0.5	3.629	0.133	341.6
	6-31G*	0.7	3.231	0.119	383.8
	6-31G*	1.5	2.980	0.110	416.1
CR-EOM-CCSD(T)	6-31G*	0.5	3.590	0.132	345.4
	6-31G*	0.7	3.150	0.116	393.6
	6-31G*	1.5	2.810	0.103	441.2
Experiment			2.655	0.098	467

Active-space approach:

Only orbitals with $-\tau < \epsilon < \tau$ are correlated.

$\tau=1.5$ corresponds to frozen-core approximation.

NWChem implementation (TCE)

Timings on 256 nodes of Chinook:

Procedure	wall time (s)
SCF total time	57
four-index transformation	192
one CCSD iteration	157
one EOM-CCSD iteration	252
CR-EOM-CCSD(T) evaluation	6301
Total time	12510

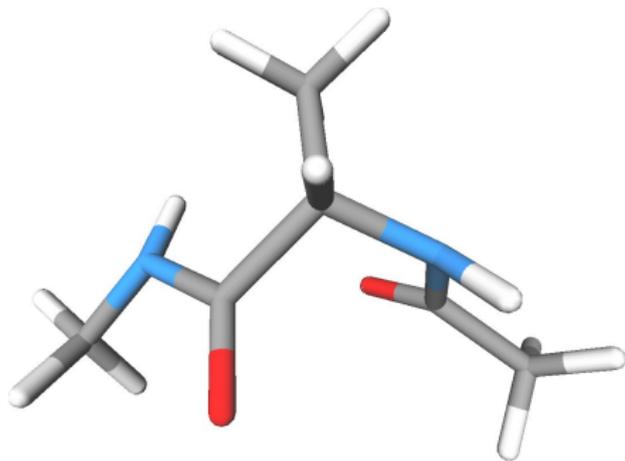
Even though this calculation is trivial with NWChem, it is still impossible with single-node codes because of the memory wall.

Karol is using 25K+ cores of Hopper2 for coronene+porphyrin.

Force-fields from first-principles?

Joint work with Karl Freed (UC), Benoît Roux (UC/ANL), Alex MacKerell (Maryland)

The protein prototype — dialanine



- Debatable if dialanine represents the real torsional potential.
- Many FF potentials use MP2 dialanine results.
- Useful for calibrating methods without pollution of cooperative effects.
- Computationally tractable for CCSD(T) (whole ϕ - ψ map).

Evaluating models with CCSD(T)

- Used OPLS-AA geometries to prevent bias.
- 6-311++G** basis set (aug-cc-pVTZ desirable).
- Difference between 6-31+G* and 6-311++G**:
SCF=1.18, MP2=2.48 (MUE in kJ/mol).
- Difference between 6-311++G** and aug-cc-pVTZ:
SCF=0.89, MP2=1.70 (MUE in kJ/mol).
- 350 configurations (30° grid everywhere, 10° in basins).
- CCSD(T)/6-311++G** takes approximately 1 hour per job on 64 nodes (whole set could run in an afternoon).

Evaluating models with CCSD(T)

Method	MUE	Max	Method	MUE	Max
MP2	0.436	1.304	M05	1.551	8.389
CCSD	0.577	1.426	TPSS+D	1.593	9.258
B2PLYP	0.913	4.690	FT97	1.711	11.047
M06	1.137	4.874	CAMB3LYP	1.747	6.268
Becke97+D	1.177	5.981	M06-2X	1.757	5.812
Becke98	1.287	7.526	BB1K	1.773	7.310
TPSS	1.312	10.691	B2LYP	1.913	6.514
B3LYP+D	1.327	6.379	HCTH120	2.119	10.141
TPSSh	1.330	9.525	BOP	2.614	9.118
M06-L	1.378	6.657	M06-HF	2.884	12.286
Becke97	1.391	7.486	SCF	3.066	11.076
PBE+D	1.404	9.812	HCTH407	3.168	9.678
X3LYP	1.430	7.747	HCTH	3.330	9.788
B3LYP	1.456	7.884	CAMPBE0	3.348	10.676
PBE0	1.506	8.041			

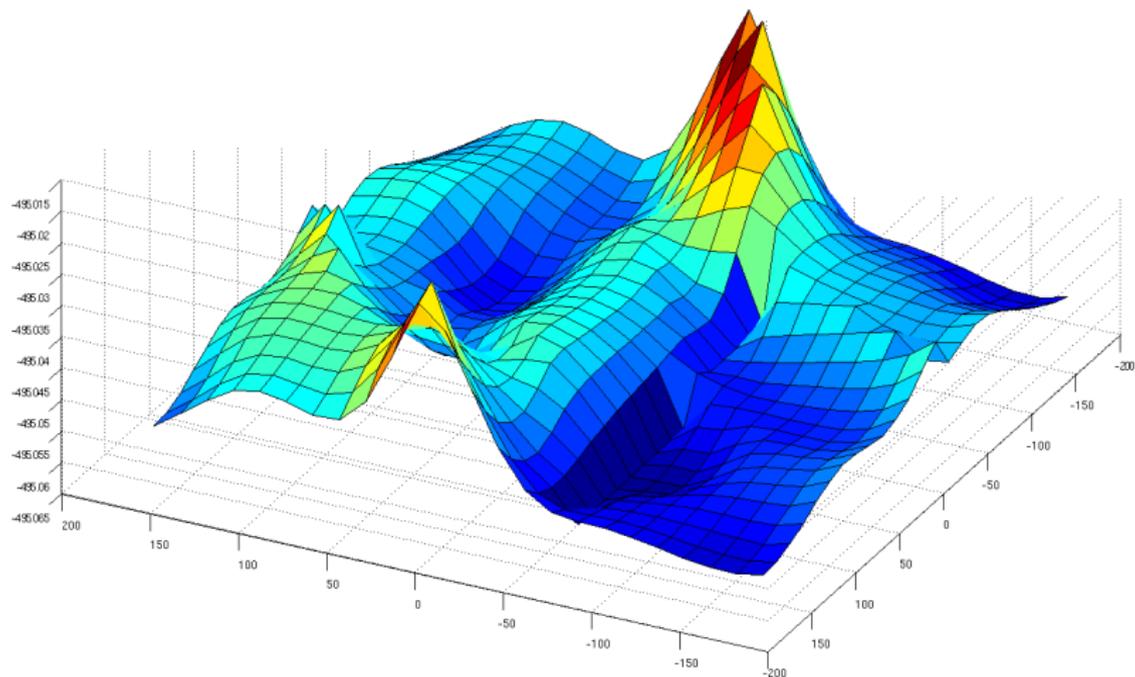
(All in NWChem 6.0)

Observations

- Justified using MP2 for fitting torsional parameters.
- Approximate functionals are getting better with time.
- DFT+D improves results in most cases.
- Unlikely that a density functional better than MP2 exists.

Now that we know MP2 is the best cheap method. . .

CCSD(T)/cc-pVTZ energies at MP2/cc-pVTZ geometries.



Exascale computational chemistry

State of HPC for Chemistry

Both classical and ab initio molecular dynamics have essentially reached algorithmic maturity. Most research is fighting Amdahl's law and related concepts (FFT does not scale), e.g. DEShaw has turned classical molecular dynamics into an engineering problem.

Quantum many-body methods are far from algorithmic maturity because they have been constrained to tiny systems so the N -body problem is hidden behind dense linear algebra.

Dense linear algebra is great for Gordon Bell Prizes but terrible for science.

PETAFLOPS II



February 15-19, 1999 Doubletree Hotel Santa Barbara, CA

2nd Conference on Enabling Technologies for Peta(f)ops Computing

Call for Participation and Papers

February 15 - 19, 1999

Doubletree Hotel

Santa Barbara, California

Conference Chair: Paul Messina

Caltech Program Chair: Thomas Sterling, Caltech/JPL

Steering Committee Chair: Paul H. Smith, DOE

Sponsors: DARPA, NASA, NSF, DOE, NSA

The 2nd Conference on Enabling Technologies for Peta(f)ops Computing is the first major open forum to treat the diversity of technical issues through a series of in-depth workshops and sponsored studies conducted to explore the factors that will determine the ultimate path to realizing such capabilities. The conference will focus on understanding of Petaflops scale computing approaches and determine directions for future research leading to practical Petaflops performance. The conference will address a wide range of issues and foster detailed discussion across conventional discipline boundaries. The conference will engage the interests of researchers in areas associated with petaflops scale computing and beyond include but are not limited to:

Is this exciting? Yes!

- s **Solution of longstanding problems**
 - Detailed chemical reaction dynamics in solution
 - Molecular level environmental chemistry
 - Ab initio design of catalysts
 - Ab initio molecular dynamics
- s **New “Grand Challenges”**
- s **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?

Did we make good on our promises last time?

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

ARTICLES

Kemp elimination catalysts by computational enzyme design

Daniela Röthlisberger^{1*}, Olga Khersonsky^{4*}, Andrew M. Wollacott^{1*}, Lin Jiang^{1,2}, Jason DeChancie⁶, Jamie Betker³, Jasmine L. Gallaher³, Eric A. Althoff¹, Alexandre Zanghellini^{1,2}, Orly Dym⁵, Shira Albeck⁵, Kendall N. Houk⁶, Dan S. Tawfik⁴ & David Baker^{1,2,3}

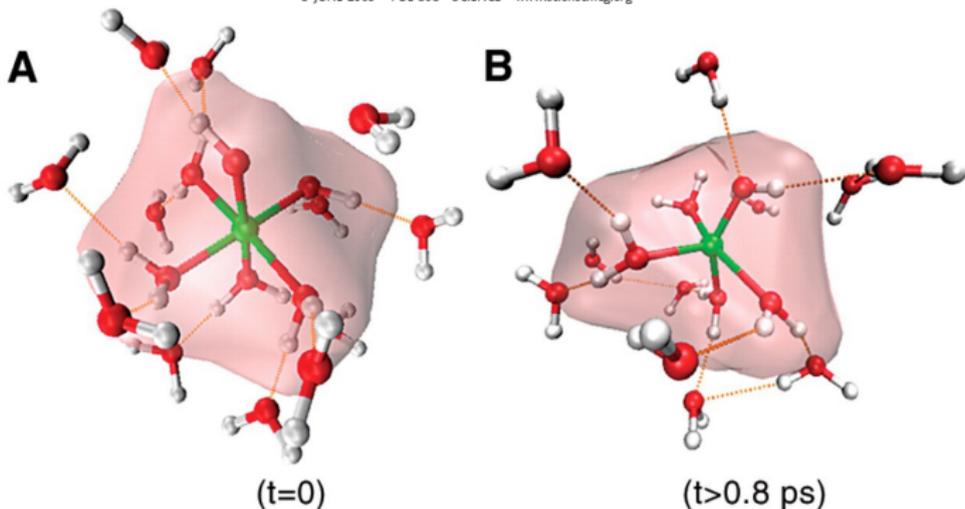
The design of new enzymes for reactions not catalysed by naturally occurring biocatalysts is a challenge for protein engineering and is a critical test of our understanding of enzyme catalysis. Here we describe the computational design of eight enzymes that use two different catalytic motifs to catalyse the Kemp elimination—a model reaction for proton transfer from carbon—with measured rate enhancements of up to 10^5 and multiple turnovers. Mutational analysis confirms that catalysis depends on the computationally designed active sites, and a high-resolution crystal structure suggests that the designs have close to atomic accuracy. Application of *in vitro* evolution to enhance the computational designs produced a >200-fold increase in k_{cat}/K_m (k_{cat}/K_m of $2,600 \text{ M}^{-1}\text{s}^{-1}$ and $k_{\text{cat}}/k_{\text{uncat}}$ of $>10^6$). These results demonstrate the power of combining computational protein design with directed evolution for creating new enzymes, and we anticipate the creation of a wide range of useful new catalysts in the future.

Kinetic Evidence for Five-Coordination in $\text{AlOH}(\text{aq})^{2+}$ Ion

Thomas W. Swaddle,¹ Jörgen Rosenqvist,² Ping Yu,³ Eric Bylaska,⁶
Brian L. Phillips,⁷ William H. Casey^{2,4,5*}

Trivalent aluminum ions are important in natural bodies of water, but the structure of their coordination shell is a complex unsolved problem. In strong acid ($\text{pH} < 3.0$), Al^{3+} exists almost entirely as the octahedral $\text{Al}(\text{H}_2\text{O})_6^{3+}$ ion, whereas in basic conditions ($\text{pH} > 7$), a tetrahedral $\text{Al}(\text{OH})_4^-$ structure prevails. In the biochemically and geochemically critical pH range of 4.3 to 7.0, the ion structures are less clear. Other hydrolytic species, such as $\text{AlOH}(\text{aq})^{2+}$, exist and are traditionally assumed to be hexacoordinate. We show, however, that the kinetics of proton and water exchange on aqueous Al^{3+} , coupled with Car-Parrinello simulations, support a five-coordinate $\text{Al}(\text{H}_2\text{O})_5\text{OH}^{2+}$ ion as the predominant form of $\text{AlOH}(\text{aq})^{2+}$ under ambient conditions. This result contrasts Al^{3+} with other trivalent metal aqua ions, for which there is no evidence for stable pentacoordinate hydrolysis products.

3 JUNE 2005 VOL 308 SCIENCE www.sciencemag.org



Hard scaling challenges for *ab initio* molecular dynamics capabilities in NWChem: Using 100,000 CPUs per second

Eric J. Bylaska¹, Kevin Glass¹, Doug Baxter¹, Scott B. Baden² and John H. Weare³

¹ Environmental Molecular Sciences Laboratory, Pacific Northwest National Laboratory, P.O. Box 999, Richland, WA 99354

² Department of Computer Science and Engineering, University of California, San Diego
9500 Gilman Drive, # 0404, La Jolla, CA 92093-0404

³ Department of Chemistry and Biochemistry, University of California, San Diego
9500 Gilman Drive, # 0303, La Jolla, CA 92093-0303

Email: Eric.Bylaska@pnl.gov

Abstract. An overview of the parallel algorithms for *ab initio* molecular dynamics (AIMD) used in the NWChem program package is presented, including recent developments for computing exact exchange. These algorithms make use of a two-dimensional processor geometry proposed by Gygi *et al.* for use in AIMD algorithms. Using this strategy, a highly scalable algorithm for exact exchange has been developed and incorporated into AIMD. This new algorithm for exact exchange employs an incomplete butterfly to overcome the bottleneck associated with exact exchange term, and it makes judicious use of data replication. Initial testing has shown that this algorithm can scale to over 20,000 CPUs even for a modest size simulation.

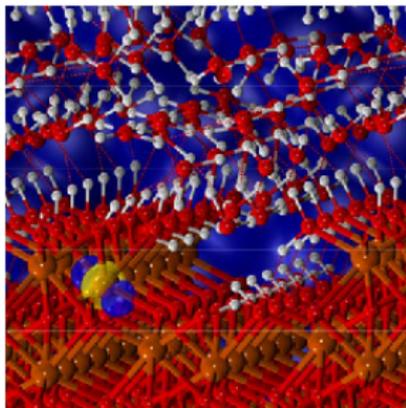
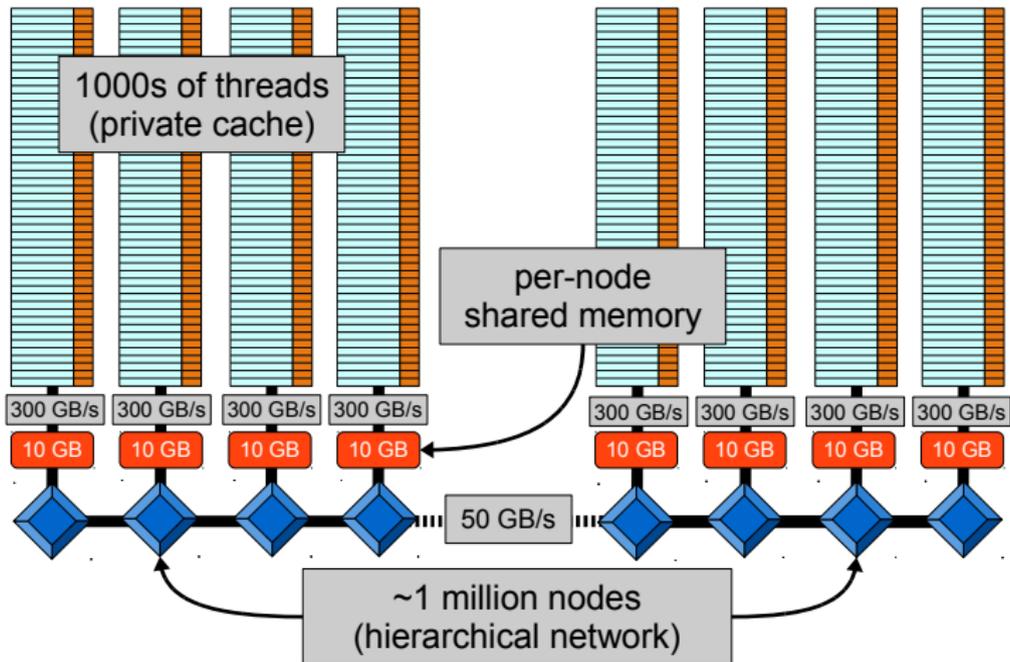


Figure 2. Illustration of a localized electron (i.e. polaron) on the surface of hematite calculated with a higher and more expensive level (e.g. hybrid DFT) of *ab initio* molecular dynamics. Lower levels of *ab initio* molecular dynamics will predict a delocalized electron.

Exascale Architecture



Coupled-cluster theory on GPUs

Joint work with Eugene DePrince (Argonne)

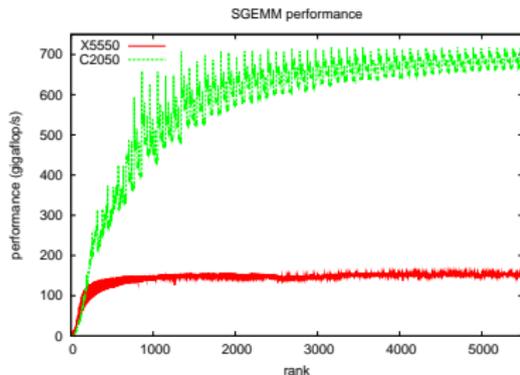
Hardware Details

	CPU		GPU	
	X5550	2 X5550	C1060	C2050
processor speed (MHz)	2660	2660	1300	1150
memory bandwidth (GB/s)	32	64	102	144
memory speed (MHz)	1066	1066	800	1500
ECC available	yes	yes	no	yes
SP peak (GF)	85.1	170.2	933	1030
DP peak (GF)	42.6	83.2	78	515
power usage (W)	95	190	188	238

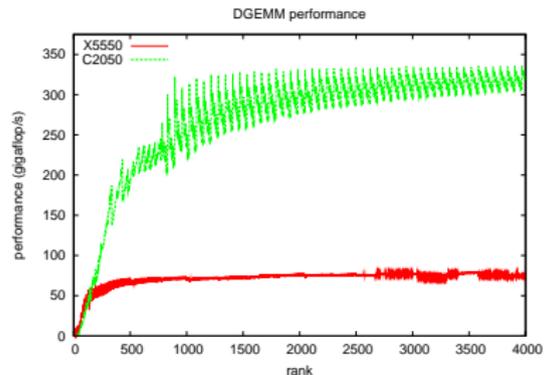
Note that power consumption is apples-to-oranges since CPU does not include DRAM, whereas GPU does.

Relative Performance of GEMM

GPU versus SMP CPU (8 threads):



Maximum:
CPU = 156.2 GF
GPU = 717.6 GF



Maximum:
CPU = 79.2 GF
GPU = 335.6 GF

We expect roughly 4-5 times speedup based upon this evaluation.

Molecule	o	v
C ₈ H ₁₀	21	63
C ₁₀ H ₈	24	72
C ₁₀ H ₁₂	26	78
C ₁₂ H ₁₄	31	93
C ₁₄ H ₁₀	33	99
C ₁₄ H ₁₆	36	108
C ₂₀	40	120
C ₁₆ H ₁₈	41	123
C ₁₈ H ₁₂	42	126
C ₁₈ H ₂₀	46	138

- 6-31G basis set
- C₁ symmetry
- F and V from PSI3.
- GPU code now runs from PSI3.

Without committing to anything, the intent is to release the code under GPL with PSI4.

Can integrate ideas from this code into TCE for NWChem. Kowalski et al. have (T) running on GPUs already.

CCD Performance Results

Iteration time in seconds

	our DP code			X5550		
	C2050	C1060	X5550	Molpro	TCE	GAMESS
C_8H_{10}	0.3	0.8	1.3	2.3	5.1	6.2
$C_{10}H_8$	0.5	1.5	2.5	4.8	10.6	12.7
$C_{10}H_{12}$	0.8	2.5	3.5	7.1	16.2	19.7
$C_{12}H_{14}$	2.0	7.1	10.0	17.6	42.0	57.7
$C_{14}H_{10}$	2.7	10.2	13.9	29.9	59.5	78.5
$C_{14}H_{16}$	4.5	16.7	21.6	41.5	90.2	129.3
C_{20}	8.8	29.9	40.3	103.0	166.3	238.9
$C_{16}H_{18}$	10.5	35.9	50.2	83.3	190.8	279.5
$C_{18}H_{12}$	12.7	42.2	50.3	111.8	218.4	329.4
$C_{18}H_{20}$	20.1	73.0	86.6	157.4	372.1	555.5

Numerical Precision versus Performance

Iteration time in seconds

molecule	C1060		C2050		X5550	
	SP	DP	SP	DP	SP	DP
C ₈ H ₁₀	0.2	0.8	0.2	0.3	0.7	1.3
C ₁₀ H ₈	0.4	1.5	0.2	0.5	1.3	2.5
C ₁₀ H ₁₂	0.7	2.5	0.4	0.8	2.0	3.5
C ₁₂ H ₁₄	1.8	7.1	1.0	2.0	5.6	10.0
C ₁₄ H ₁₀	2.6	10.2	1.5	2.7	8.4	13.9
C ₁₄ H ₁₆	4.1	16.7	2.4	4.5	12.1	21.6
C ₂₀	6.7	29.9	4.1	8.8	22.3	40.3
C ₁₆ H ₁₈	9.0	35.9	5.0	10.5	28.8	50.2
C ₁₈ H ₁₂	10.1	42.2	5.6	12.7	29.4	50.3
C ₁₈ H ₂₀	17.2	73.0	10.1	20.1	47.0	86.6

This the apples-to-apples CPU vs. GPU.

Hybrid CCSD

	Iteration time in seconds						
	Hybrid	CPU	Molpro	NWChem	PSI3	TCE	GAMESS
C ₈ H ₁₀	0.6	1.4	2.4	3.6	7.9	8.4	7.2
C ₁₀ H ₈	0.9	2.6	5.1	8.2	17.9	16.8	15.3
C ₁₀ H ₁₂	1.4	4.1	7.2	11.3	23.6	25.2	23.6
C ₁₂ H ₁₄	3.3	11.1	19.0	29.4	54.2	64.4	65.1
C ₁₄ H ₁₀	4.4	15.5	31.0	49.1	61.4	90.7	92.9
C ₁₄ H ₁₆	6.3	24.1	43.1	65.0	103.4	129.2	163.7
C ₂₀	10.5	43.2	102.0	175.7	162.6	233.9	277.5
C ₁₆ H ₁₈	10.0	38.9	84.1	117.5	192.4	267.9	345.8
C ₁₈ H ₁₂	14.1	57.1	116.2	178.6	216.4	304.5	380.0
C ₁₈ H ₂₀	22.5	95.9	161.4	216.3	306.9	512.0	641.3

Hybrid CCSD

molecule	Basis	o	v	Hybrid	CPU	Molpro	CPU	Molpro
CH ₃ OH	aTZ	7	175	2.5	4.5	2.8	1.80	1.12
benzene	aDZ	15	171	5.1	14.7	17.4	2.88	3.41
C ₂ H ₆ SO ₄	aDZ	23	167	9.0	33.2	31.2	3.69	3.47
C ₁₀ H ₁₂	DZ	26	164	10.7	39.5	56.8	3.69	5.31
C ₁₀ H ₁₂	6-31G	26	78	1.4	4.1	7.2	2.93	5.1

Hybrid code uses multicore OpenMP implementation plus GPU implementation with some DAG scheduling for load-balancing.

Supercomputers and parallel software were critical to the accurate study of:

- 1 di-8-ANEPPS ion channel probe
- 2 dialanine protein model

Proof-of-concept for exascale-oriented node-level parallelism:

- 1 GPU implementation of coupled-cluster theory

Acknowledgments



Argonne: Eugene DePrince, Pavan Balaji, Jim Dinan

PNNL: Karol Kowalski, Niri Govind

UChicago: Benoît Roux, Karl Freed, Glen Hocky

The fast multipole method at petascale

FMM by Ivo Kabadshow and Holger Dachsel (Jülich)
With help from Sreeram Potluri (OSU) and Pavan Balaji (Argonne)

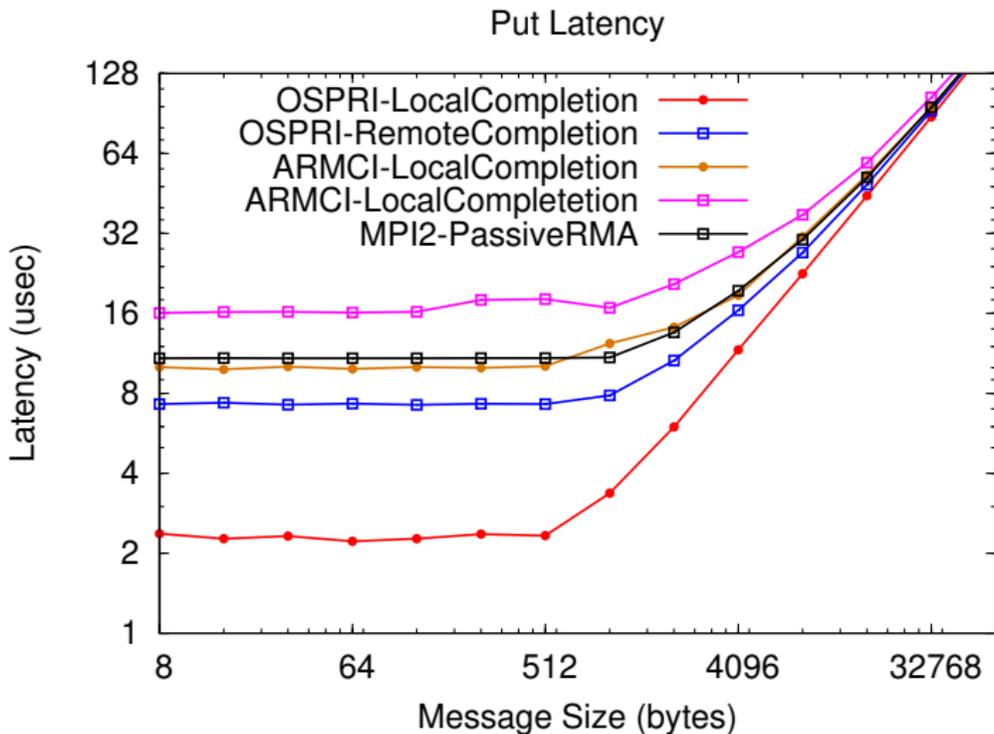
Why does FMM matter?

- N -body with $1/R$ interactions is ubiquitous: gravity, bioMD Coulomb, QC Coulomb, ...
- FMM gets rid of numerical error from summing zeroes: **marriage of speed and accuracy**
- FMM meets exascale hardware requirements: hierarchical, localized communication, high *flop/mop*, ...
- Hierarchical paradigm can be templated onto other local interactions

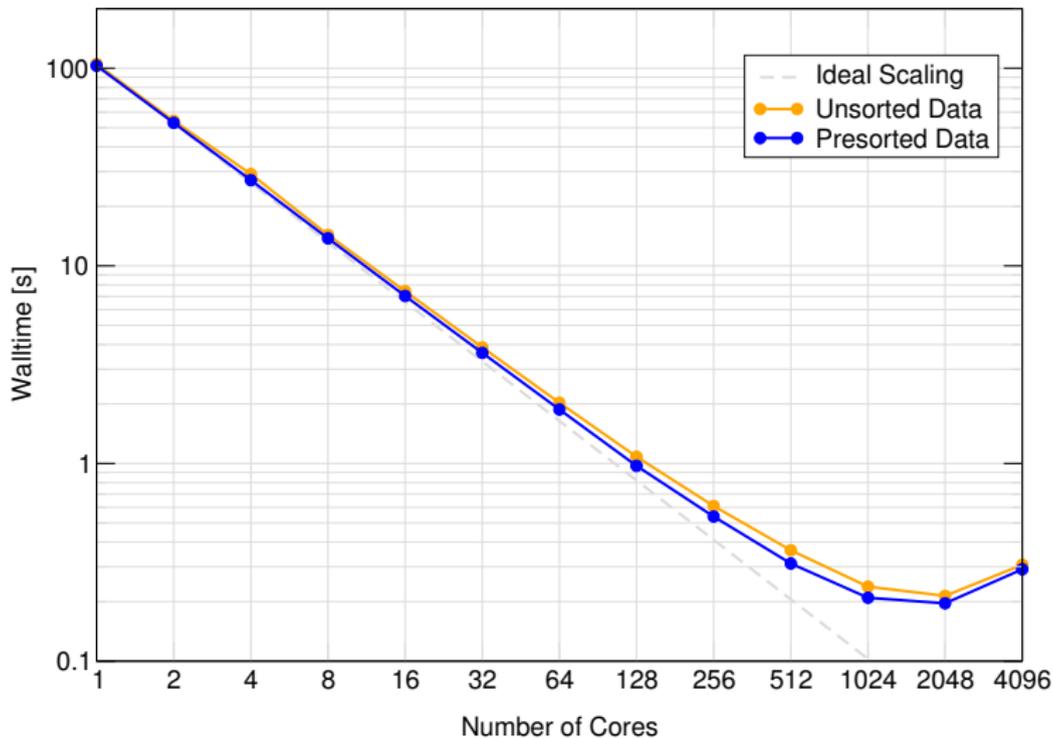
Implementation Details

- One-sided communication (Put) — latency bound
- Minimal synchronization (Notify-Wait)
- BG/P DMA can do Put in hardware
- ARMCI implementation suboptimal on BG/P
- A1 (Argonne 1-sided) aka OSPRI (One-Sided PRI mitives) designed for petascale architectures starting with BG/P
- A1/OSPRI exploits BG/P, XE6, PERCS, BG/Q hardware

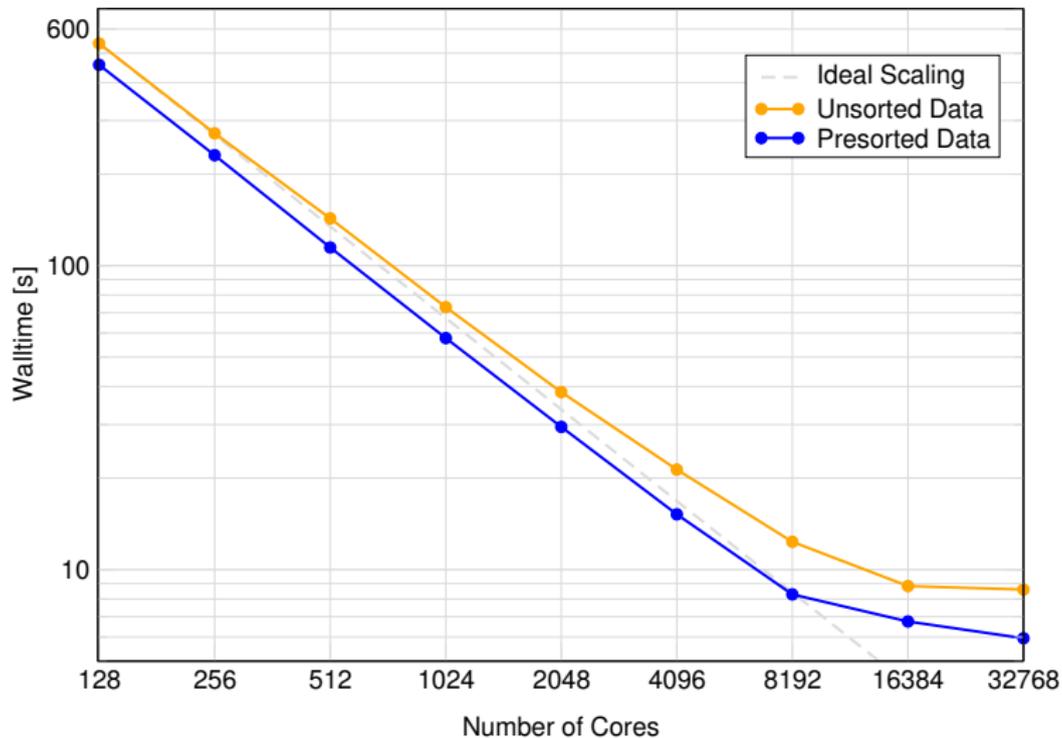
OSPRI versus ARMCI versus MPI-2 on BG/P



$128^3 = 2\text{M}$ particles on BG/P



$1024^3 = 1\text{B}$ particles on BG/P



Without hardware-optimized communication software,
FMM is non-starter on BG/P.

Full JSC machine is 292K cores — 1 PF peak.

First-ever trillion-particle Coulomb solve on whole machine.

Low-overhead error control: *J. Chem. Phys.* **131**, 244102 (2009).