

Quantum Monte Carlo Calculations at the Petascale

David Ceperley

Department of Physics, NCSA

University of Illinois at Urbana-Champaign

- Why quantum Monte Carlo?
- Various QMC methods
- The Endstation Project
- Challenges at the petascale
- Some results in hydrogenic physics
- General considerations -- the sign problem

"The general theory of quantum mechanics is now almost complete. The underlying physical laws necessary for the mathematical theory of a large part of physics and the whole of chemistry are thus completely known, and the difficulty is only that the exact application of these laws leads to equations much too complicated to be soluble."

Dirac, 1929

$$\hat{H} = -\sum_{i=1}^N \frac{\hbar^2}{2m_i} \nabla_i^2 + \sum_{i<j} \frac{e_i e_j}{r_{ij}} \quad \hat{H}\Psi(r_1, r_2, \dots) = E\Psi(r_1, r_2, \dots)$$

Maxwell, Boltzmann and Schrödinger gave us the model (at least for condensed matter physics.) Hopefully, all we must do is numerically solve the mathematical problem and determine the properties. (*first principles or ab initio methods*) Without numerical calculations, the predictive power of quantum mechanics is limited.

The curse of dimensionality: Direct quantum methods are slow!

Suppose we represent the complete **N**-body wavefunction and treat it as strictly a problem in linear algebra—find the exact solution.

If each dimension takes **100** complex numbers

Then **N** particles in 3 dimensions will take **10^{6N}** numbers..

Even with computer time and memory increasing exponentially, the size of system we can treat will only grow linearly in time.

2 particle scattering was done on earliest computers ~1950

3 particle scattering (Fadeev eq., coupled channels) after 1980's.

4 particle scattering (generic 12 dimensional problem) is still very hard.

Expect progress on 5 particle scattering in 2020.

There is no way out of this argument--except to change the problem.

We don't always need the wavefunction! Experimentalists can't measure it!

For example, stochastic methods (simulations) don't scale this way.

Nature sets a very high standard for accuracy!

1eV=11,600K Room temperature = 0.025eV

Modeling of processes relevant to materials, chemistry, biology, ...
needs to be accurate at the level of 0.01eV ~ 0.4mK

Examples for small molecules: error in binding energy

$O_3, H_2O_2, C_2, F_2, Be_2, \dots$

Si_2, P_2, S_2, Cl_2

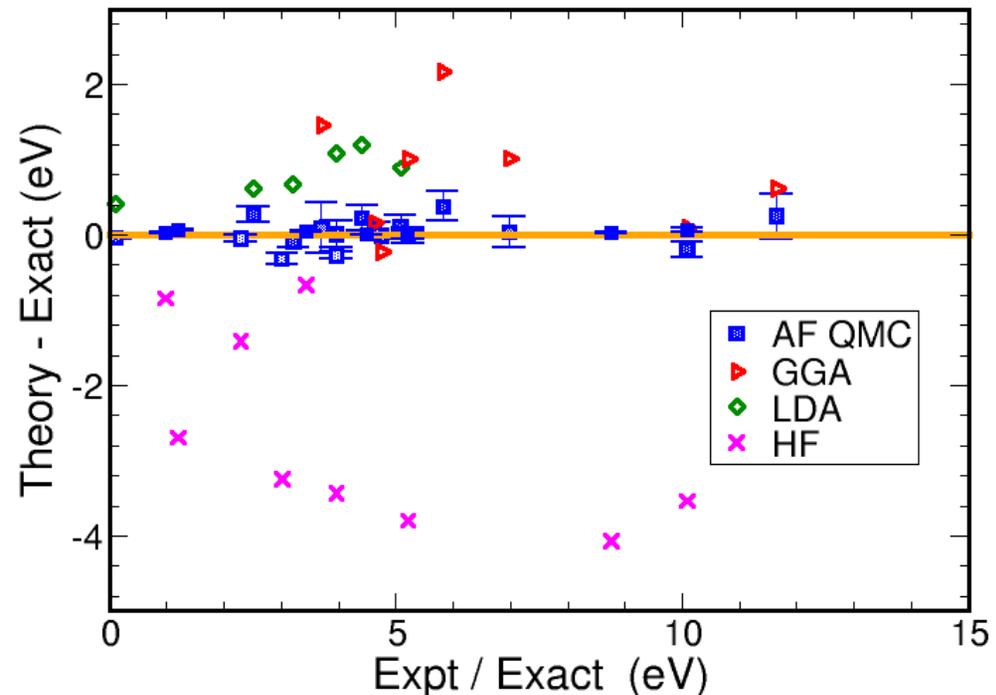
As_2, Br_2, Sb_2

TiO, MnO

Current methods are inadequate

QMC typically accurate to

0.1eV



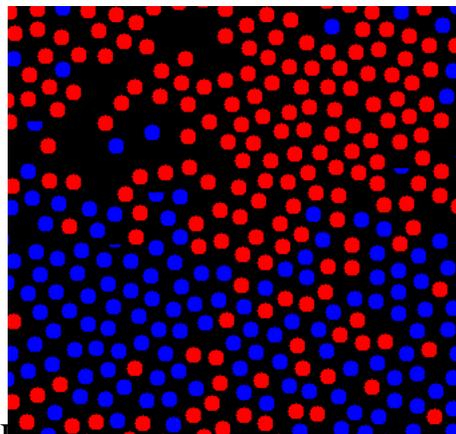
Quantum Monte Carlo

- Premise: need to use simulation techniques to “solve” many-body quantum problems just as you need them classically.
- Both the wavefunction and expectation values are determined by the simulations. Correlation built in from the start.
- QMC gives most accurate method for general quantum many-body systems.
- QMC electronic energy is a standard for approximate DFT calculations. (3rd largest citation in PRL.)
- provide a new understanding of quantum phenomena and a practical tool
- A continuum of stochastic methods:
 - Variational Monte Carlo (VMC)
 - Projector Monte Carlo methods for $T=0$:
 - Diffusion Monte Carlo (DMC)
 - Reptation MC (RQMC)
 - Auxiliary field QMC (AFQMC)
 - Path Integral Monte Carlo for $T>0$ (PIMC)
 - Coupled electron-ion Monte Carlo $T>0$ (CEIMC)

Goal is NOT large N, but higher accuracy and new capabilities

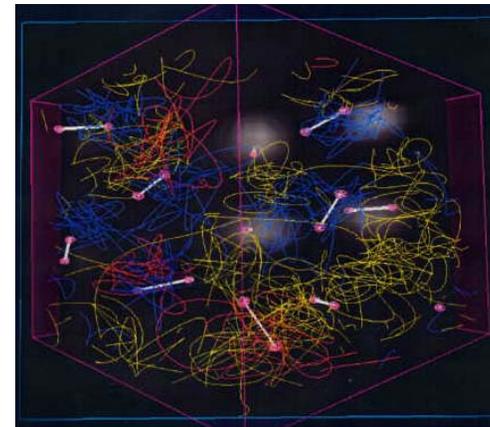
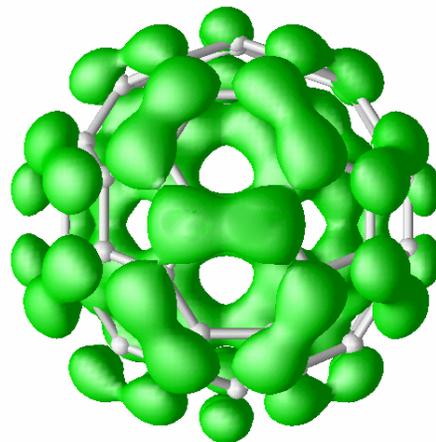
Breakthrough Quantum Monte Carlo simulations:

- Hard-core bosons on a CDC 6600 (1974)
- Electron gas on CRAY-1 (1980)
- Superfluid helium (1984)
- Ground state of solid hydrogen at high pressures, CRAY XMP and CYBER 205 (1987)
- Electronic and structure properties of carbon/silicon clusters on HP 9000/715 cluster and Cray Y-MP (1995)
- Coupled Electron-Ion Monte Carlo simulations of dense hydrogen on Linux Clusters (2000s)



ENDSTATION Project

FallC



7

Variational Monte Carlo (VMC)

(McMillan 1965)

- Put correlation directly into the wavefunction.
- Integrals are hard to do: need MC.
- Take sequence of increasingly better wavefunctions. Stochastic optimization is important!
- **Can we make arbitrarily accurate functions?** Method of residuals says how to do this.
- Recent progress with "backflow"
- No sign problem, and with classical complexity.

- Posit a wavefunction $\phi(\mathbf{R}, \mathbf{a})$
- sample $|\phi(\mathbf{R}, \mathbf{a})|^2$ with random walk.
- minimize energy or variance of $\phi(\mathbf{R}, \mathbf{a})$ with respect to \mathbf{a}

$R \equiv (\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N) = \text{"walker"}$

$$\Psi_2(R) = \text{Det}\{\phi_i(\mathbf{r}_j)\} e^{-\sum_{i<j} u_{ij}(r_{ij})}$$

$$\Psi_{n+1}(R) \approx \Psi_n(R) e^{-\underbrace{\langle \phi_n^{-1} H \phi_n \rangle}_{\text{smoothing}}}$$

smoothing

One-body orbitals

- Fixed-node approximation in DMC/RMC uses the “nodes” of

$$\Psi_{AS}(\{\mathbf{R}\}) = D^\dagger D^\downarrow$$

- Elementary operation: random access to wf tables
- What method is used to obtain $\{\phi_k(\mathbf{r})\}$ matters for accuracy!

–Hartree-Fock? DFT? Hybrid ? or beyond HF/DFT

–How to represent $\{\phi_k(\mathbf{r})\}$ matters for efficiency!

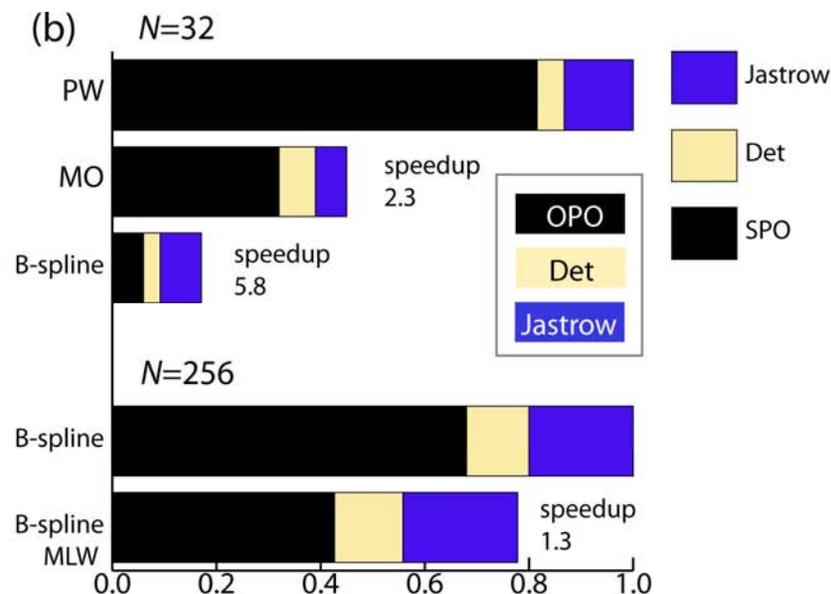
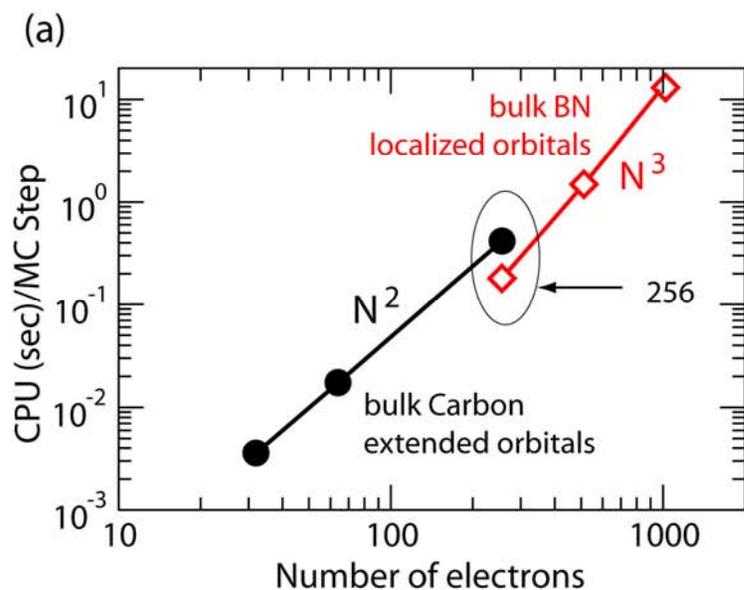
- Plane-wave basis set
- Molecular orbitals in a localized atomic basis set: Gaussian- and Slater-type orbitals and numerical orbitals
- Real-space grid

–we have to calculate $\phi_k(\mathbf{r})$, $\nabla\phi_k(\mathbf{r})$, and $\nabla^2\phi_k(\mathbf{r})$ at \mathbf{r}

- Can use any combination of mixed basis sets that are optimized for the performance and computing resources.

-This is the time consuming step today!

Cost of Evaluations of Wave Function



- Molecular Orbitals $\psi_i(\mathbf{r}) = \sum_l \sum_\alpha c_i^{l,\alpha} \phi_{i,\alpha}(\mathbf{r} - \mathbf{R}_l) \sim N^3$
- Plane-wave $\psi_i(\mathbf{r}) = \sum_{\mathbf{G}} c_i^{\mathbf{G}} \exp^{i\mathbf{G}\cdot\mathbf{r}} \sim N^3$
- B-spline $\psi_i(\mathbf{r}) = \sum_l c_i^l f_l(\mathbf{r}) \sim N^2$

Projector Monte Carlo

e.g. Diffusion Monte Carlo (DMC)

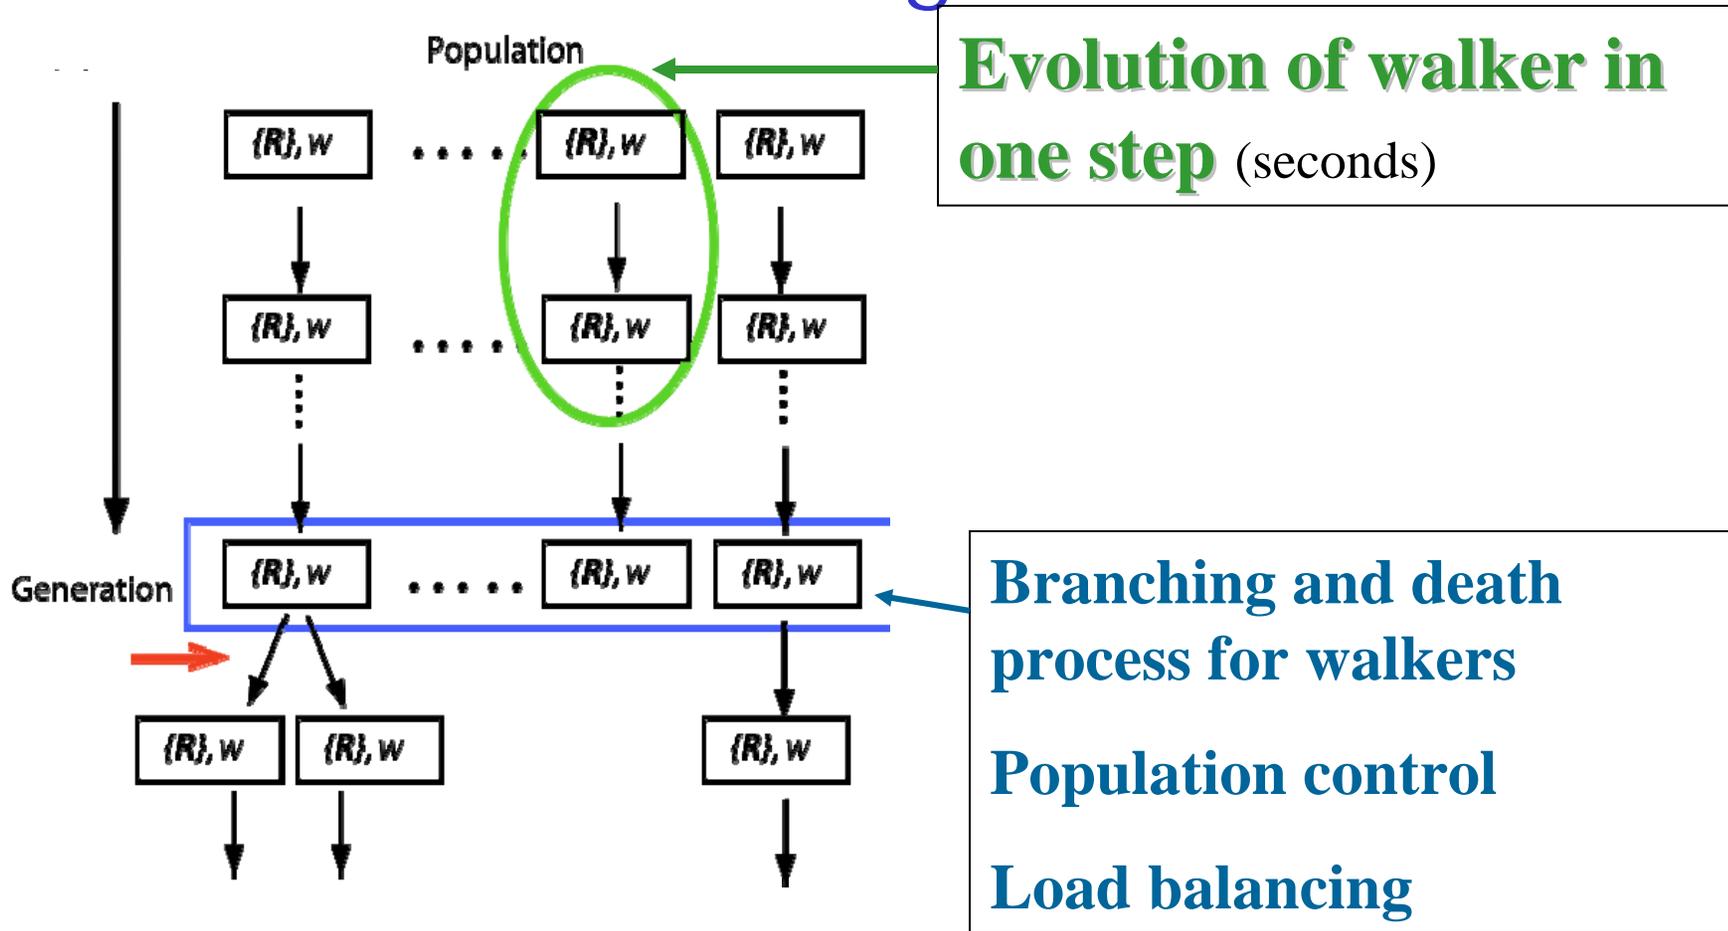
- Automatic way to get better wavefunctions.
- Project single state using the Hamiltonian

$$\phi(t) = e^{-(H-E)t} \phi(0)$$

- This is a diffusion + branching operator.
- Very scalable: each walker gets a processor.
- **But is this a probability?**
- **Yes!** for bosons since ground state can be made real and non-negative. **But** all excited states must have sign changes.
- In **exact** methods one carries along the sign as a weight and samples the modulus. This leads to the famous sign problem

$$\phi(t) = e^{-(H-E)t} \text{sign}(\phi(R,0)) |\phi(R,0)|$$

DMC algorithm



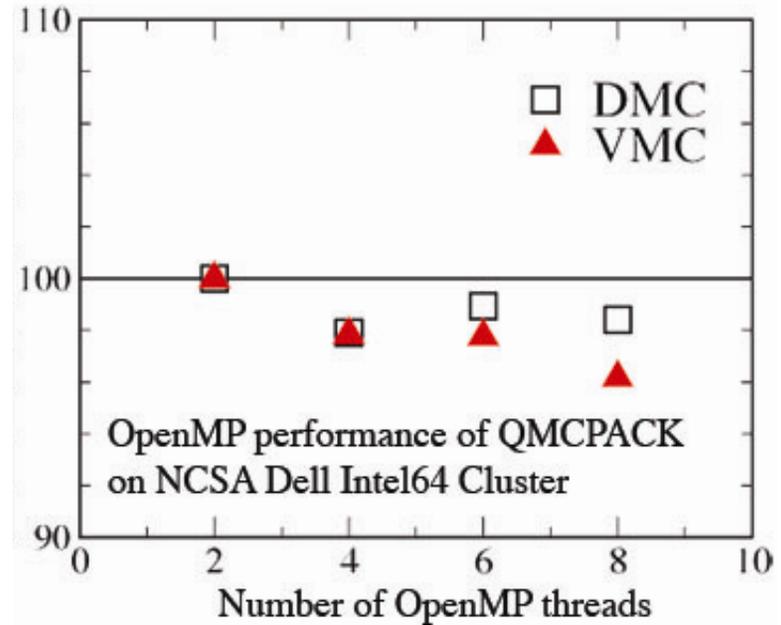
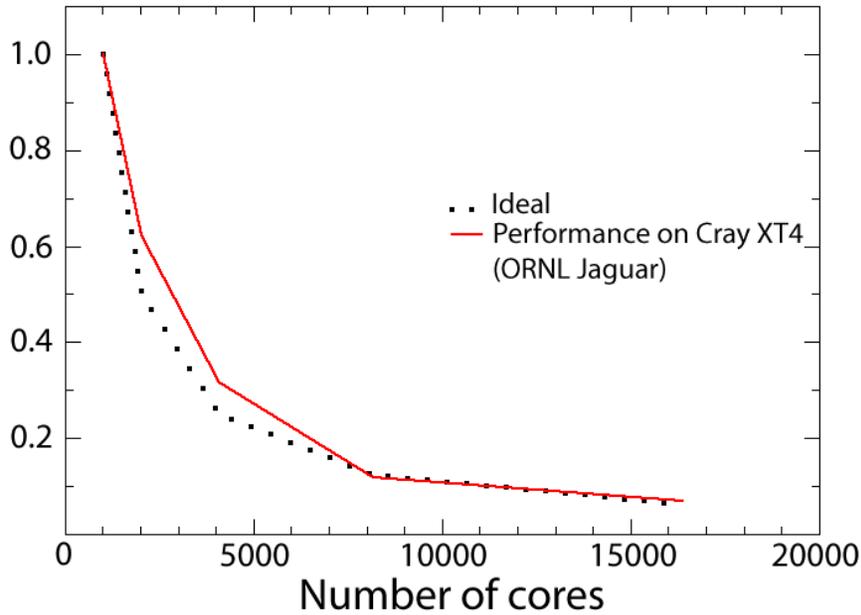
QMCPACK

- Open-source library and application package to perform Quantum Monte Carlo (QMC) Simulations
- Implements various QMC algorithms: VMC, DMC, RQMC
 - Generic representations of the physical entities and models
 - Object-oriented implementation of QMC algorithms (C++)
 - Generic programming of computational Kernels
- Designed for **large-scale QMC simulations** of molecules, solids and nanostructures: OpenMP/MPI Hybrid parallelization, effective for multi-core systems
- Standard open-source libraries and utilities for development, compilation and executions
- Adopts XML/HDF5 for I/O
- Developed at MCC and NCSA <http://www.mcc.uiuc.edu/qmc/>

Principal author: Jeongnim Kim UIUC

DMC scales well

Wall time to solution (normalized)

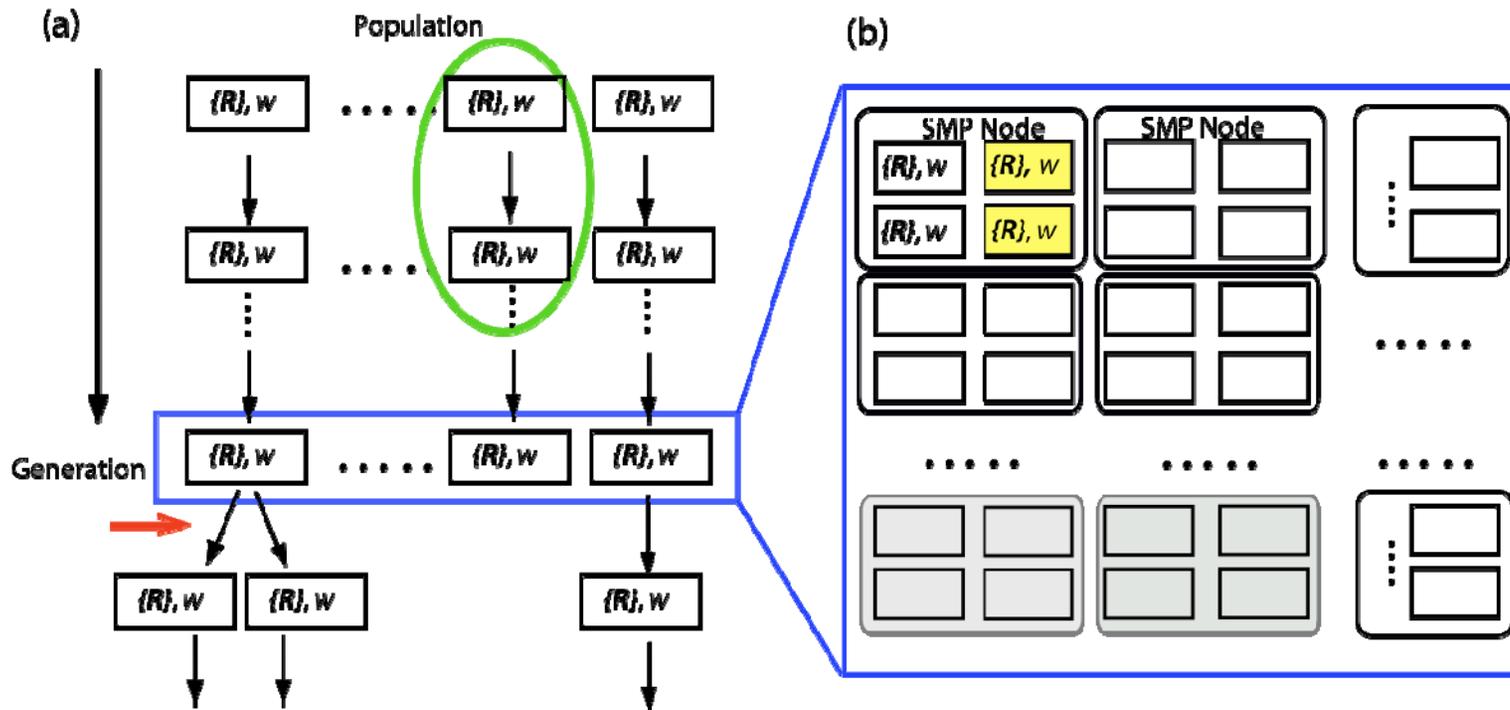


Current work: extend production codes to $> 30,000$ nodes

What happens when we scale up? Say by a factor of 100

- Can one distribute more walkers across nodes?
 - Will reduce statistical error by factor of 10
 - But wall clock time will be the same
 - Systematic errors will be the same
 - Load balancing/fault tolerance problems
 - For most problems, we will not achieve the science goals
- Need to find more parallelism to keep errors the same but reduce wall clock time and reduce systematic errors
- Vulnerable to software and hardware errors?
- Are existing pseudo-random number generators adequate at the petascale? Tested SPRNG for 28K random number streams with 10^{10} numbers/stream.

DMC algorithm and OpenMP/MPI parallelization

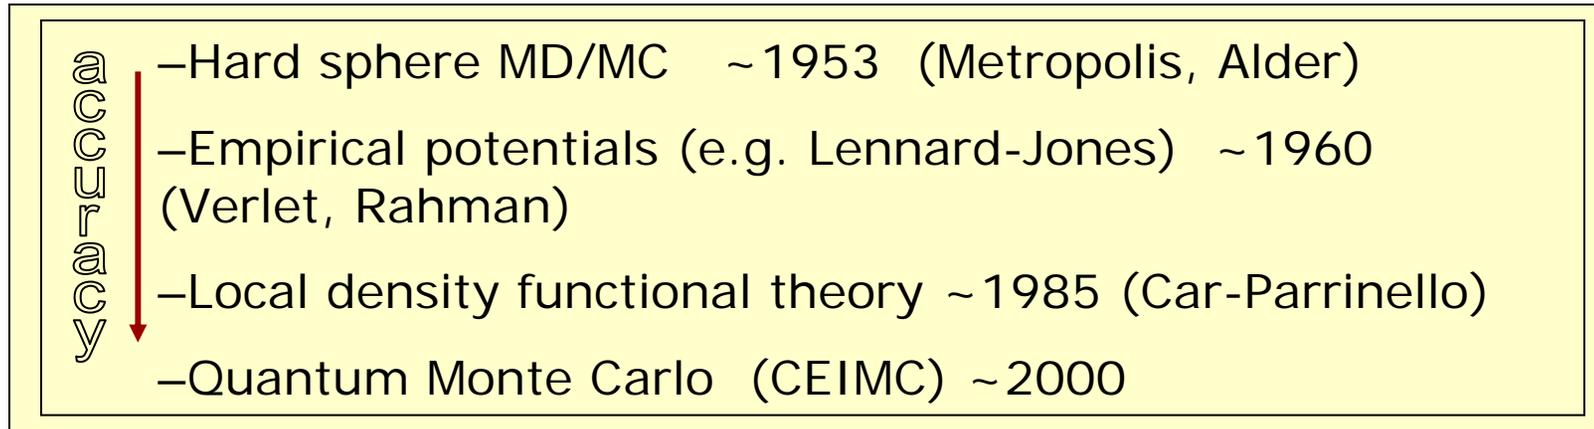


- Memory limitations of the current MPI parallelization
- Will total memory per SMP node will grow with more cores: 8-64 GB will be sufficient for many problems.
- One or more walker $\{R\}$ per core (OpenMP thread)
- Read-only large-scale data (e.g., wavefunctions) are shared but others objects are allocated per thread.

Load Balancing for QMC

- Multi-level scheme of managers, sub-managers, workers.
- Walker point of view:
 - After a walker completes an elementary step, it sends averages, and branching information up the tree
 - Looks for new configurations to send or receive.
- Manager
 - Decides on load balancing and sends transfer information down the tree
 - Provides overall stability of population, checkpointing and I/O of averages
 - When processors die, those branches are pruned
- Eliminate need for synchronization, blocking calls,...
- Can run in an unstable environment

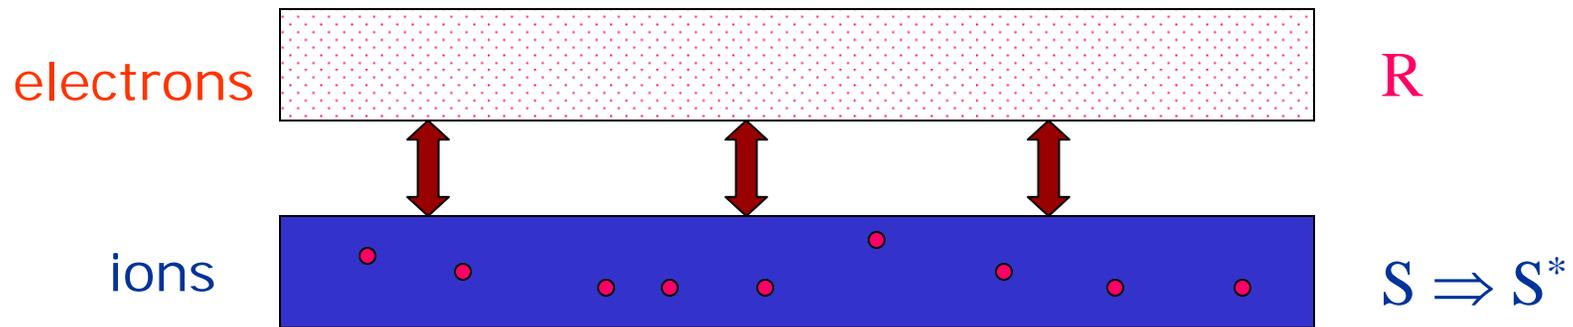
MD and MC Simulations



- Initial simulations used semi-empirical potentials.
- Much progress with “ab initio” molecular dynamics simulations where the effects of electrons are solved for each step.
- However, the potential surface as determined by density functional theory is not always accurate enough
- **QMC+MD =CEIMC is a candidate for petascale “killer app”**

Coupled Electron-Ionic Monte Carlo: CEIMC

1. Do Path Integrals for the ions at $T > 0$.
2. Let electrons be at zero temperature, a reasonable approximation for room temperature simulations.
3. Use Metropolis MC to accept/reject moves based on QMC computation of electronic energy



The “noise” coming from electronic energy can be treated without approximation: the penalty method.

A feature of Monte Carlo?

averages are almost free.

Suppose we have an extra parameter “q” to sum over.

$$E(s) = \frac{1}{M} \sum_{i=1}^M E(s; q_i)$$

- In a deterministic calculation, this will multiply CPU time by M.
- This extra parameter will not slow down the calculation by Monte Carlo: it is just one more variable to average over.
- We start up M calculations on M separate processors for different values of q: they all serve to reduce the error bar.
- The only slow down occurs comes from “start up” costs: e. g. Metropolis warm-up or initializations.

Twist averaged boundary conditions

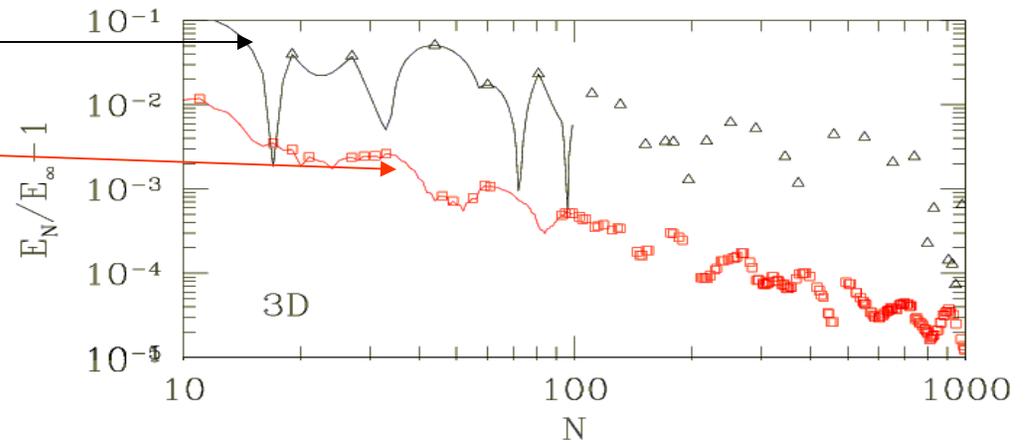
- In periodic boundary conditions, the wavefunction is periodic \Rightarrow Large finite size effects for metals because of fermi surface.
- In twist averaged BC, we use an arbitrary phase θ as $r \rightarrow r+L$
- Integrate over all phases, i.e. Brillouin zone integration.
- Momentum distribution changes from a lattice of k-vectors to a fermi sea.
- Eliminates single-particle finite-size effects.

$$\Psi(x + L) = e^{i\theta} \Psi(x)$$

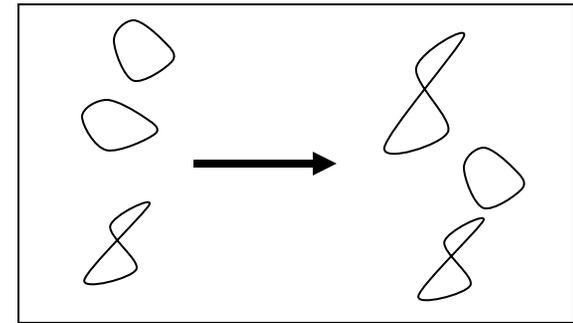
$$\bar{A} = \frac{1}{(2\pi)^3} \int_{-\pi}^{\pi} d^3\theta \langle \Psi_{\theta} A \Psi_{\theta} \rangle$$

Error with PBC
Error with TABC

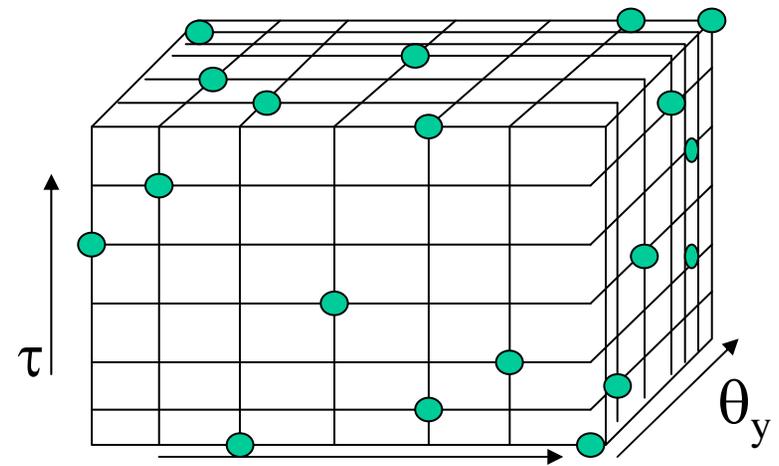
Error is zero in the grand canonical ensemble at the mean field level.



- Make a move of the protonic paths
- Partition the 4D lattice of boundary conditions ($\theta_x \theta_y \theta_z$) and imaginary time (τ) in such a way that each variable is uniformly sampled (stratified)
- Send them all out to M separate processes
- Do DMC to get energy differences and variances
- Combine to get global difference and variance.



$$R_\tau \rightarrow R'_\tau$$



$$\Delta E_{BO} = \frac{1}{M} \sum E_{\theta, \tau}$$

$$\sigma^2 = \frac{1}{M^2} \sum \sigma_{\theta, \tau}^2$$

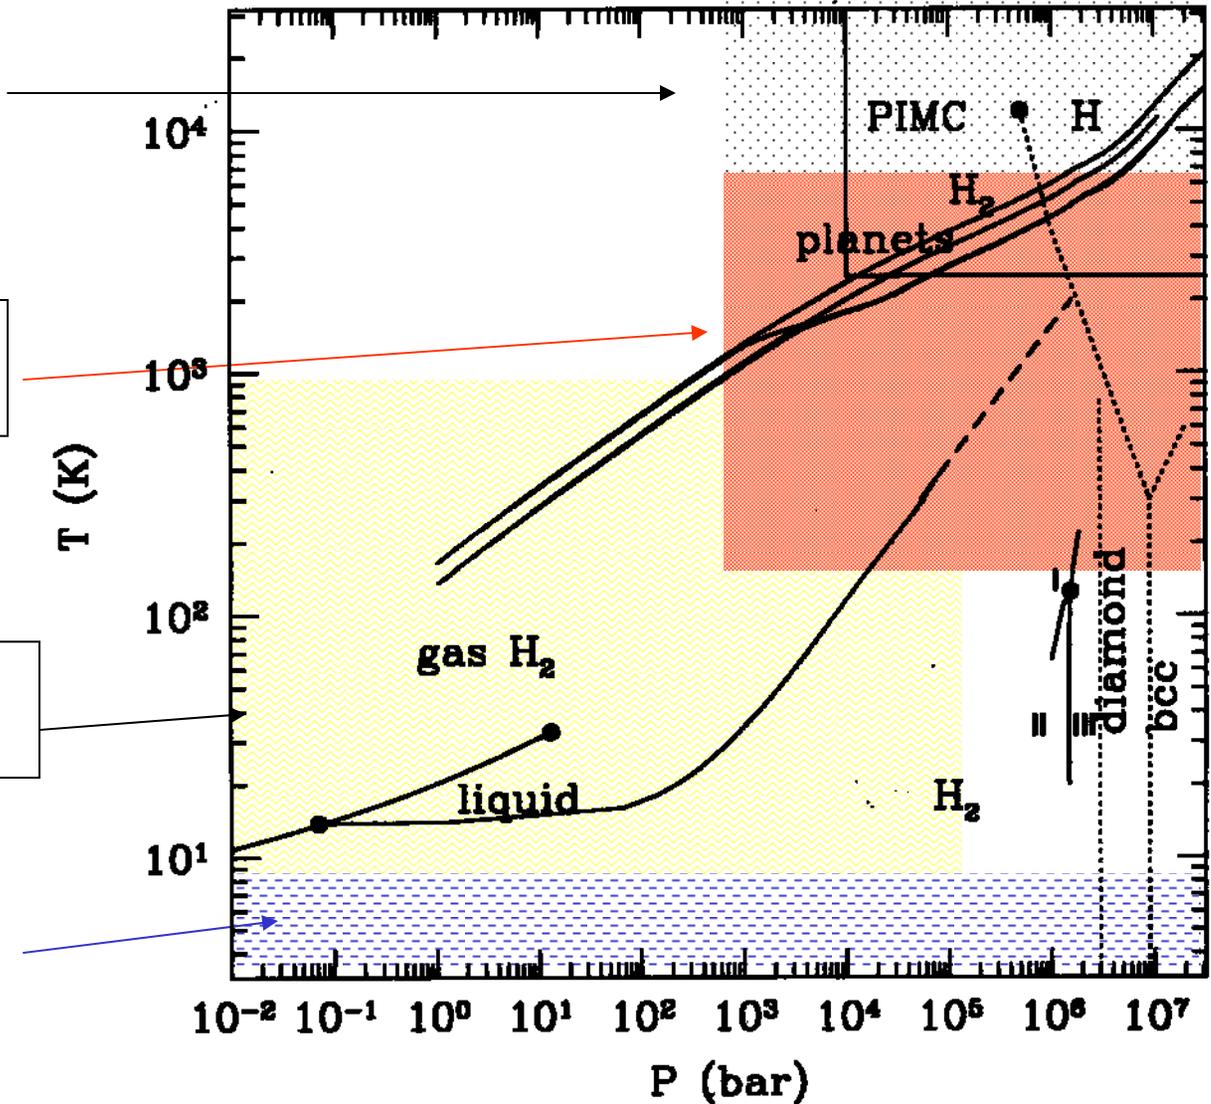
QMC methods for Hydrogen

Path Integral MC for
 $T > E_F/10$

Coupled-electron Ion
 MC

Path Integral MC with
 an effective potential

Diffusion MC $T=0$

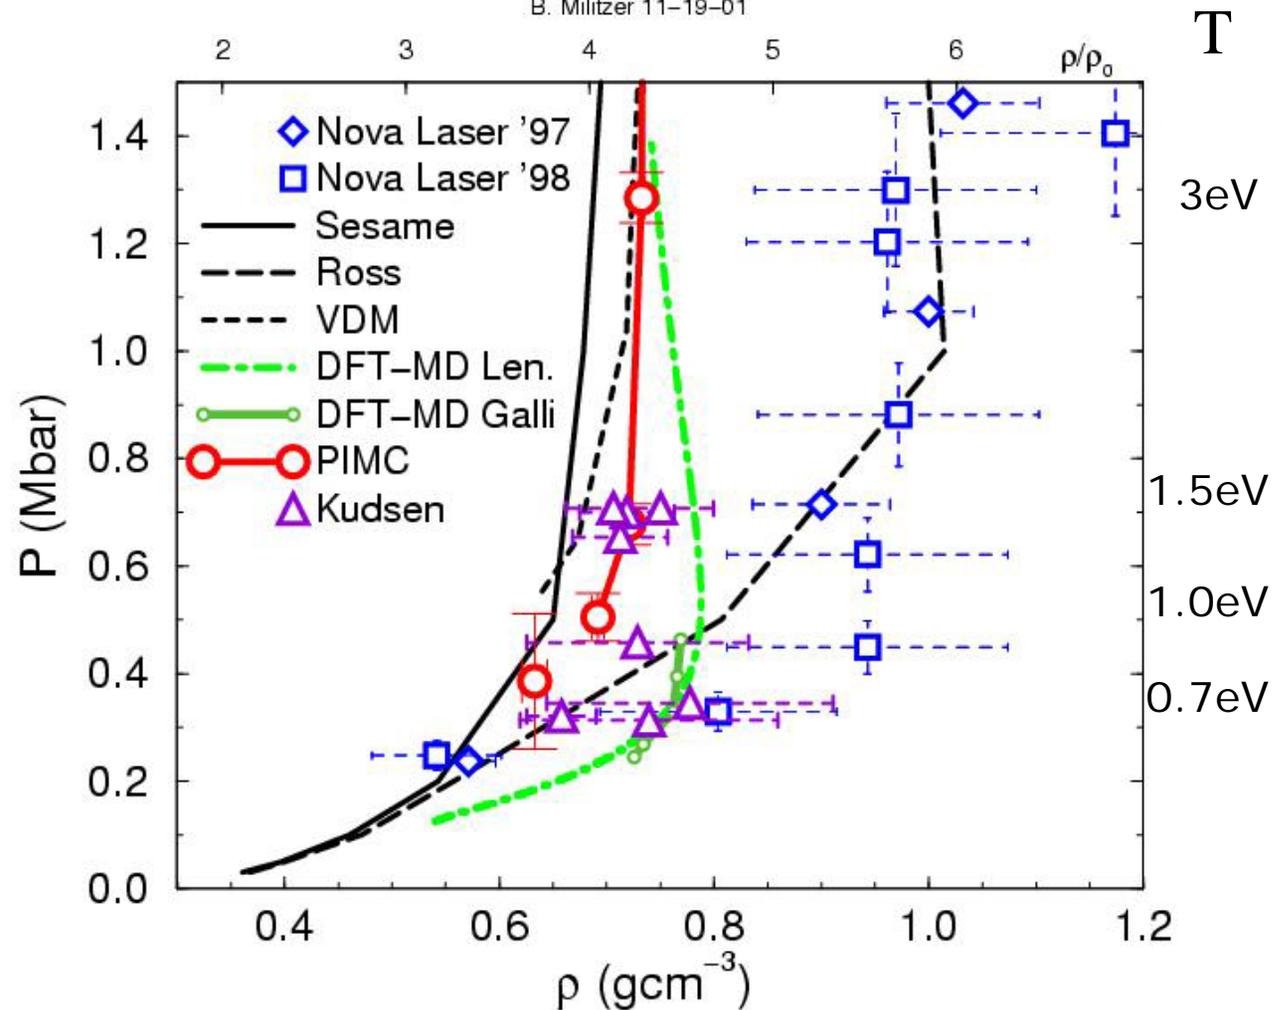


Experiment vs PIMC/DFT simulations

Deuterium Hugoniot

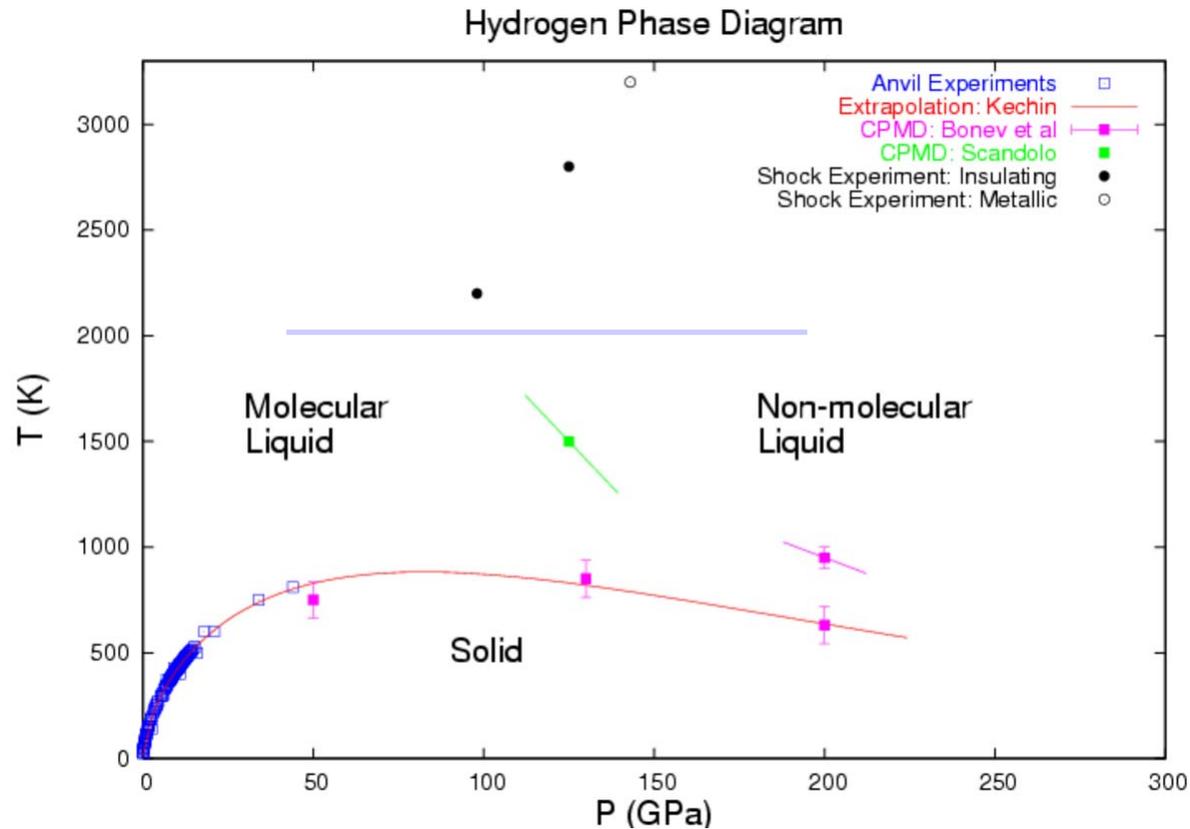
B. Militzer 11-19-01

- Older laser (NOVA) shocks are incompatible with microscopic theory.
- Chemical models are not predictive in this regime.
- Z-pinch experiments of *Knudsen et al., PRL 87, 225501 (2001)*
- QMC predicted results that were later vindicated!



Plasma Phase Transition

- Study nature of transition from molecular to non-molecular fluid using CEIMC
- Simulations at $T=2000K$ with $P=50-200GPa$



Atomic-Molecular Transition

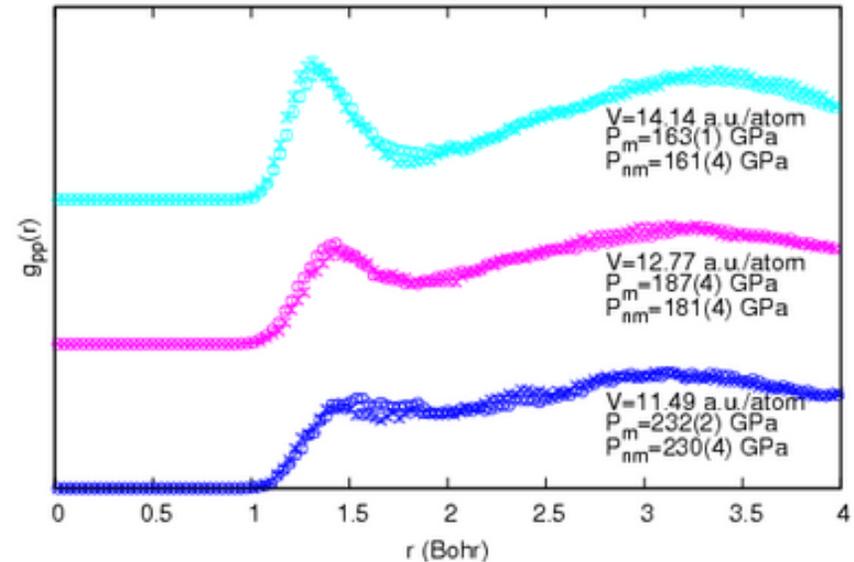
- VMC: Hysteresis; probably 1st order transition.
- RQMC: No hysteresis; continuous transition.
- VMC trial function has difficulty with the mixed H₂-H state.

$$N_p = N_e = 54$$

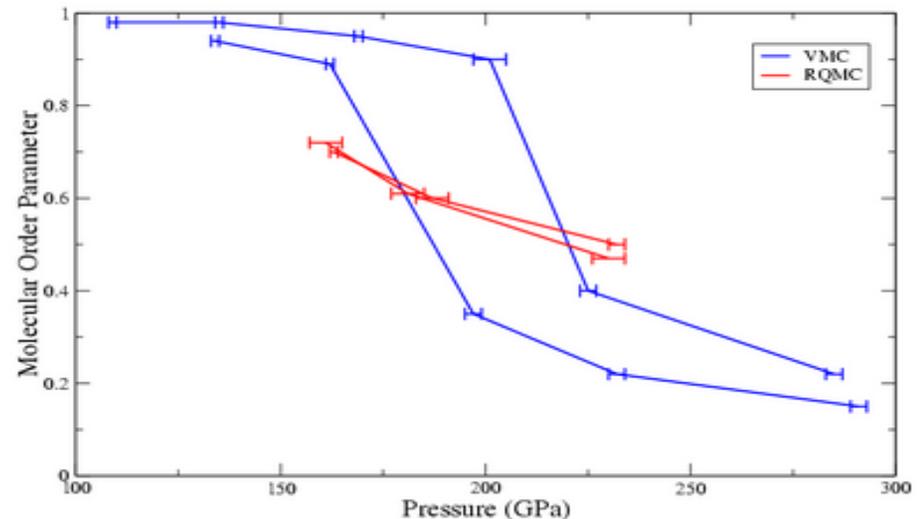
H₂ order parameter

$$g(r) = \lambda g_{mol}(r) + (1 - \lambda) g_{nonmol}(r)$$

Collection of $g_{pp}(r)$ s for T=2000K using RQMC



Molecular to Non-molecular Fluid Phase Transition



General statement of the “fermion problem”

- Given a system with **N** fermions and a known Hamiltonian and a property **O** (usually the energy):
 - How much time **T** will it take to estimate **O** to an accuracy ϵ ?
 - How does **T** scale with **N** and ϵ ?
- If you can map the quantum system onto an equivalent problem in classical statistical mechanics then:

$$T \propto N^\alpha \epsilon^{-2} \quad \text{With } 0 < \alpha < 3$$

This would be a “solved” quantum problem!

- All approximations must be controlled!
- Algebraic scaling in N!

e.g. properties of Boltzmann or Bose systems in equilibrium.

Which problems can be done with the direct fermion methods at the petascale? e.g. the 3DEG (1980)

SUMMARY

- No existing methods are perfect but QMC today is competitive with other methods and usually much more accurate.
- Progress in *ab initio* simulations in last 40 years, coming from both
 - Computer power
 - Algorithmic power
- We are now in position to do much more accurate simulations
- Our petascale goal: water and transition metal oxides (e.g. MnO), ...
- Ab initio computation of electronic system is a great problem to solve. Intellectually **and** technologically very important. More work needed in algorithms to get higher accuracy, treat larger systems, heavier elements allowing:
 - benchmarking to validate cheaper approaches
 - replace more approximate approaches.

The ENDSTATION TEAM

- Jeongnim Kim, David Ceperley, *University of Illinois at Urbana-Champaign*
- Henry Krakauer, Shiwei Zhang, *College of William and Mary*
- Paul Kent, *Oak Ridge National Laboratory*
- Richard Hennig, Cyrus Umrigar, *Cornell University*
- Lubos Mitas, *North Carolina State University*
- Ashok Srinivasan, *Florida State University*

+ many students, postdocs and other collaborators

The Endstation project is supported by: DOE-DE-FG05-08OR23336

Computer time through the INCITE program