The QMC Petascale Project

Richard G. Hennig

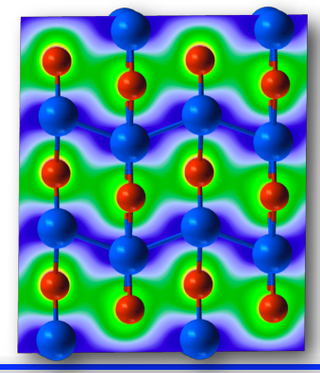

What will a petascale computer look like?
What are the limitations of current QMC algorithms for petascale computers?
How can Quantum Monte Carlo algorithms take advantage of petascale?

- 1 Petaflop = 10^{15} floating point operations per second
- Next generation of massively parallel computers: 10^5 – 10^6 cores
- Limited memory and bandwidth
- Requires novel algorithm to obtain optimal performance
- QMC algorithms for petascale
- Comparison and benchmarking of VMC/DMC codes

Supported by DOE & NSF
Computational resources provided by NCCS, CCNI, NCSA, NERSC and OSC

Cornell University
College of Engineering
Materials Science and Engineering

rhennig@cornell.edu

Quantum Monte Carlo Petascale Endstation

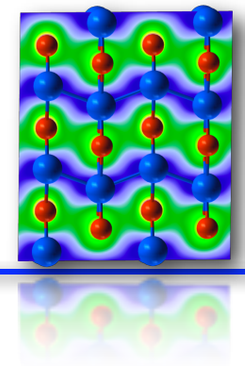

Motivation:

- Breakthrough QMC simulations enabled by the state-of-art, high-performance computers of every generation
  - Hard-core bosons – CDC 6600 (Kalos '74)
  - Ground state of hydrogen at high pressures – CRAY XMP & CYBER 205 (Ceperley '87)
  - Carbon/silicon clusters – HP 9000/715 cluster & Cray Y-MP (Williamson '01)
  - Continuous DMC-MD – LLNL teraflop cluster (Grossman '05)
  - Coupled Electron-Ion Monte Carlo simulations of hydrogen – Cray XT3/4 at ORNL (Ceperley)
- Novel QMC algorithms
  - Better wave functions (backflow, geminals, pfaffians)
  - Efficient optimization methods for all parameters (Jastrow, CSF, orbital, basis)
  - Efficient computation of forces and energy differences
  - Better estimators of observable (zero-variance & zero-bias principle)
  - Improved scaling (localized orbitals, B-splines)
  - Efficient finite size corrections

Team:

Ceperley & Kim (UIUC), Mitas (NC), Hennig & Umrigar (Cornell), Zhang & Krakauer (W&M), Kent & Schulthess (ORNL), Srinivasan (Florida State)

Cornell University
College of Engineering
Materials Science and Engineering

rhennig@cornell.edu

Processors for Petascale

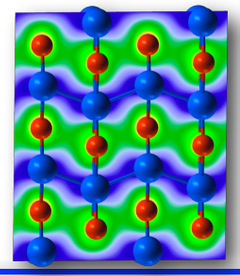

Moore's law

- Number of transistors on a chip doubles every 2 years

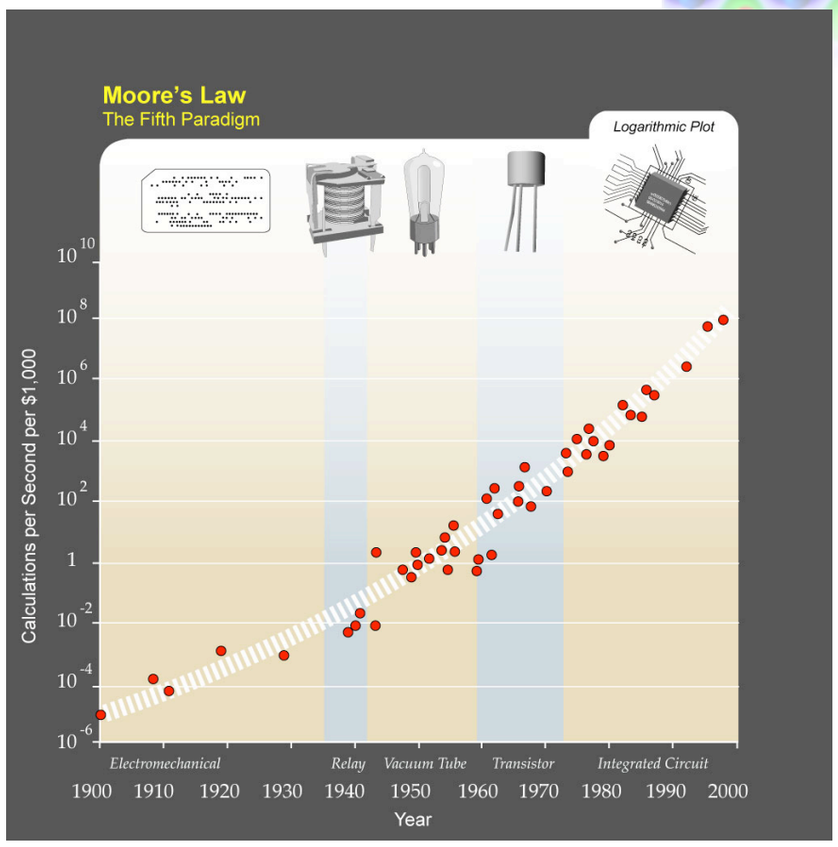

Homogeneous Multicore Architectures (Intel)

- Intel Teraflop Research Chip
  - 80 cores, 3.16 GHz, 62 Watt power consumption
- Compare to ASCI Red
  - 1 TFlop in 1996, 10,000 processors
  - 500 kW power + 500 kW cooling

Heterogeneous Multicore Architectures (AMD)

- Combination of cores for specific tasks
- Graphics (GPU's), network (FGPA's) & compute cores

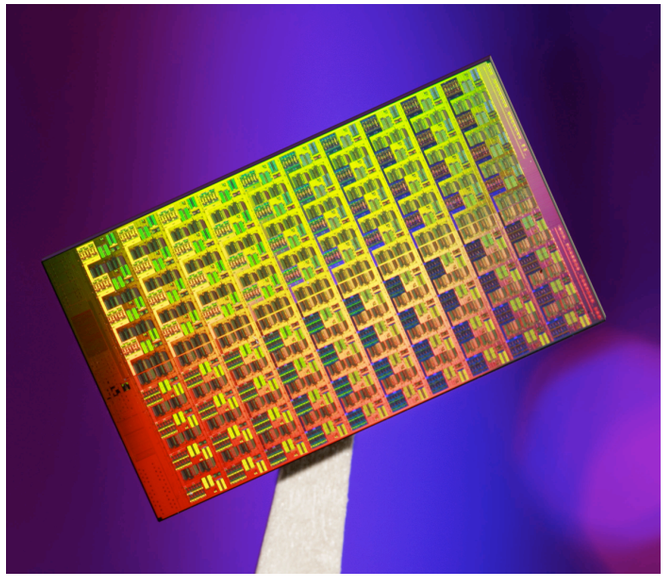

Low Power Consumption

- IBM Power 4

Tera → Petascale Supercomputers

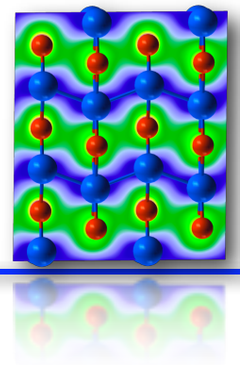

Massively Parallel

- Over 100,000 cores ⇒ Small memory footprint
- Multilevel parallelism
- Multi (32) core processors ⇒ Limited memory bandwith

Examples of Existing Systems

Machine	Vendor	Performance	Top500	Processors	Memory
DOE/ OakRidge	Cray XT4	205 TFlops	5	7,832 quad core 2.1 GHz AMD Opteron	2 GB/core
NERSC/ LBNL	Cray XT4	85 TFlops	15	9,660 dual core 2.6 GHz AMD Opteron	2 GB/core
CCNI/RPI	IBM Bluegene/L	73 TFlops	22	32,768 700 MHz IBM Power 440	256/512 MB/proc
NCSA	Dell	68 TFlops	23	2,400 quad core 2.3GHz Intel Xeon	1 GB/core

Scaling of Current QMC Algorithms

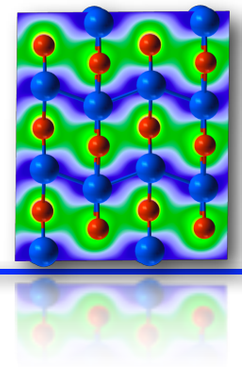

Computational effort in DMC dominated by wave function evaluation

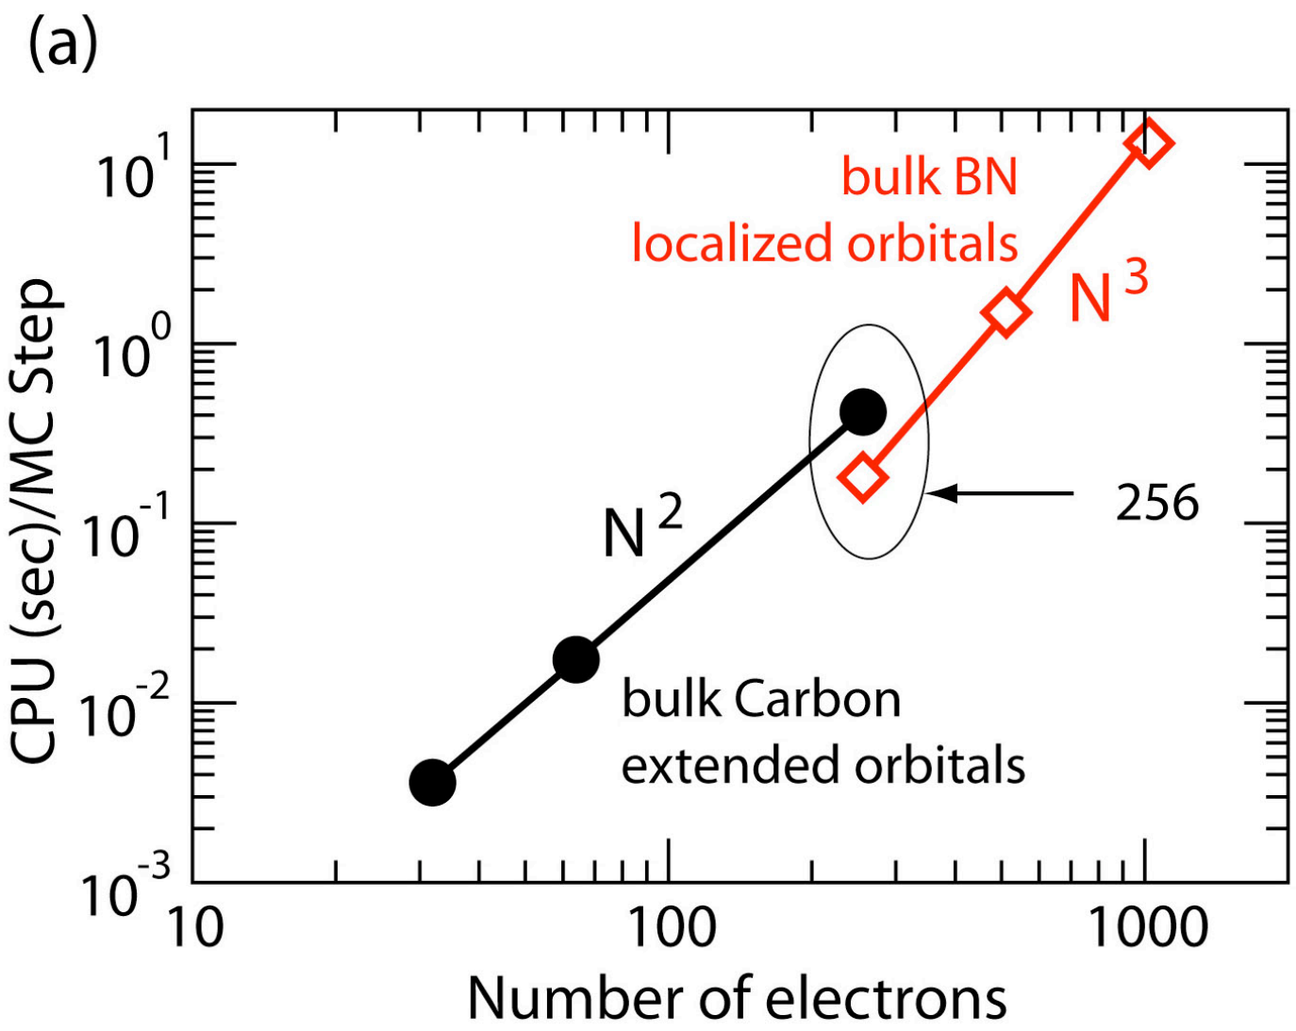

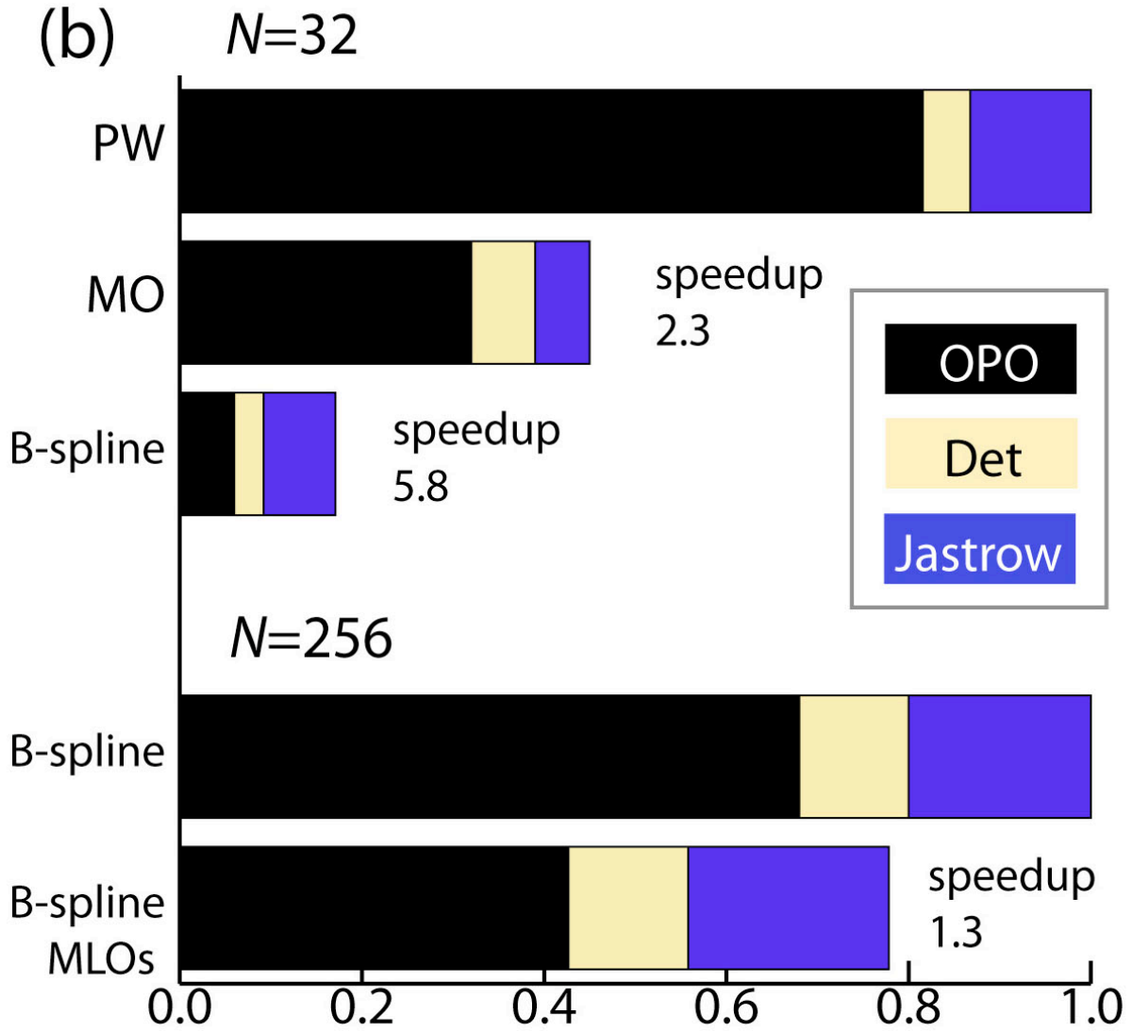

O(N^3) dominates above ≈1,000 electrons

B-spline cost dominates
Jastrow and determinant part increasingly important for larger systems

Cornell University
College of Engineering
Materials Science and Engineering

rhennig@cornell.edu

Parallelism in Quantum Monte Carlo

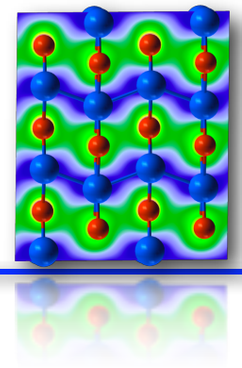

Walker distribution

- Stochastic propagation of walker ensemble ⇒ Distribution of walkers across nodes

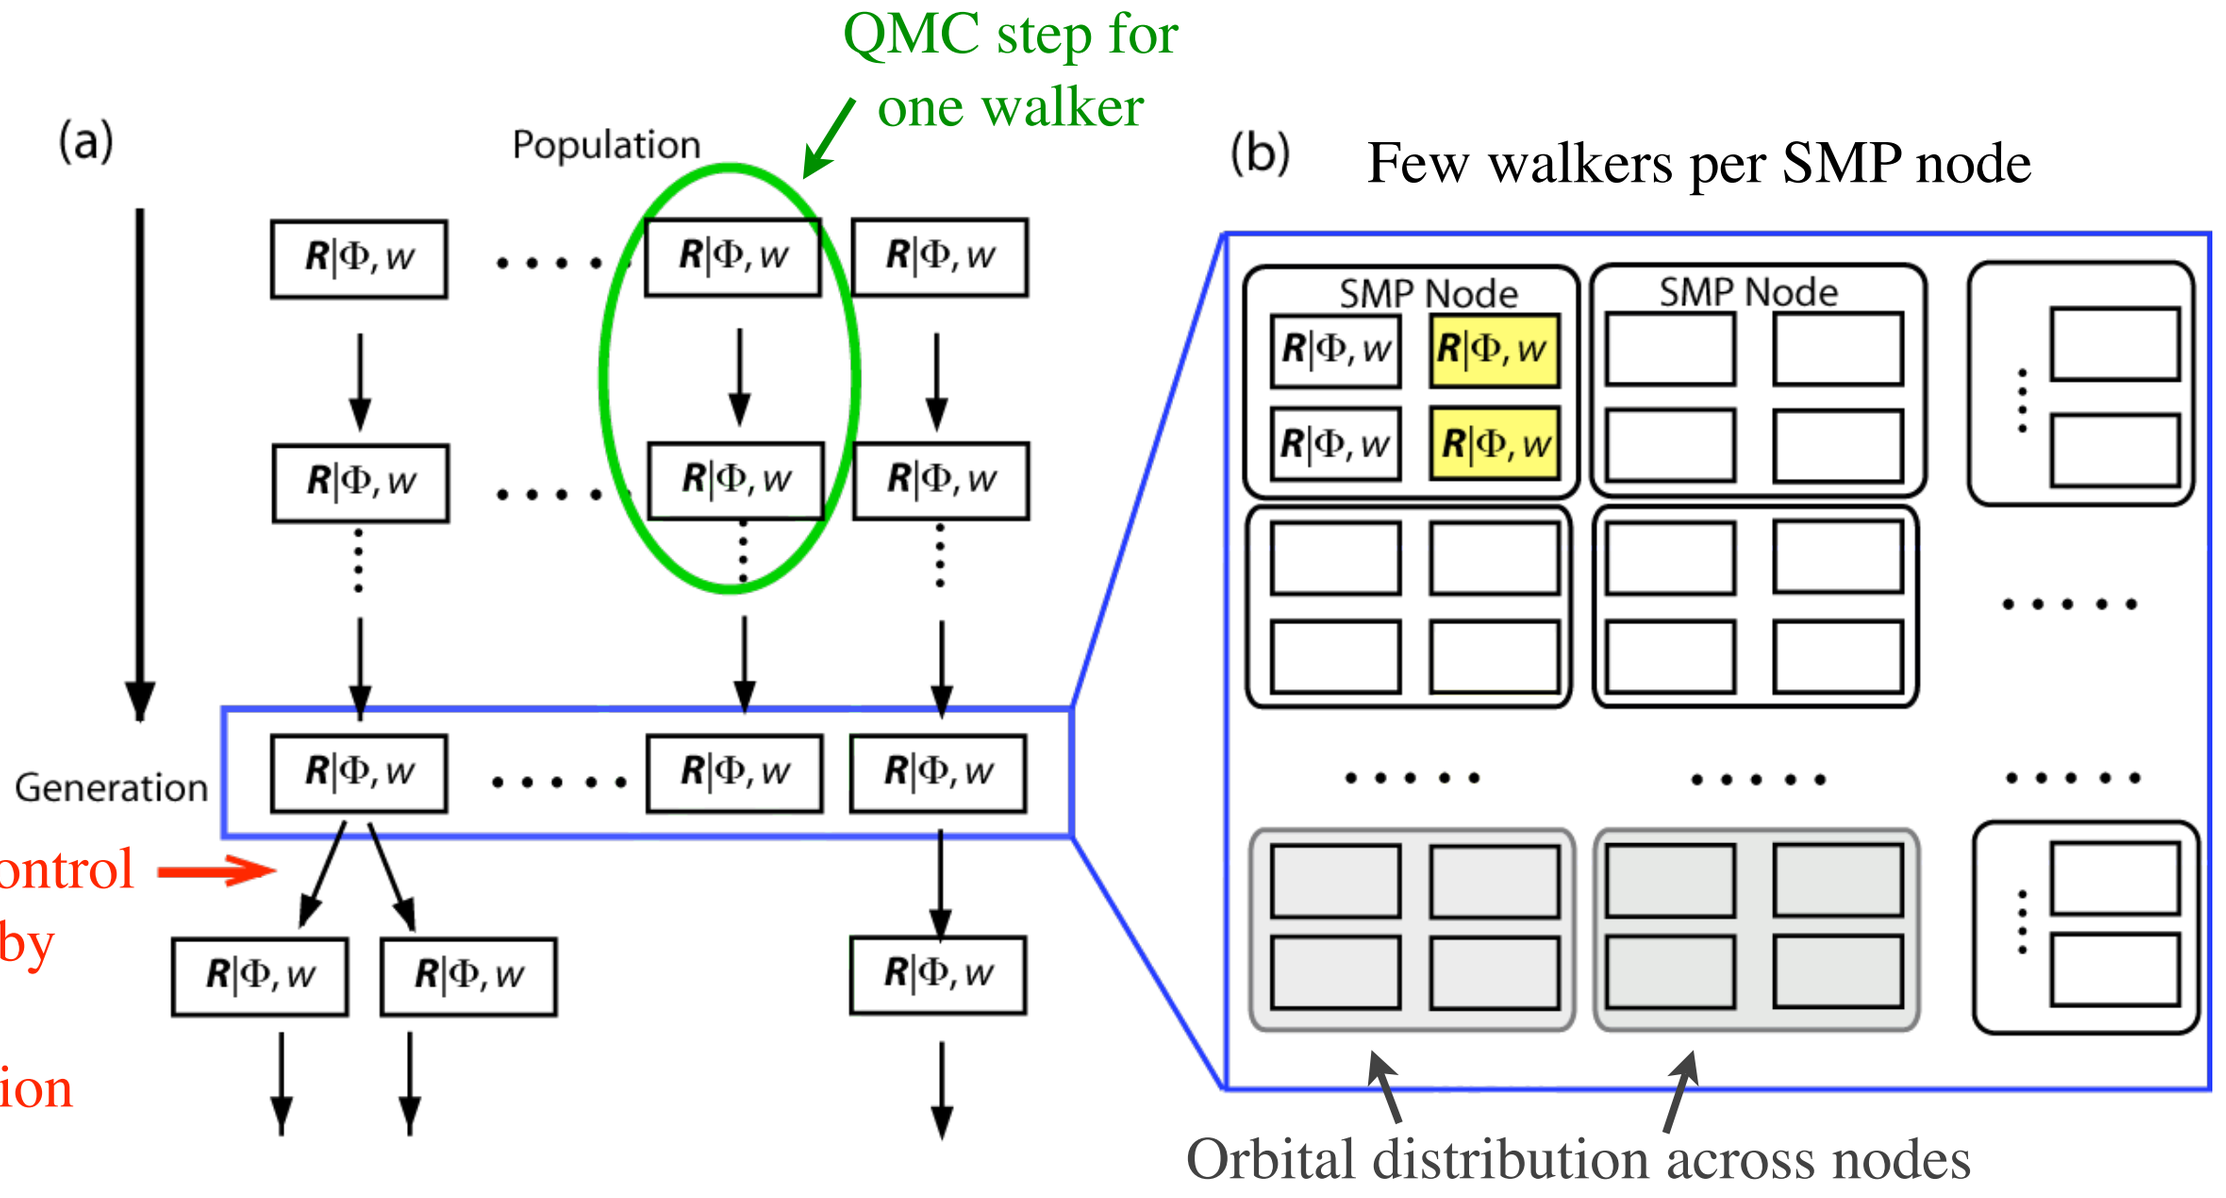

Cornell University
College of Engineering
Materials Science and Engineering

rhennig@cornell.edu

Potential Break-Through Simulations

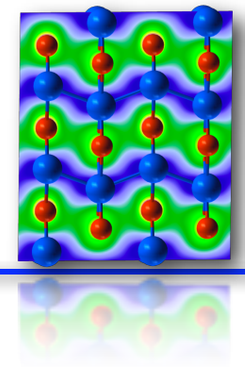

Coupled electron-ion Monte Carlo/Molecular Dynamics

- Free energies of solid/liquid phases, chemical reactions/catalysis
- Reduce wall clock time per MD time step by
  using minimum number of walkers required for accuracy
- Limit of 1 walker/core is reached for ≈ 10,000–100,000 cores

Transition Metal Oxides

- Magnetism, Mott transition, defects
- Requires system size of 10,000 electrons (large memory for wave function)
- Does not necessarily need increased number of walkers

Limit number of walkers to what is needed
Reduce memory requirements

Algorithms for Petascale

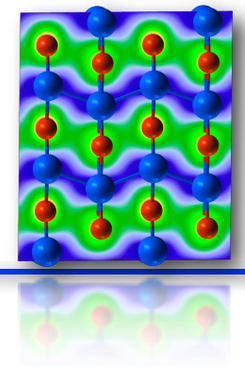

Exploit shared-memory capability of single node

- *Increase* in number of cores per node with *constant* memory per node
- *Fixed* average number of walkers per node

  ⇒ Multithreading of trial wave function evaluation
  - Single particle orbitals, dense/sparse determinant updates, Jastrow evaluation
  - Careful scheduling of tasks with varying computation loads

Multi-level parallelism

- Averages of set of variables or parameters distributed of processor groups
  - Twist averaging, nudged-elastic band calculations, *etc*.

Challenges

- Load balancing (remove synchronization barriers and blocking communication)
- Fault tolerance (invalidate failed nodes and remove walkers from pool)
- Orbital storage for very large systems
  (share read-only data on SMP nodes, distribute orbitals between nodes, overlay communication with computation)
- Pseudorandom number generation (test of random number streams in SPRNG)

Load-Balancing for Petascale

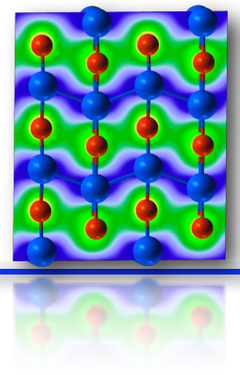

Simple and effective implementation by manager/worker model

- Ceperley, Kim *et al*.
- At set wall-clock intervals, nodes send their number of walkers and properties
- Nodes continue propagating local walker population
- Manager decides on optimal walker redistribution and send messages
- Workers check occasionally for posted messages
- Pointwise transfers of walkers between pairs of nodes

- For 32 cores per nodes, single walker per core should be achievable

Quantum Monte Carlo Codes

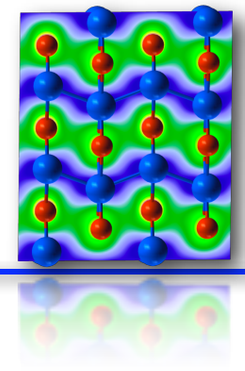

Set of five QMC codes with different approximations, algorithms etc.

- Plane-wave/spline based QMC codes:
  - (1) **Casino**
  - (2) **Champ**
  - (3) **QMCPack**
- Gaussian-based QMC code: (4) **QWalk**
- Auxiliary-field QMC code: (5) **AFQMC**

Goals

- Benchmark **accuracy** and **performance** of algorithms and codes
- Identify standard examples and procedures for petascale algorithm development

Quantum Monte Carlo Benchmarks

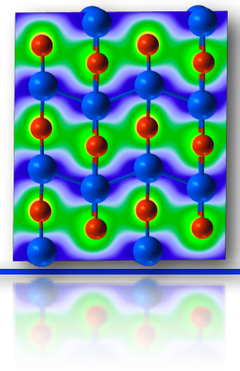

Calculations

- Accuracy of VMC sampling (no Jastrow)
- Efficiency of Jastrow factor (ee, en, een, plane-wave Jastrow)
- Accuracy and performance of QMC algorithm (time step)

Systems

- **Silicon bulk**
  - 2 and 54 atom cells
  - DF-PP (Trail & Needs, moderate cutoff) and LDA-PP (small cutoff)
- **MnO**
  - LDA-PP (moderate cutoff)
- Single water molecule
- Electron gas

QMC Benchmarks for Silicon

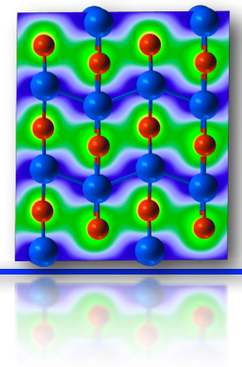

- **Silicon bulk**
  - 2 atom cell, DF-PP (Trail & Needs, moderate cutoff)
  - AFQMC uses Kleinman-Bylander form of pseudopotential (shift in energy)

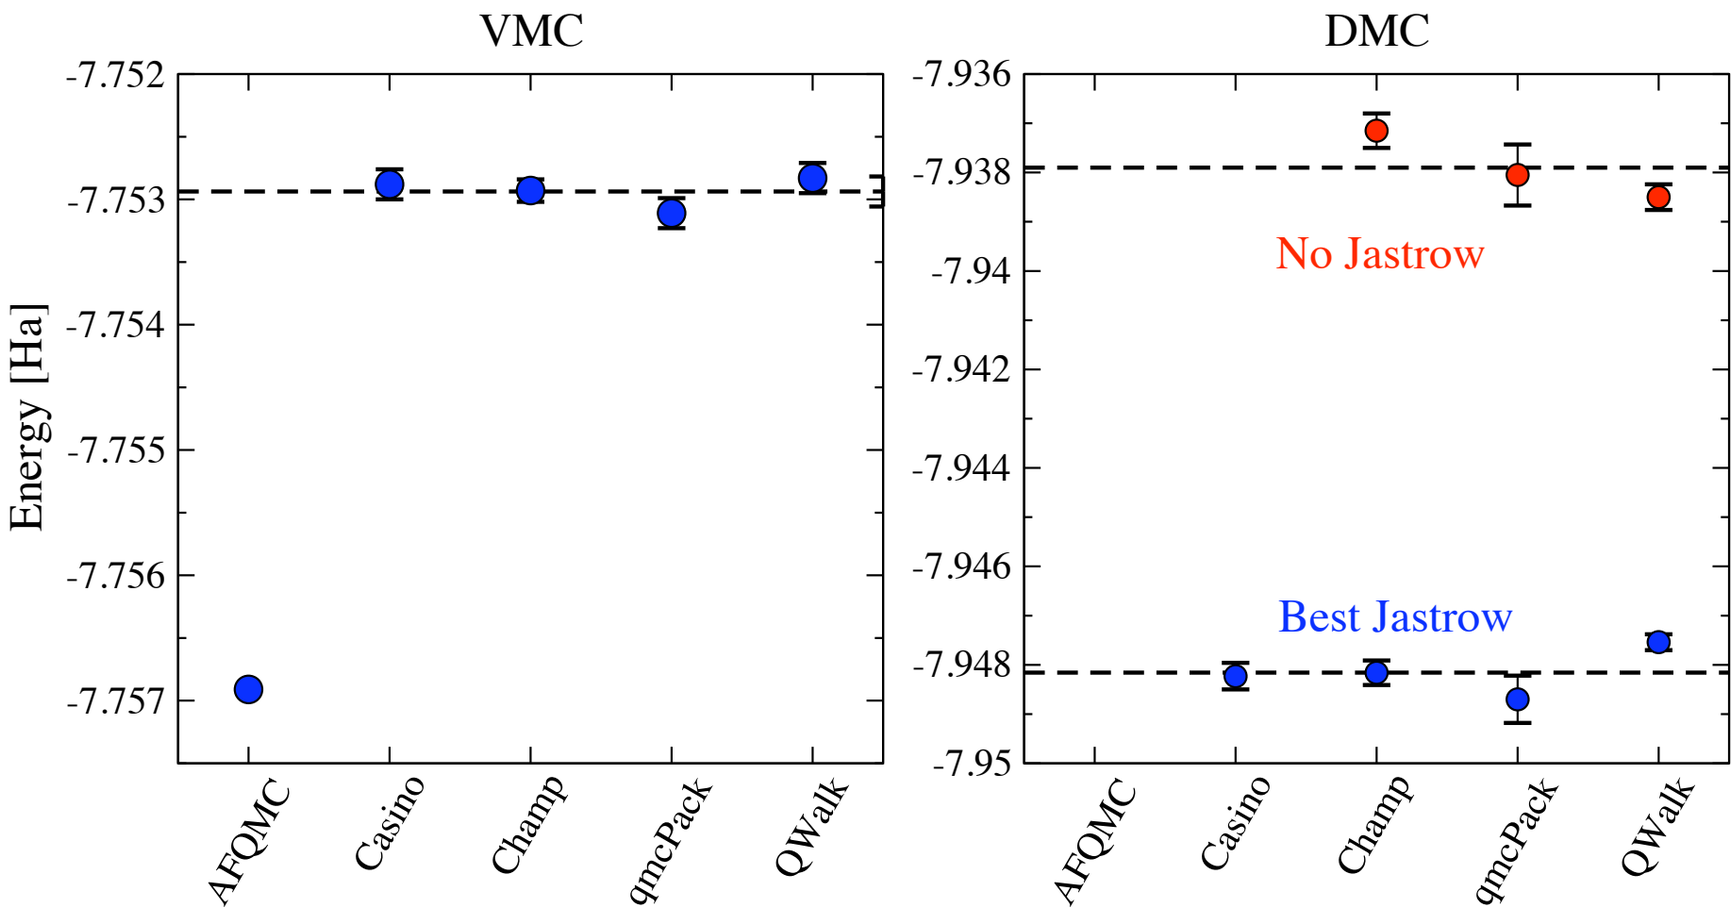

Excellent agreement of VMC and DMC energies and efficient Jastrows

QMC Benchmarks for MnO

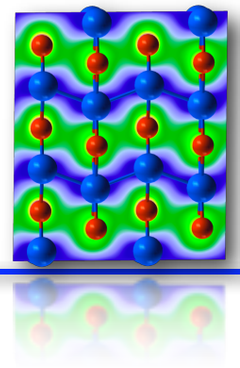

MnO

- Ferromagnetic rocksalt structure, LDA-PP (moderate cutoff)

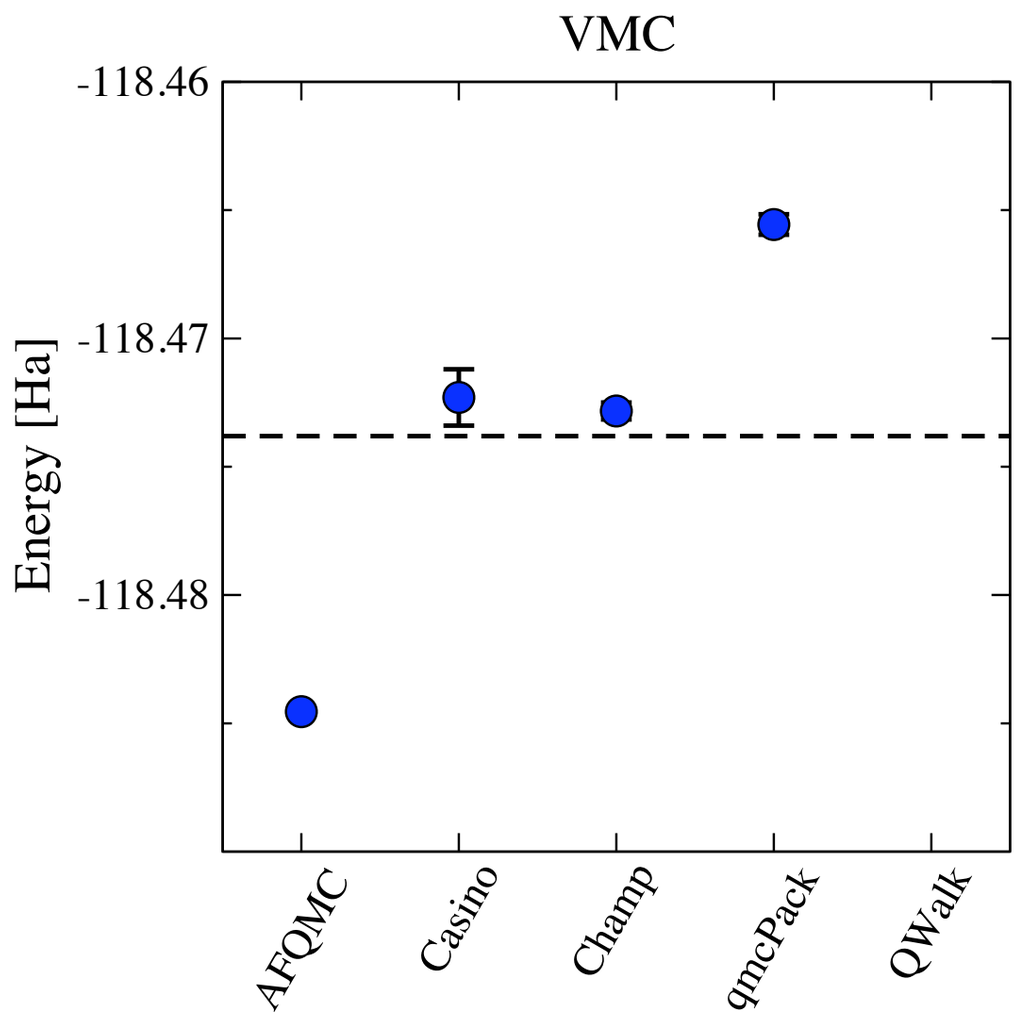

Difference in VMC energies appears to originate in pseudopotential energy

Cornell University
College of Engineering
Materials Science and Engineering

rhennig@cornell.edu

QMC Application for Molecules

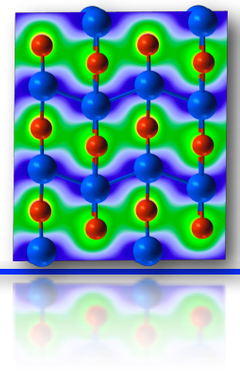

Optimization of large number of wave function parameters

- Energy optimization of Jastrow, CSF and orbital coefficient of trial wave function

The Benzene Dimer

- Simplest prototype of π-π interactions, very weak binding of 1-3 kcal/mol
- Difficult for quantum chemistry (MP2, CCSD(T), basis set convergence)
- Parallel displaced geometry shows strongest binding

Previous QMC calculations

- Diedrich, Luchow and Grimme (2005)
  PD: -3.00(38) TS: -3.58(0.38) kcal/mol
- Korth, Luchow, Grimme (2008):
  PD -1.65(42) kcal/mol
- Sorella et al.: PD 2.2(3) kcal/mol

QMC Application for Benzene

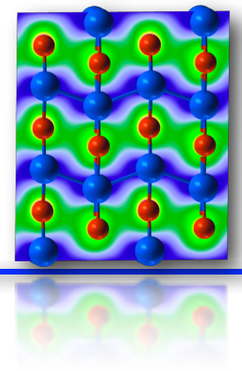

Optimization of large number of wave function parameters

- Energy optimization of Jastrow, CSF and orbital coefficient of trial wave function

Parallel Benzene Dimer

- Burkatzki, Fillipi, Dolg PP, Double zeta basis, up to 1,000 parameters
- Test case demonstrates **importance of Jastrow, CSF & orbital optimization**

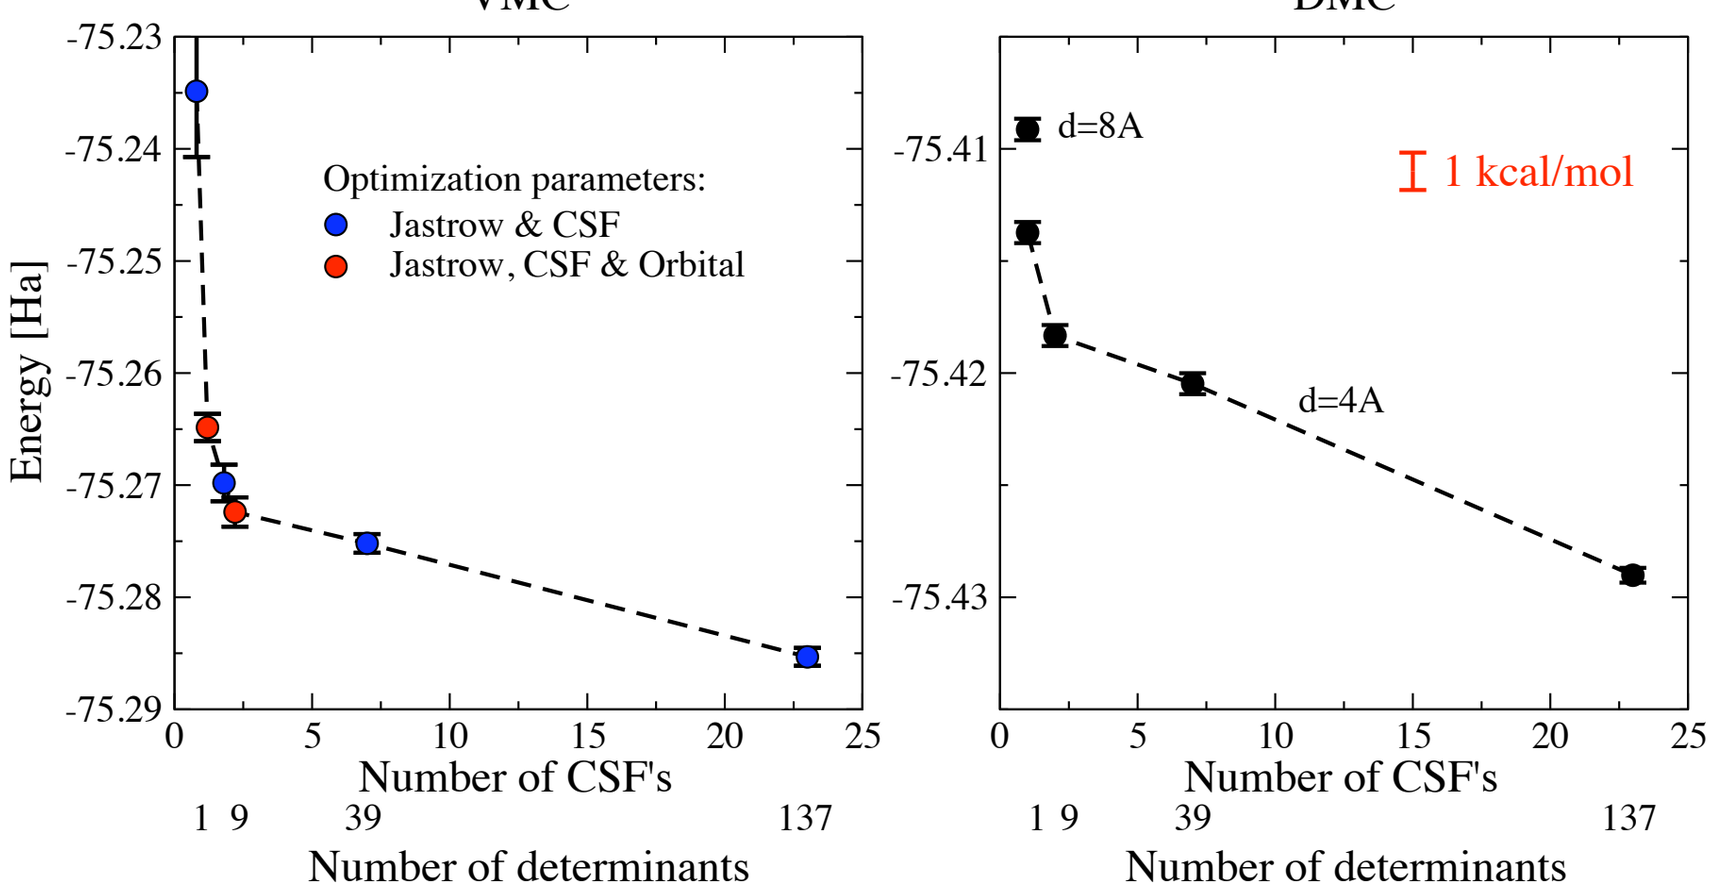

Cornell University
College of Engineering
Materials Science and Engineering

rhennig@cornell.edu

The QMC Petascale Project

Richard G. Hennig

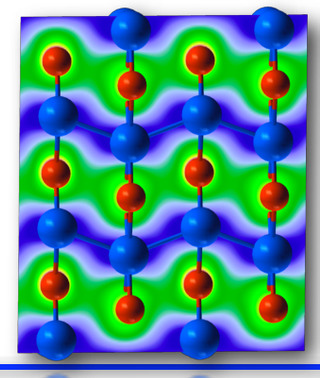

What will a petascale computer look like?
What are the limitations of current QMC algorithms for petascale computers?
How can Quantum Monte Carlo algorithms take advantage of petascale?

- 1 Petaflop = 10^{15} floating point operations per second
- 10^5 – 10^6 compute cores, 32 cores per node, limited memory and bandwidth
- QMC algorithms for petascale ⇒ Single walker per node
  - Exploit shared-memory capability of cores on SMP node
  - Multi-level parallelism & Load balancing
  - Memory reduction ⇒ Orbital distribution
- Comparison and benchmarking of VMC/DMC codes
- Importance of orbital optimization

Supported by DOE & NSF
Computational resources provided by NCCS, CCNI, NCSA, NERSC and OSC

Cornell University
College of Engineering
Materials Science and Engineering

rhennig@cornell.edu