Quantum Monte Carlo Methods

Cyrus Umrigar

Physics Department, Cornell University, Ithaca.

Email: CyrusUmrigar@cornell.edu

April - May, 2015, Université de Paris, Jussieu

Lecture 1
Basics of MC

1. What is quantum Monte Carlo?
2. Essence of variational and projector Monte Carlo methods
3. Early history of MC and QMC.
4. Central limit theorem, Chebyshev inequality
5. Monte Carlo vs. deterministic integration
    1. Importance Sampling
6. Pseudo-random vs. quasi-random numbers
7. Sampling nonuniform probability densities
    1. Transformation method
    2. Rejection method
    3. Metropolis-Hastings algorithm (part of next lecture)
8. Unbiased estimators

Lecture 2

Variational Monte Carlo and Metropolis-Hastings Algorithm

1. Ingredients needed for accurate and efficient variational Monte Carlo
2. Metropolis-Hastings Algorithm
    1. Markov chains/matrices/kernels
    2. Stationarity condition and detailed balance condition
    3. Choice of acceptance matrix/kernel
    4. Choice of proposal matrix/kernel
    5. Various observations and connection to heat-bath/Gibbs sampler
    6. Optimizing the Markov matrix for finite spaces
3. Estimating errors of autocorrelated variables
4. Example of a variational wavefunction

Optimization of Many-body Wavefunctions (if there is time)

1. Measures of goodness of many-body wave functions
2. Pros and cons of optimizing energy versus variance of energy
3. Optimization methods
    1. Newton method
    2. linear method
    3. augmented hessian method
    4. perturbation theory
    5. stabilization
    6. optimization of linear combination of energy and variance

Lecture 3

1. Different requirements for trial versus guiding wave functions
2. Outline of variational and projector Monte Carlo methods
3. Importance sampling in projector Monte Carlo methods
4. Forms of projector Monte Carlo (PMC)
5. Manifestation of sign problem in different projector Monte Carlo methods
6. Full configuration interaction QMC (FCIQMC) and semistochastic QMC (SQMC)
7. Fixed-node approximation in discrete-space projector Monte Carlo methods
8. Diffusion Monte Carlo without and with importance sampling
9. Expectation values of various classes of operators
10. Weighted branching random walks
11. Population control error and how to remove it
12. Pure state versus finite temperature MC methods. Schematic of VMC, PMC, PIMC and reptation MC walks
13. Path-integral Monte Carlo (PIMC)
14. Reptation Monte Carlo (a hybrid between PMC and PIMC)

References

There are no books or review papers with a strong overlap with what I will cover, but here are some that may be useful.

Review papers

1. *Quantum Monte Carlo Simulations of Solids*, W. M. C. Foulkes, L. Mitas, R. J. Needs and G. Rajagopal, Rev. Mod Phys. **73**, 33, (2001).
2. *Applications of quantum Monte Carlo methods in condensed systems*, Jindrich Kolorenc and Lubos Mitas, Reports on Progress in Physics, **74**, 026502 (2011).
3. *Path-integrals in the theory of condensed helium*, D. M. Ceperley, Rev. Mod Phys. **67**, 2, (1995).

Books

1. *Monte Carlo Methods*, M.H. Kalos and P. A. Whitlock, Wiley-Interscience, 2^{nd} edition (2008).
2. *Monte Carlo Methods in Ab-initio Quantum Chemistry*, B.L. Hammond, W.A. Lester, and P.J. Reynolds, World Scientific (1994).
3. *Quantum Monte Carlo Methods in Physics and Chemistry*, ed. by M. P. Nightingale and C. J. Umrigar, NATO ASI Ser. C 525, Kluwer (1999).

What is Quantum Monte Carlo?

Stochastic implementation of the power method for projecting out the dominant eigenvector of a matrix or integral kernel.

"Dominant state" means state with largest absolute eigenvalue.

If we repeatedly multiply an arbitrary vector, not orthogonal to the dominant state, by the matrix, we will eventually project out the dominant state.

QMC methods are used only when the number of states is so large ($> 10^{10}$) that it is not practical to store even a single vector in memory. Otherwise use exact diagonalization method, e.g., Lanczos. So, at each MC generation, only a sample of the states are stored.

QMC methods are used not only in a large discrete space but also in a continuously infinite space. Hence "matrix or integral kernel" above. In the interest of brevity I will use either discrete or continuous language (sums and matrices or integrals and integral kernels), but much of what is said will apply to both situations.

Definitions

Given a complete or incomplete basis: $\{|\phi_i\rangle\}$, either discrete or continuous

$$\text{Exact} \quad |\Psi_0\rangle = \sum_i e_i|\phi_i\rangle, \quad \text{where,} \quad e_i = \langle\phi_i|\Psi_0\rangle$$

$$\text{Trial} \quad |\Psi_T\rangle = \sum_i t_i|\phi_i\rangle, \quad \text{where,} \quad t_i = \langle\phi_i|\Psi_T\rangle$$

$$\text{Guiding} \quad |\Psi_G\rangle = \sum_i g_i|\phi_i\rangle, \quad \text{where,} \quad g_i = \langle\phi_i|\Psi_G\rangle$$

(If basis incomplete then "exact" means "exact in that basis".)

Ψ_T used to calculate variational and mixed estimators of operators $\hat{A}$, i.e., $\langle\Psi_{\rm T}|\hat{A}|\Psi_{\rm T}\rangle/\langle\Psi_{\rm T}|\Psi_{\rm T}\rangle$, $\langle\Psi_{\rm T}|\hat{A}|\Psi_0\rangle/\langle\Psi_{\rm T}|\Psi_0\rangle$

Ψ_G used to alter the probability density sampled, i.e., $\Psi_{\rm G}^2$ in VMC, $\Psi_{\rm G}\Psi_0$ in PMC.

$\Psi_{\rm G}$ must be such that $g_i \neq 0$ if $e_i \neq 0$. If $\Psi_{\rm T}$ also satisfies this condition then $\Psi_{\rm G}$ can be chosen to be $\Psi_{\rm T}$. Reasons to have $\Psi_{\rm G} \neq \Psi_{\rm T}$ are: a) rapid evaluation of "local energy", b) have finite-variance estimators. To simplify expressions, we use $\Psi_{\rm G} = \Psi_{\rm T}$ or $\Psi_{\rm G} = 1$ in what follows.

Variational MC

$$E_V = \frac{\langle \Psi_{\rm T}|\hat{H}|\Psi_{\rm T}\rangle}{\langle \Psi_{\rm T}|\Psi_{\rm T}\rangle} = \frac{\sum_{ij}^{N_{\rm st}} \langle \Psi_{\rm T}|\phi_i\rangle \ \langle \phi_i|\hat{H}|\phi_j\rangle \ \langle \phi_j|\Psi_{\rm T}\rangle}{\sum_{i}^{N_{\rm st}} \langle \Psi_{\rm T}|\phi_k\rangle \ \langle \phi_k|\Psi_{\rm T}\rangle}$$

$$= \frac{\sum_{ij}^{N_{\rm st}} t_i H_{ij} t_j}{\sum_{k}^{N_{\rm st}} t_k^2} = \sum_{i}^{N_{\rm st}} \frac{t_i^2}{\sum_{k}^{N_{\rm st}} t_k^2} \frac{\sum_{j}^{N_{\rm st}} H_{ij} t_j}{t_i}$$

$$= \sum_{i}^{N_{\rm st}} \frac{t_i^2}{\sum_{k}^{N_{\rm st}} t_k^2} E_{\rm L}(i) = \frac{\left[\sum_{i}^{N_{\rm MC}} E_{\rm L}(i)\right]_{\Psi_{\rm T}^2}}{N_{\rm MC}} \xrightarrow{\Psi_G \neq \Psi_T} \frac{\left[\sum_{i}^{N_{\rm MC}} \left(\frac{t_i}{g_i}\right)^2 E_{\rm L}(i)\right]_{\Psi_{\rm G}^2}}{\left[\sum_{k}^{N_{\rm MC}} \left(\frac{t_k}{g_k}\right)^2\right]_{\Psi_{\rm G}^2}}$$

Sample probability density function $\frac{g_i^2}{\sum_k^{N_{\rm st}} g_k^2}$ using Metropolis-Hastings, if $\Psi_{\rm G}$ complicated.

Value depends only on $\Psi_{\rm T}$. Statistical error depend on $\Psi_{\rm T}$ and $\Psi_{\rm G}$.

Energy bias and statistical error vanish as $\Psi_{\rm T} \to \Psi_0$.

For fixed Ψ_T, $\Psi_G = \Psi_T$ does not minimize statistical fluctuations!

In fact need $\Psi_G \neq \Psi_T$ at times to get finite variance. $\Psi_G = \Psi_T$ does give unbiased estimator.

Projector MC

Pure and Mixed estimators for energy are equal: $E_0 = \dfrac{\langle \Psi_0 | \hat{H} | \Psi_0 \rangle}{\langle \Psi_0 | \Psi_0 \rangle} = \dfrac{\langle \Psi_0 | \hat{H} | \Psi_{\rm T} \rangle}{\langle \Psi_0 | \Psi_{\rm T} \rangle}$

Projector: $|\Psi_0\rangle = \hat{P}(\infty)\,|\Psi_{\rm T}\rangle = \lim_{n\to\infty} \hat{P}^n(\tau)\,|\Psi_{\rm T}\rangle$

$$E_0 = \frac{\langle \Psi_0 | \hat{H} | \Psi_{\rm T} \rangle}{\langle \Psi_0 | \Psi_{\rm T} \rangle} = \frac{\sum_{ij}^{N_{\rm st}} \langle \Psi_0 | \phi_i \rangle \, \langle \phi_i | \hat{H} | \phi_j \rangle \, \langle \phi_j | \Psi_{\rm T} \rangle}{\sum_{k}^{N_{\rm st}} \langle \Psi_0 | \phi_k \rangle \, \langle \phi_k | \Psi_{\rm T} \rangle}$$

$$= \frac{\sum_{ij}^{N_{\rm st}} e_i H_{ij} t_j}{\sum_{k}^{N_{\rm st}} e_k t_k} = \sum_{i}^{N_{\rm st}} \frac{e_i t_i}{\sum_{k}^{N_{\rm st}} e_k t_k} \, \frac{\sum_{j}^{N_{\rm st}} H_{ij} t_j}{t_i}$$

$$= \sum_{i}^{N_{\rm st}} \frac{e_i t_i}{\sum_{k}^{N_{\rm st}} e_k t_k} \, E_{\rm L}(i) = \frac{\left[\sum_{i}^{N_{\rm MC}} E_{\rm L}(i)\right]_{\Psi_{\rm T}\Psi_0}}{N_{\rm MC}} \xrightarrow{\Psi_G \neq \Psi_T} \frac{\left[\sum_{i}^{N_{\rm MC}} \left(\frac{t_i}{g_i}\right) E_{\rm L}(i)\right]_{\Psi_{\rm G}\Psi_0}}{\left[\sum_{k}^{N_{\rm MC}} \left(\frac{t_k}{g_k}\right)\right]_{\Psi_{\rm G}\Psi_0}}$$

Sample $e_i g_i / \sum_k^{N_{\rm st}} e_k g_k$ using projector.

For exact PMC, value indep. of $\Psi_{\rm T}$, $\Psi_{\rm G}$, statistical error depends on $\Psi_{\rm T}$, $\Psi_{\rm G}$.
(For FN-PMC, value depends on $\Psi_{\rm G}$, statistical error on $\Psi_{\rm T}, \Psi_{\rm G}$.)
(For FN-DMC, value depends on nodes of $\Psi_{\rm G}$, statistical error on $\Psi_{\rm T}, \Psi_{\rm G}$.)
Statistical error vanishes as $\Psi_{\rm T} \to \Psi_0$.
For fixed Ψ_T, $\Psi_G = \Psi_T$ does not minimize statistical fluctuations!

Variational and Projector MC

$$E_V = \frac{\left[\sum_i^{N_{\rm MC}} \left(\frac{t_i}{g_i}\right)^2 E_{\rm L}(i)\right]_{\Psi_{\rm G}^2}}{\left[\sum_k^{N_{\rm MC}} \left(\frac{t_k}{g_k}\right)^2\right]_{\Psi_{\rm G}^2}} \quad \text{(Value depends on } \Psi_{\rm T}\text{, error } \Psi_{\rm T}, \Psi_{\rm G}\text{)}$$

$$E_0 = \frac{\left[\sum_i^{N_{\rm MC}} \left(\frac{t_i}{g_i}\right) E_{\rm L}(i)\right]_{\Psi_{\rm G}\Psi_0}}{\left[\sum_k^{N_{\rm MC}} \left(\frac{t_k}{g_k}\right)\right]_{\Psi_{\rm G}\Psi_0}} \quad \text{(Value exact}^\dagger\text{. Error depends on } \Psi_{\rm T}, \Psi_{\rm G}\text{.)}$$

$$E_{\rm L}(i) = \frac{\sum_j^{N_{\rm st}} H_{ij} t_j}{t_i}$$

In both VMC and PMC weighted average of the *configuration value of* $\hat{H}$ aka *local energy,* $E_{\rm L}(i)$, but from points sampled from different distributions.

This is practical for systems that are large enough to be interesting if

1. $t_i = \langle \phi_i | \Psi_{\rm T} \rangle$ can be evaluated in polynomial time, say N^3
2. the sum in $E_{\rm L}(i)$ can be done quickly, i.e., $\hat{H}$ is sparse (if space discrete) or semi-diagonal (if space continuous).

† In practice, usually necessary to make approximation (e.g. FN) and value depends on $\Psi_{\rm G}$.

Variational Monte Carlo in Real Space

W. L. McMillan, Phys. Rev. **138**, A442 (1965)

Monte Carlo is used to perform the many-dimensional integrals needed to calculate quantum mechanical expectation values. e.g.

$$
\begin{aligned}
E_T &= \frac{\int d\mathbf{R}\, \Psi_{\rm T}^*(\mathbf{R})\, \mathcal{H}\, \psi_{\rm T}(\mathbf{R})}{\int d\mathbf{R}\, \psi_{\rm T}^2(\mathbf{R})} \\
&= \int d\mathbf{R}\, \frac{\psi_{\rm T}^2(\mathbf{R})}{\int d\mathbf{R}\, \psi_{\rm T}^2(\mathbf{R})}\, \frac{\mathcal{H}\psi_{\rm T}(\mathbf{R})}{\psi_{\rm T}(\mathbf{R})} \\
&= \frac{1}{N}\sum_i \frac{\mathcal{H}\Psi_{\rm T}(\mathbf{R}_i)}{\Psi_{\rm T}(\mathbf{R}_i)} = \frac{1}{N}\sum_i E_L(\mathbf{R}_i)
\end{aligned}
$$

Energy is obtained as an arithmetic sum of the *local energies* $E_L(\mathbf{R}_i)$ evaluated for configurations sampled from $\psi_{\rm T}^2(\mathbf{R})$ using a generalization of the Metropolis method. If $\psi_{\rm T}$ is an eigenfunction, the $E_L(\mathbf{R}_i)$ do not fluctuate. Accuracy of VMC depends crucially on the quality of $\psi_{\rm T}(\mathbf{R})$. Diffusion MC does better by projecting onto ground state.

Rest of this lecture

Now that you know the essence of quantum Monte Carlo methods, for the rest of this lecture we will discuss basic concepts that underlie both classical and quantum Monte Carlo methods, e.g., the central limit theorem, techniques for sampling various distributions, importance sampling for reducing statistical error, calculation of unbiased estimator, ...

Then in the rest of the lectures we will continue our study of quantum Monte Carlo methods.

When to use Monte Carlo Methods

Monte Carlo methods: A class of computational algorithms that rely on repeated random sampling to compute results.
A few broad areas of applications are:

1. physics
2. chemistry
3. engineering
4. finance and risk analysis

When are MC methods likely to be the methods of choice?

1. When the problem is many-dimensional and approximations that factor the problem into products of lower dimensional problems are inaccurate.
2. A less important reason is that if one has a complicated geometry, a MC algorithm may be simpler than other choices.

Obvious drawback of MC methods: There is a statistical error.
Frequently there is a tradeoff between statistical error and systematic error and one needs to find the best compromise.

Physics/Chemistry applications of Quantum Monte Carlo

Some systems to which they have been applied are:

- strongly correlated systems (Hubbard, Anderson, t-J, ... models)
- quantum spin systems (Ising, Heisenberg, xy, ... models),
- liquid and solid helium, liquid-solid interface, droplets
- energy and response of homogeneous electron gas in 2-D and 3-D
- nuclear structure
- lattice gauge theory
- atomic clusters
- electronic structure calculations of atoms, molecules, solids, quantum dots, quantum wires

- both to zero temperature (pure states) and finite temperature problems, but in these lectures we will mostly discuss zero temperature methods

MC Simulations versus MC calculations

One can distinguish between two kinds of algorithms:

1. The system being studied is stochastic and the stochasticity of the algorithm mimics the stochasticity of the actual system. e.g. study of neutron transport and decay in nuclear reactor by following the trajectories of a large number of neutrons. Such problems are suitable for MC algorithms in a very obvious way.
2. Much more interesting are applications where the system being studied is not stochastic, but nevertheless a stochastic algorithm is the most efficient, or the most accurate, or the only feasible method for studying the system. e.g. the solution of a PDE in a large number of variables, e.g., the solution of the Schrödinger equation for an N-electron system, with say $N = 100$ or 1000. (Note: The fact that the wavefunction has a probabilistic interpretation has *nothing* to do with the stochasticity of the algorithm. The wavefunction itself is perfectly deterministic.)

I prefer to use the terminology that the former are MC simulations whereas the latter are MC calculations but not everyone abides by that terminology.

Early Recorded History of Monte Carlo

1777 Comte de Buffon: If a needle of length L is thrown at random onto a plane ruled with straight lines a distance $d (d > L)$ apart, then the probability P of the needle intersecting one of those lines is $P = \frac{2L}{\pi d}$.
Laplace: This could be used to compute π (inefficiently).

1930s First significant scientific application of MC: Enrico Fermi used it for neutron transport in fissile material.
Segre: "Fermi took great delight in astonishing his Roman colleagues with his "too-good-to-believe" predictions of experimental results."

1940s Monte Carlo named by Nicholas Metropolis and Stanislaw Ulam

1953 Algorithm for sampling any probability density
Metropolis, Rosenbluth, Rosenbluth, Teller and Teller
(generalized by Hastings in 1970)

1962,1974 First PMC calculations, Kalos, and, Kalos, Levesque, Verlet.

1965 First VMC calculations (of liquid He), Bill McMillan.

Central Limit Theorem

de Moivre (1733), Laplace (1812), Lyapunov (1901), Pólya (1920)

Let $X_1, X_2, X_3, \cdots, X_N$ be a sequence of N independent random variables sampled from a probability density function with a finite expectation value, μ, and variance σ^2. The central limit theorem states that as the sample size N increases, the probability density of the sample average of these random variables approaches the normal distribution, $\frac{1}{\sqrt{2\pi}\sigma} e^{-(x-\mu)^2/(2\sigma^2/N)}$, with a mean μ, and variance σ^2/N, irrespective of the original probability density function.

Law of Large Numbers

Cardano, Bernouli, Borel, Cantelli, Kolmogorov, Khinchin

Even if the variance is infinite, if the expected value is finite, the sample means will converge to the expected value but usual error estimates go down slower than $1/\sqrt{N}$ and do not imply usual confidence intervals. Beware of skewed densities that have ∞ variance!

Chebychev Inequality

The Central Limit Theorem by itself does not tell you how quickly the averages converge to a Gaussian distribution.

For an arbitrary distribution with mean μ and variance σ^2, we have much weaker bounds given by Chebychev's inequality:

The probability of a variable lying between $\mu - n\sigma$ and $\mu + n\sigma$ is $> 1 - 1/n^2$.

Prob. of being within 1σ of μ is $\geq$ 0% versus 68.3% for Gaussian
Prob. of being within 2σ of μ is $\geq$ 75% versus 95.4% for Gaussian
Prob. of being within 3σ of μ is $\geq$ 89% versus 99.7% for Gaussian

The worst case occurs for a distribution with probability $1 - 1/n^2$ at μ and probability $1/2n^2$ at $\mu - n\sigma$ and $\mu + n\sigma$.

Monte Carlo versus Deterministic Integration methods

Deterministic Integration Methods:

Integration Error, ϵ, using $N_{\rm int}$ integration points:

1-dim Simpson rule: $\epsilon \propto N_{\rm int}^{-4}$, (provided derivatives up to 4^{th} exist)

d-dim Simpson rule: $\epsilon \propto N_{\rm int}^{-4/d}$, (provided derivatives up to 4^{th} exist)

So, for a given error, N and so the computer time increases exponentially with d, since $N \propto (\frac{1}{\epsilon})^{d/4}$.

Monte Carlo:

$\epsilon \sim \sigma(T_{\rm corr}/N_{\rm int})^{1/2}$, independent of dimension!, according to the central limit theorem since width of gaussian decreases as $(T_{\rm corr}/N_{\rm int})^{1/2}$ provided that the variance of the integrand is finite. ($T_{\rm corr}$ is the autocorrelation time.)

Roughly, Monte Carlo becomes advantageous for $d > 8$.

For a many-body wavefunction $d = 3N$ and can be a few thousand!

Scaling with number of electrons

Simpson's rule integration

$$\epsilon = \frac{c}{N_{\rm int}^{4/d}} = \frac{c}{N_{\rm int}^{4/3N_{\rm elec}}}$$

$$N_{\rm int} = \left(\frac{c}{\epsilon}\right)^{\frac{3N_{\rm elec}}{4}} \qquad \text{exponential in } N_{\rm elec}$$

Monte Carlo integration

$$\epsilon = \sigma\sqrt{\frac{N_{\rm elec}}{N_{\rm MC}}}$$

$$N_{\rm MC} = \left(\frac{\sigma}{\epsilon}\right)^2 N_{\rm elec} \qquad \text{linear in } N_{\rm elec}$$

(For both methods, computational cost is higher than this since the cost of evaluating the wavefunction increases with $N_{\rm elec}$, e.g., as $N_{\rm elec}^3$, (better if one uses "linear scaling"; worse if one increases $N_{\rm det}$ with $N_{\rm elec}$.))

Monte Carlo Integration

$$I = \int_V f(x)dx = V\,\overline{f} \pm V\sqrt{\frac{\overline{f^2} - \overline{f}^2}{N-1}}$$

where $\overline{f} = \frac{1}{N}\sum_i^N f(x_i), \quad \overline{f^2} = \frac{1}{N}\sum_i^N f^2(x_i)$

and the points x_i are sampled uniformly in V.

Importance sampling

$$I = \int_V g(x)\frac{f(x)}{g(x)}dx = \overline{\left(\frac{f}{g}\right)} \pm \sqrt{\frac{\overline{\left(\frac{f}{g}\right)^2} - \overline{\left(\frac{f}{g}\right)}^2}{N-1}}$$

where the probability density function $g(x) \geq 0$ and $\int_V g(x)dx = 1$.
If $g(x) = 1/V$ in V then we recover original fluctuations but if $g(x)$ mimics $f(x)$ then the fluctuations are much reduced. Optimal g is $|f|$. Need: a) $g(x) \geq 0$, b) know integral of $g(x)$, and, c) be able to sample it.

Importance sampling can turn an $\infty-$variance estimator into a finite variance one!

Illustration of Importance Sampling

$f(x)$ is the function to be integrated. $g(x)$ is a function that is "similar" to $f(x)$ and has the required properties: a) $g(x) \geq 0$, b) we know integral of $g(x)$, and, c) we know how to sample it. $\int f(x)dx$ can be evaluated efficiently by sampling $g(x)$ and averaging $f(x)/g(x)$.

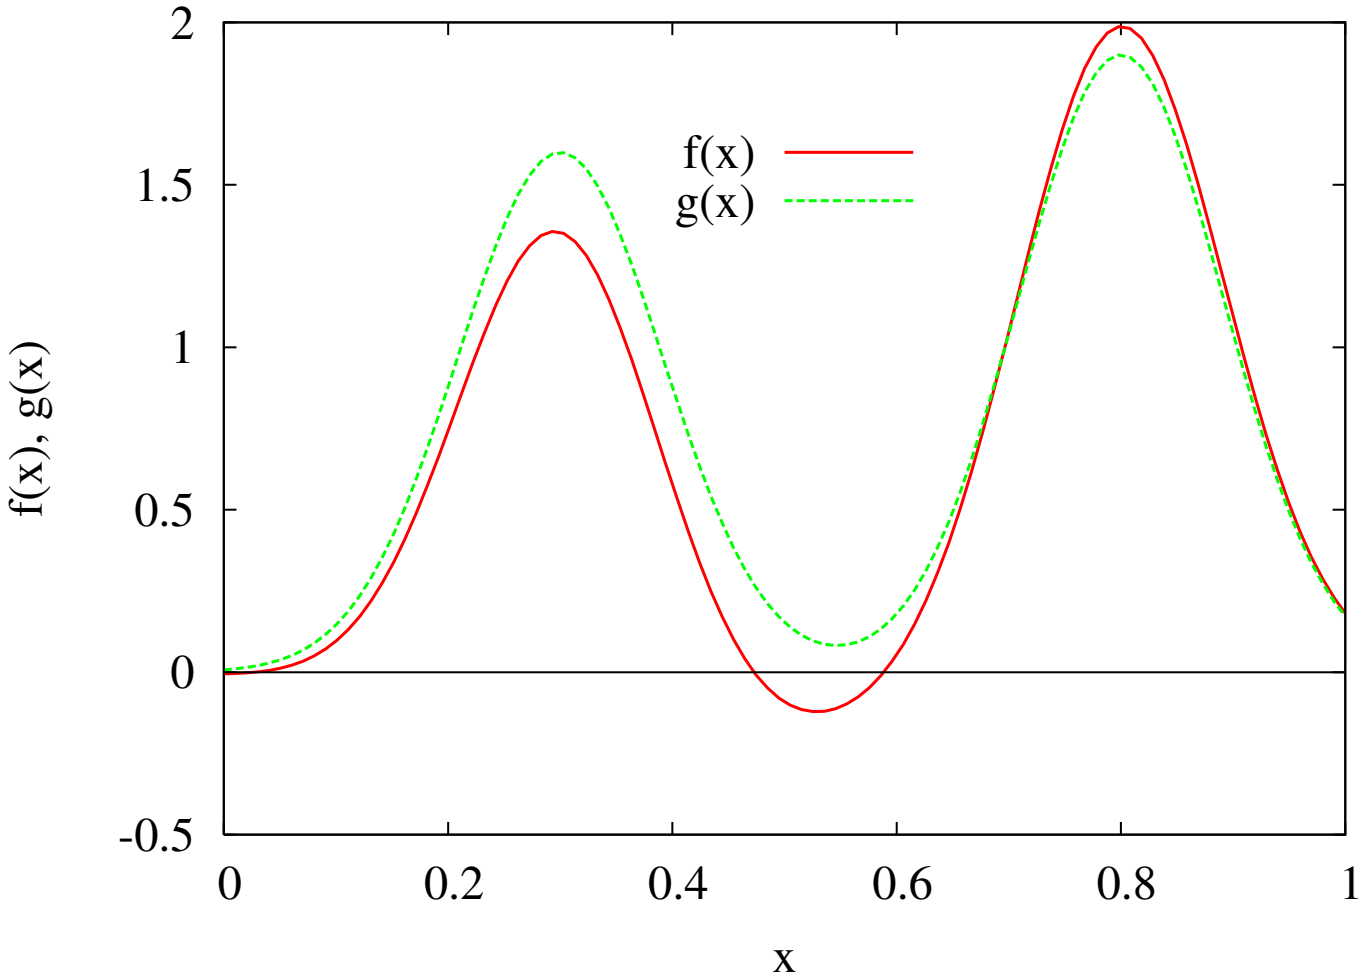

Infinite variance estimators

When variance σ^2 is finite, by the central limit theorem the average

$$F_N = \frac{\sum_{i=1}^{N} f(x_i)}{N}$$

converges for increasing N to a gaussian of width $\sigma_N = \sigma/\sqrt{N}$.
Since we have a gaussian distribution the probability of F_N lying between $\mu - n\sigma_N$ and $\mu + n\sigma_N$ is $\mathrm{erf}(n/\sqrt{2})$
F_N being within $1\sigma_N$ of the true mean is 68.3%
F_N being within $2\sigma_N$ of the true mean is 95.4%
F_N being within $3\sigma_N$ of the true mean is 99.7%.

What if the population variance $\sigma^2 = \infty$ but we do not know that beforehand? The computed sample variance will ofcourse always be finite. The practical signature of an infinite variance estimator is that the estimated σ increases with sample size, N and tends to have upward jumps. So the estimated error of the sample mean, $\sigma_N = \sigma/\sqrt{N}$, goes down more slowly than $\frac{1}{\sqrt{N}}$, or even does not go down at all.

Pseudo-random vs quasi-random numbers

Terrible misnomers!

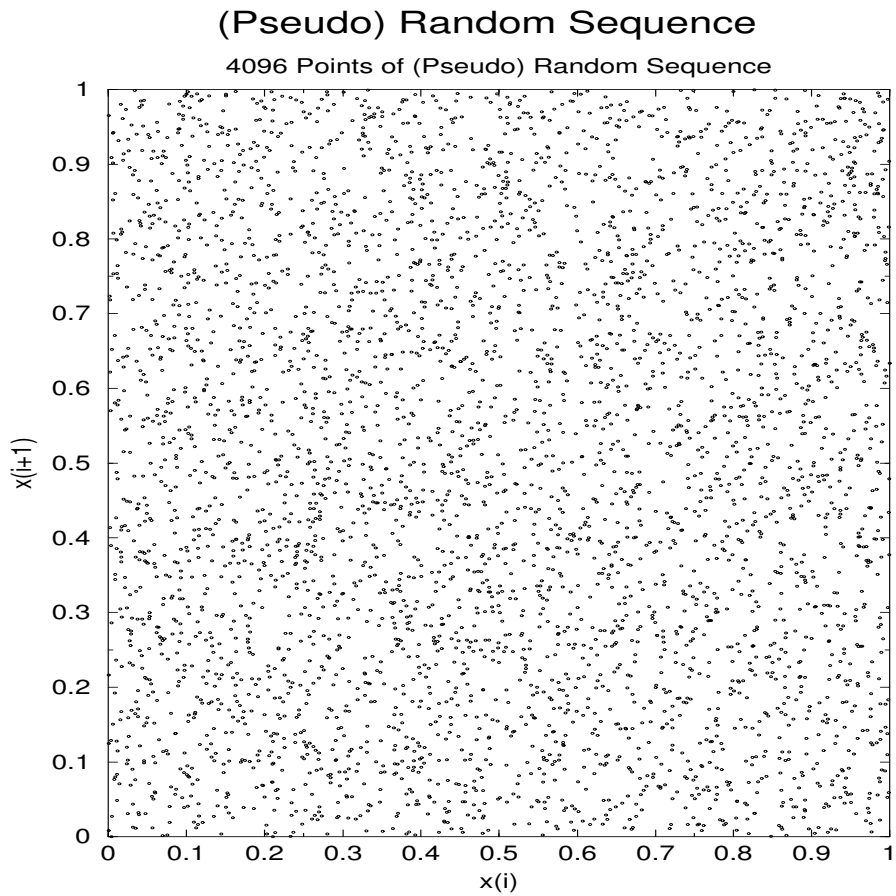

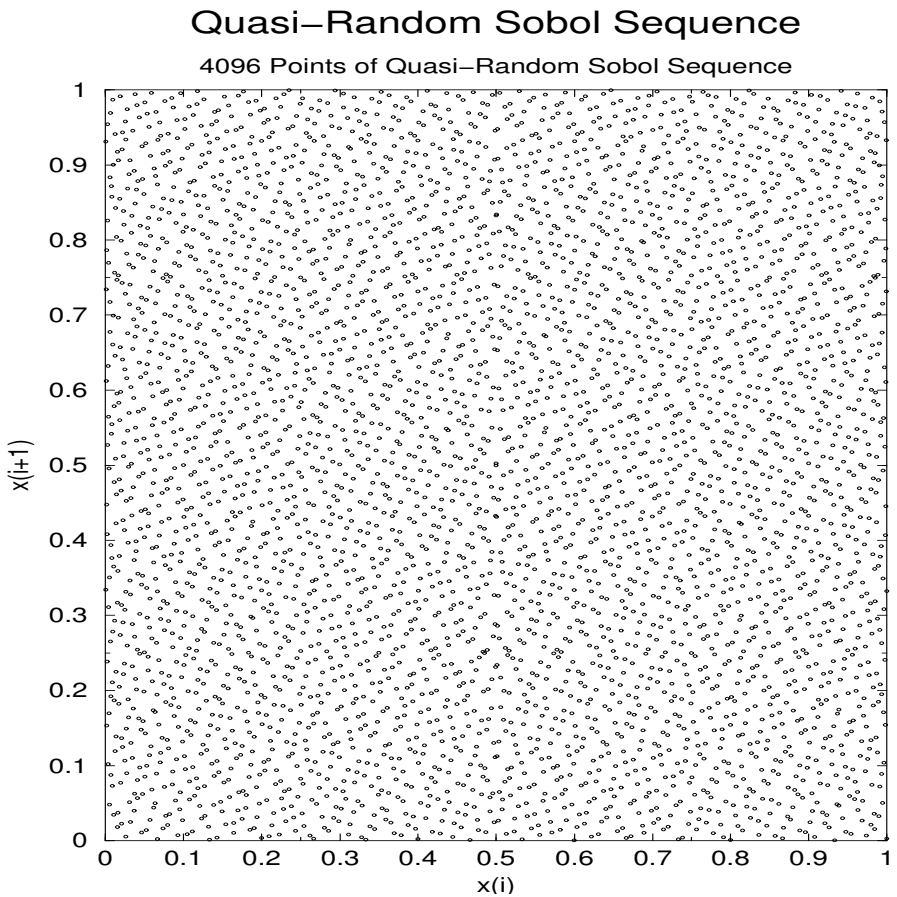

Reason why uniform grid is inefficient: Projection of $N = n^d$ points in d dimensions onto a line maps n^{d-1} points onto a single point.

Reason why quasi-MC is more efficient than pseudo-MC in intermediate # of dimensions (e.g. finance applications): Quasi-MC avoids clusters and voids.

Negatives for quasi-MC: Difficult to combine with importance sampling (needed for spiky functions), cannot choose # of MC points freely.

Sampling of arbitrary probability density functions

Infinite-variance estimators can be replaced by finite-variance estimators by sampling the MC points from an appropriate probability density functions.

Techniques for sampling arbitrary probability density functions employ standard random numbers generators that sample a uniform distribution in $[0, 1]$. We study 3 techniques for sampling nonuniform distributions:

1. transformation method
2. rejection method
3. Metropolis-Hastings method

but first we say a few words about random number generators.

Random Number Generators

Conventional random number generators generate random numbers uniformly distributed on [0,1).

Of course no computer generated sequence of random numbers is truly random. For one, the random numbers must repeat after a finite (though hopefully very large) period. Also, if N bits are used to represent the random numbers, then the number of different numbers generated can by no larger than 2^N.

Note however, that the period can be (and typically is for the better generators) much larger than 2^N.

Many different algorithms exist for generating random numbers, e.g., linear congruential generators (with or without an additive constant), linear feedback shift register, lagged Fibonacci generator, XORshift algorithm etc. They are typically subjected to a battery of statistical tests, e.g., the Diehard tests of Marsaglia. Of course no random number generator can pass all the tests that one can invent, but hopefully the random number generator used does not have correlations that could significantly impact the system being studied.

Random Number Generators

For many MC calculations it is the short-ranged correlations that matter most, but one has to think for each application what is important. For example, if one were studying an Ising model with a power of two number of spins, it would be problematic to have random number generator that generated numbers with bits that repeat at an interval of 2^N.

In the old days, there were quite a few calculations that produced inaccurate results due to bad random number generators. For example, the standard generators that came with UNIX and with C were badly flawed. In the 1980s a special purpose computer was built at Santa Barbara to study the 3-D Ising model. However, at first it failed to reproduce the known exact results for the 2-D Ising model and that failure was traced back to a faulty random number generator. Fortunately, these days the standard random number generators are much more reliable.

Sampling random variables from nonuniform probability density functions

We say x is sampled from $f(x)$ if for any a and b in the domain,

$$\mathrm{Prob}[a \le x \le b] \quad = \quad \int_a^b dx' \; f(x')$$

1) Transformation method (For many simple functions)
2) Rejection method (For more complicated functions)
3) Metropolis-Hastings method (For any function)

1) Transformation method: Perform a transformation $x(\xi)$ on a uniform deviate ξ, to get x sampled from desired probability density $f(x)$.

$$|\mathrm{Prob}(\xi) d\xi| \quad = \quad |\mathrm{Prob}(x) dx| \quad \text{conservation of probability}$$

If we have sampled ξ from a uniform density ($\mathrm{Prob}(\xi) = 1$) and we wish x to be sampled from the desired density, $f(x)$, then setting $\mathrm{Prob}(x) = f(x)$,

$$\left| \frac{d\xi}{dx} \right| \quad = \quad f(x)$$

Solve for $\xi(x)$ and invert to get $x(\xi)$, i.e., invert the cumulative distrib.

Examples of Transformation Method

Example 1: $f(x) = ae^{-ax}$, $x \in [0, \infty)$

$$\left|\frac{d\xi}{dx}\right| = ae^{-ax}, \quad \text{or,} \quad \xi = e^{-ax}, \quad \text{i.e.,} \quad \boxed{x = \frac{-\ln(\xi)}{a}}$$

Example 2: $f(x) = \frac{x^{-1/2}}{2}$, $x \in [0, 1]$

$$\left|\frac{d\xi}{dx}\right| = \frac{x^{-1/2}}{2}, \quad \text{or} \quad \xi = x^{1/2}, \quad \text{i.e.,} \quad \boxed{x = \xi^2}$$

Note that in this case we are sampling a probability density that is infinite at 0, but that is OK!

Example 3: $f(x) = xe^{-x^2/2}$, $x \in [0, \infty)$

$$\left|\frac{d\xi}{dx}\right| = xe^{-x^2/2}, \quad \text{or,} \quad \xi = e^{-x^2/2}, \quad \text{i.e.,} \quad \boxed{x = \sqrt{-2\ln(\xi)}}$$

Examples of Transformation Method

Example 4a: $f(x) = \frac{e^{-x^2/2}}{\sqrt{2\pi}}$, $x \in (-\infty, \infty)$ (using Box-Müller method)

$$\frac{1}{2\pi} e^{-(\frac{x_1^2}{2}+\frac{x_2^2}{2})}\, dx_1\, dx_2 \;=\; \left(r\, e^{-\frac{r^2}{2}}\, dr\right)\left(\frac{d\phi}{2\pi}\right)$$

$$r = \sqrt{-2\log(\xi_1)}, \qquad \phi = 2\pi\xi_2$$

$$\boxed{x_1 = \sqrt{-2\log(\xi_1)}\cos(2\pi\xi_2)}, \qquad \boxed{x_2 = \sqrt{-2\log(\xi_1)}\sin(2\pi\xi_2)}$$

(x_1 and x_2 are uncorrelated)

Example 4b: $f(x) \approx \frac{e^{-x^2/2}}{\sqrt{2\pi}}$, $x \in (-\infty, \infty)$ (using central-limit theorem)

Since σ^2 for uniform distribution about 0 is $\displaystyle\int_{-1/2}^{1/2} dx\, x^2 \;=\; \frac{1}{12}$

$$x \;=\; \lim_{N\to\infty} \sqrt{\frac{12}{N}}\left(\sum_{i=1}^{N} \xi_i - \frac{N}{2}\right) \;\approx\; \boxed{\sum_{i=1}^{12} \xi_i - 6}$$

(avoids log, sqrt, cos, sin, but, misses tiny tails beyond ± 6)

Rejection Method

We wish to sample $f(x)$.
Find a function $g(x)$ that can be sampled by another method (say transformation) and that preferably mimics the behaviour of $f(x)$.
Let C be an upper bound to the maximum value of $f(x)/g(x)$.
Let $C \geq \max(f(x)/g(x))$.
Then $f(x)$ is sampled by sampling $g(x)$ and keep the sampled points with probability

$$P = \frac{f(x)}{Cg(x)}$$

The efficiency of the method is the fraction of the sampled points that are kept.

$$\begin{aligned} Eff &= \int dx \frac{f(x)}{Cg(x)} g(x) \\ &= \frac{1}{C} \end{aligned}$$

Drawback: It is often hard to know C and a "safe" upperbound choice for C may lead to low efficiency. An alternative is to associate weights with the sampled points.

Importance Sampling for computing integrals efficiently

Now that we know how to sample simple probability density functions, we study how to use *importance sampling* to compute integrals more efficiently.

Example of Importance Sampling to Calculate Integrals More Efficiently

Suppose we wish to compute

$$\int_0^1 dx\, f(x) \;=\; \int_0^1 dx\, \frac{1}{x^p + x} \qquad \left(= \frac{log(2)}{1-p}, \text{ but pretend not known}\right)$$

Note that

$$\int_0^1 dx (f(x))^2 \;=\; \infty, \quad (\text{for } p \geq 0.5)$$

so if we estimate the integral by sampling points uniformly in $[0, 1]$ then this would be an infinite variance estimator and the error of the estimate will go down more slowly than $N^{-1/2}$. However, we can instead sample points from the density

$$g(x) \;=\; \frac{1-p}{x^p}$$

Now the variance of $f(x)/g(x)$ is finite and the error decreases as $N^{-1/2}$, and, with a small prefactor. (Still would not use this in 1D.)

Homework Problem 1

Compute

$$I = \int_0^1 dx\, f(x) = \int_0^1 dx\, \frac{1}{x^p + x} \quad \left(= \frac{log(2)}{1-p}, \text{ but pretend not known}\right) \approx \frac{1}{N_{\rm MC}} \sum_{k=1}^{N_{\rm MC}} \frac{1}{\xi_k^p + \xi_k}$$

with/without importance sampling, using for the importance sampling function

$$g(x) = \frac{(1-p)}{x^p}$$

To sample $g(x)$: $\left|\frac{d\xi}{dx}\right| = (1-p)x^{-p}$, i.e., $\xi = x^{1-p}$, i.e., $\boxed{x = \xi^{\frac{1}{1-p}}}$

$$\int_0^1 dx\, f(x) = \int_0^1 dx\, g(x)\frac{f(x)}{g(x)} = \int_0^1 dx\, \frac{1-p}{x^p}\frac{1}{(1-p)(1+x^{1-p})}$$

$$\approx \frac{1}{N_{\rm MC}(1-p)} \sum_{k=1}^{N_{\rm MC}} \frac{1}{(1+x_k^{1-p})} = \frac{1}{N_{\rm MC}(1-p)} \sum_{k=1}^{N_{\rm MC}} \frac{1}{(1+\xi_k)}$$

Do this for $p = 0.25, 0.5, 0.75, 0.95$ and $N_{\rm MC} = 10^3, 10^4, 10^5, 10^6, 10^7, 10^8, 10^9$.
Plot 2 graphs, each having 8 curves (4 values of p, and, with/without importance sampling):

1. Log of estimated 1-standard deviation statistical error versus $\log(N_{\rm MC})$.
2. Actual error in I, with estimated 1-std. dev. statistical error as an error bar versus $\log(N_{\rm MC})$.

Homework Solution 1a

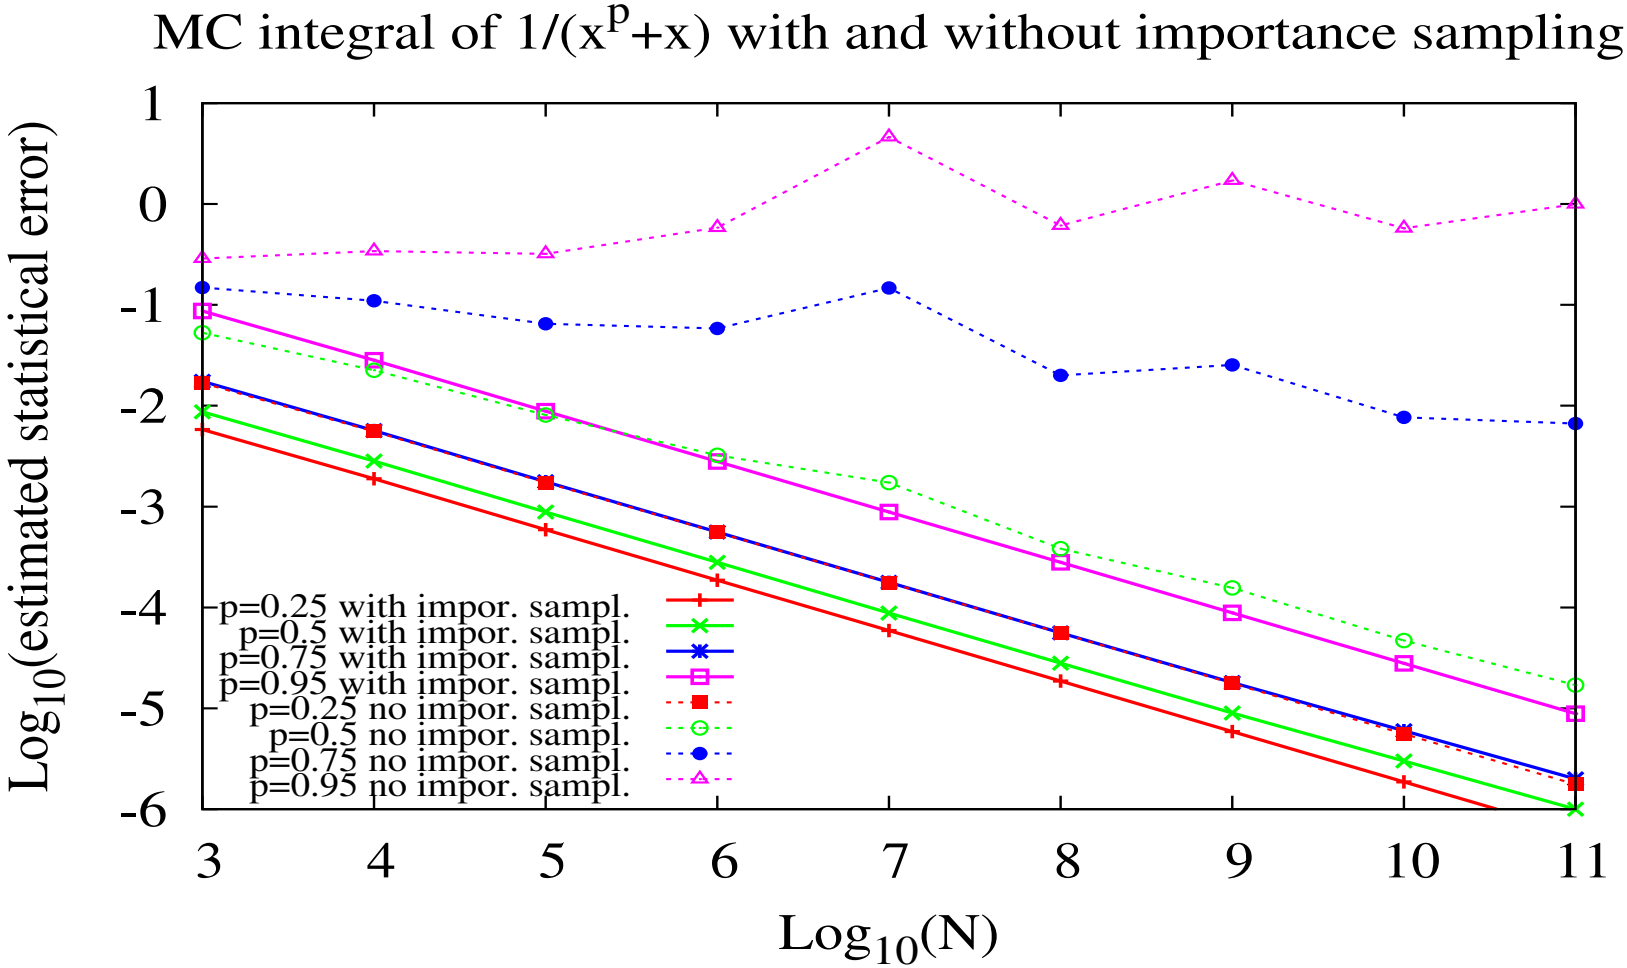

Statistical errors $\sim N_{\rm MC}{}^{-1/2}$ for all p with importance sampling but only for $p = 0.25$ without importance sampling. For $p = 1$ even the integral is infinite. For $p = 0.95$ no sign of convergence. Theorem about asymptotic convergence of little practical utility.

Homework Solution 1b

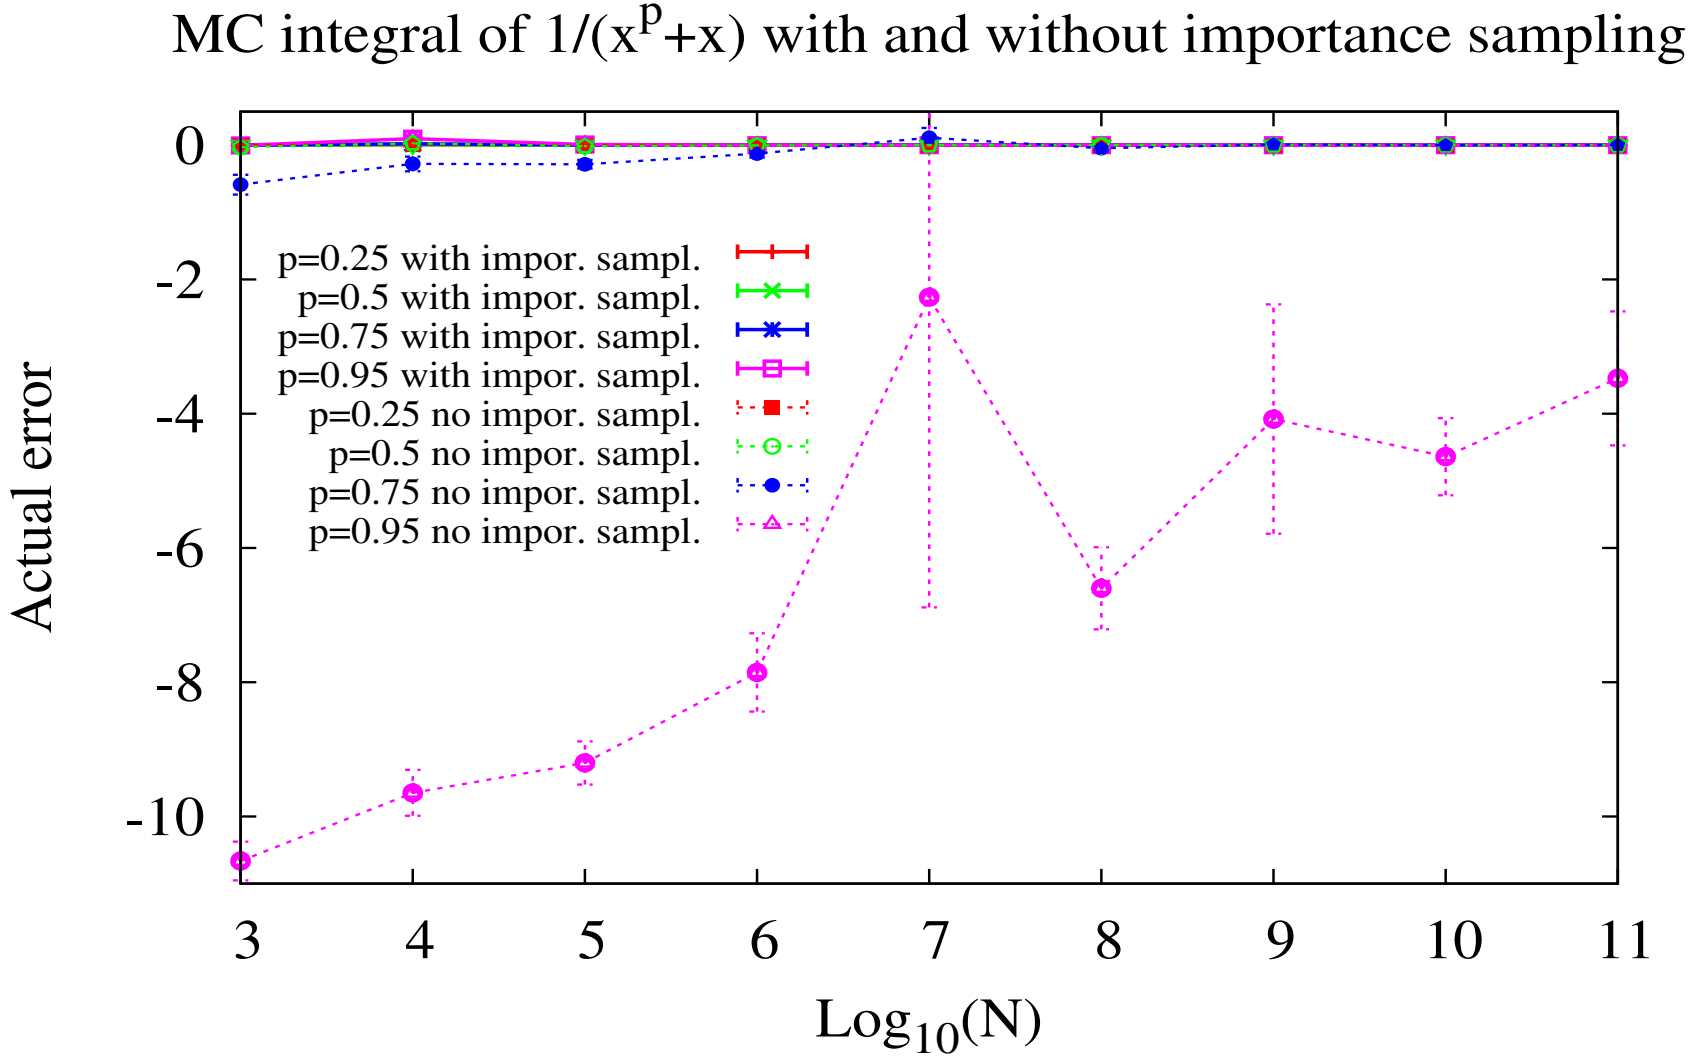

For $p = 0.95$ all of the errors are negative. Occasional large positive errors will bring mean to correct value. Usual error estimates are meaningless. The points with the larger errors tend to have the smaller estimated errors! Weighting estimates by inverse variances is bad!!

Unbiased Estimators

In this section, denote population means by brackets and sample means by bars.
Let $f(x)$ be a random variable with probability density $\rho(x)$.

$$\text{Population mean:} \quad \mu = \int dx\, f(x)\, \rho(x) \equiv \langle f \rangle_\rho$$

$$\text{Population variance:} \quad \sigma^2 = \int dx\, (f(x) - \langle f \rangle_\rho)^2\, \rho(x) = \langle f^2 \rangle_\rho - \langle f \rangle_\rho^2$$

Estimator is unbiased if averaging over an infinite number of samples of size N gives the same result as that from a single infinite sample. Can estimate mean and variance from independent finite samples of size N, but the "obvious" estimators often have $\mathcal{O}(1/N)$ errors, so we provide here estimators that are correct at least to $\mathcal{O}(1/N)$.

$$\text{Unbiased estimator for} \quad \langle f \rangle_\rho \quad : \quad \frac{1}{N} \sum_{i=1}^{N} f(x_i) \equiv \overline{f_\rho}$$

$$\text{Unbiased estimators for} \quad \sigma^2 \quad : \quad \frac{1}{N} \sum_{i=1}^{N} (f(x_i) - \langle f \rangle_\rho)^2 \quad \text{(but } \langle f \rangle_\rho \text{ not known)}$$

$$\text{and} \qquad \frac{1}{N-1} \sum_{i=1}^{N} (f(x_i) - \overline{f_\rho})^2 = \frac{N}{N-1} \left(\overline{f_\rho^2} - \overline{f_\rho}^2 \right)$$

Estimating Unbiased Variance from Uncorrelated Samples

Let $\langle f(x) \rangle$ denote the population mean and $\overline{f}(x)$ denote the sample mean. Then $\overline{f^2} - (\overline{f})^2 =$

$$\left\langle \frac{\sum_i f^2(x_i)}{N} - \left[\frac{\sum_i f(x_i)}{N} \right]^2 \right\rangle = \langle f^2 \rangle - \left\langle \frac{\sum_i f^2(x_i) + \sum_{i,j \neq i} \sum_j f(x_i) f(x_j)}{N^2} \right\rangle$$

Since $f(x_i)$ and $f(x_j)$ are independent

$$RHS = \left(1 - \frac{1}{N}\right) \langle f^2 \rangle - \frac{N(N-1)}{N^2} \langle f \rangle^2 = \frac{N-1}{N} (\langle f^2 \rangle - \langle f \rangle^2) = \frac{N-1}{N} \sigma^2$$

So, the unbiased estimate for σ^2 is

$$\boxed{\sigma^2 \approx \frac{N}{N-1} \left(\overline{f^2} - (\overline{f})^2 \right)}$$

Loss of one degree of freedom because sample variance is computed relative to sample mean rather than the true mean.

Examples of Unbiased and Biased Estimators

$$E_T = \frac{\int d\mathbf{R}\, \psi_{\rm T}(\mathbf{R})\, \mathcal{H}\, \psi_{\rm T}(\mathbf{R})}{\int d\mathbf{R}\, \psi_{\rm T}^2(\mathbf{R})} = \int d\mathbf{R}\, \frac{\psi_{\rm T}^2(\mathbf{R})}{\int d\mathbf{R}\, \psi_{\rm T}^2(\mathbf{R})}\, \frac{\mathcal{H}\psi_{\rm T}(\mathbf{R})}{\psi_{\rm T}(\mathbf{R})}$$

$$= \frac{1}{N}\sum_{i=1}^{N} \frac{\mathcal{H}\Psi_{\rm T}(\mathbf{R}_i)}{\Psi_{\rm T}(\mathbf{R}_i)} = \frac{1}{N}\sum_{i=1}^{N} E_L(\mathbf{R}_i) \qquad \text{unbiased}$$

$$E_T = \frac{\int d\mathbf{R}\, \psi_{\rm T}(\mathbf{R})\, \mathcal{H}\, \psi_{\rm T}(\mathbf{R})}{\int d\mathbf{R}\, \psi_{\rm T}^2(\mathbf{R})} = \frac{\int d\mathbf{R}\, \frac{|\psi_{\rm T}(\mathbf{R})|}{\int d\mathbf{R}\, |\psi_{\rm T}(\mathbf{R})|}\, \mathrm{sgn}(\psi_{\rm T}(\mathbf{R}))\, \mathcal{H}\psi_{\rm T}(\mathbf{R})}{\int d\mathbf{R}\, \frac{|\psi_{\rm T}(\mathbf{R})|}{\int d\mathbf{R}\, |\psi_{\rm T}(\mathbf{R})|}\, |\psi_{\rm T}(\mathbf{R})|}$$

$$= \frac{\sum_{i=1}^{N} \mathrm{sgn}(\psi_{\rm T}(\mathbf{R}))\, \mathcal{H}\Psi_{\rm T}(\mathbf{R}_i)}{\sum_{i=1}^{N} |\psi_{\rm T}(\mathbf{R})|} \qquad \mathcal{O}\left(\tfrac{1}{N}\right) \text{ bias}$$

Can do better by calculating covariances.

Unbiased Estimators to $\mathcal{O}(1/N)$ of functions of expectation values and their variance

$\langle x \rangle \equiv$ population averages of x, i.e., true expectation value
$\bar{x} \equiv$ average of x over sample of size N

Let F be a function of expectation values, $\{\langle f_i \rangle\}$.
$F\left(\{\bar{f}_i\}\right)$ is unbiased estimator for $F(\{\langle f_i \rangle\})$ iff F is linear function of $\{\langle f_i \rangle\}$.

In general

$$F(\{\langle f_i \rangle\}) = F\left(\{\bar{f}_i\}\right) - \frac{1}{2}\sum_{i,j} \frac{\partial^2 F}{\partial f_i \partial f_j} \frac{\mathrm{cov}(f_i, f_j)}{N} + \mathcal{O}\left(\frac{1}{N^2}\right)$$

$$\mathrm{var}\big(F(\{\langle f_i \rangle\})\big) = \sum_{i,j} \frac{\partial F}{\partial f_i} \frac{\partial F}{\partial f_j} \mathrm{cov}(f_i, f_j) + \mathcal{O}\left(\frac{1}{N}\right)$$

Unbiased Estimators to $\mathcal{O}(1/N)$ or better (cont)

$$\text{Estim. of mean} \quad \langle f \rangle_\rho = \overline{f_\rho}$$

$$\text{Estim. of variance} \quad \langle f^2 \rangle_\rho - \langle f \rangle_\rho^2 = \frac{N}{N-1}\left(\overline{f_\rho^2} - \overline{f_\rho}^2\right)$$

$$\text{Estim. of error of sample mean} = \sqrt{\frac{1}{N-1}\left(\overline{f_\rho^2} - \overline{f_\rho}^2\right)}$$

$$\text{Estim. of covar.} \quad \mathrm{cov}(f,g) \equiv \langle fg \rangle_\rho - \langle f \rangle_\rho \langle g \rangle_\rho = \frac{N}{N-1}\left(\overline{fg_\rho} - \overline{f_\rho}\,\overline{g_\rho}\right)$$

$$\text{Estim. of product of expec. values} \quad \langle f \rangle_\rho \langle g \rangle_\rho = \overline{f_\rho}\,\overline{g_\rho} - \frac{1}{N}\mathrm{cov}(f,g)$$

$$\text{Estim. of ratio of expec. values} \quad \frac{\langle f \rangle_\rho}{\langle g \rangle_\rho} \approx \frac{\overline{f_\rho}}{\overline{g_\rho}}\left[1 - \frac{1}{N}\left(\frac{\sigma_g^2}{\langle g \rangle_\rho^2} - \frac{\mathrm{cov}(f,g)}{\langle f \rangle_\rho \langle g \rangle_\rho}\right)\right]$$

$$\mathrm{Var}\left(\overline{f_\rho}\ \overline{g_\rho}\right) = \frac{1}{N}\langle f \rangle_\rho^2 \langle g \rangle_\rho^2 \left(\frac{\sigma_f^2}{\langle f \rangle_\rho^2} + \frac{\sigma_g^2}{\langle g \rangle_\rho^2} + 2\frac{\mathrm{cov}(f,g)}{\langle f \rangle_\rho \langle g \rangle_\rho}\right)$$

$$\mathrm{Var}\left(\frac{\overline{f_\rho}}{\overline{g_\rho}}\right) = \frac{1}{N}\frac{\langle f \rangle_\rho^2}{\langle g \rangle_\rho^2}\left[\frac{\sigma_f^2}{\langle f \rangle_\rho^2} + \frac{\sigma_g^2}{\langle g \rangle_\rho^2} - 2\frac{\mathrm{cov}(f,g)}{\langle f \rangle_\rho \langle g \rangle_\rho}\right].$$

Note that the product, $\overline{f_\rho}\,\overline{g_\rho}$ is unbiased if $\mathrm{cov}(f,g) = 0$, but the ratio $\frac{\overline{f_\rho}}{\overline{g_\rho}}$ has $\mathcal{O}(1/N)$ bias even if $\mathrm{cov}(f,g) = 0$. The ratio has no bias (and no fluctuations) when f and g are perfectly correlated. In practice replace population means by sample means on RHS.

Unbiased Estimators of autocorrelated variables

Independent samples:

Estim. for error of sample mean $$\overline{\Delta_f} = \sqrt{\frac{1}{N-1}\left(\overline{f_\rho^2} - \overline{f_\rho}^2\right)}$$

Autocorrelated samples (e.g. from Metropolis):

Estim. for error of sample mean $$\overline{\Delta_f} = \sqrt{\frac{1}{N_{\rm eff}-1}\left(\overline{f_\rho^2} - \overline{f_\rho}^2\right)}$$

where

$$N_{\rm eff} = \frac{N}{(1+2\tau_f)} \equiv \frac{N}{T_{\rm corr}}$$

$$\tau_f = \frac{\sum_{t=1}^{\infty}\left[\langle f_1 f_{1+t}\rangle_\rho - \langle f\rangle_\rho^2\right]}{\sigma_f^2}$$

If samples are indep., $\langle f_1 f_{1+t}\rangle_\rho = \langle f\rangle_\rho^2$ and integrated autocorrelation time $\tau_f = 0$. Since the relevant quantity for MC calculations is $(1+2\tau_f) \equiv T_{\rm corr}$ we will refer to it as the autocorrelation time of f, though this is not standard usage.

Lecture 2

Variational Monte Carlo and Metropolis-Hastings Algorithm

Cyrus J. Umrigar

Variational Monte Carlo in Real Space

W. L. McMillan, Phys. Rev. **138**, A442 (1965)

Monte Carlo is used to perform the many-dimensional integrals needed to calculate quantum mechanical expectation values. e.g.

$$
\begin{aligned}
E_T &= \frac{\int d\mathbf{R}\, \Psi_{\rm T}^*(\mathbf{R})\, \mathcal{H}\, \psi_{\rm T}(\mathbf{R})}{\int d\mathbf{R}\, \psi_{\rm T}^2(\mathbf{R})} \\
&= \int d\mathbf{R}\, \frac{\psi_{\rm T}^2(\mathbf{R})}{\int d\mathbf{R}\, \psi_{\rm T}^2(\mathbf{R})}\, \frac{\mathcal{H}\psi_{\rm T}(\mathbf{R})}{\psi_{\rm T}(\mathbf{R})} \\
&= \frac{1}{N}\sum_i \frac{\mathcal{H}\Psi_{\rm T}(\mathbf{R}_i)}{\Psi_{\rm T}(\mathbf{R}_i)} = \frac{1}{N}\sum_i E_L(\mathbf{R}_i)
\end{aligned}
$$

Energy is obtained as an arithmetic sum of the *local energies* $E_L(\mathbf{R}_i)$ evaluated for configurations sampled from $\psi_{\rm T}^2(\mathbf{R})$ using a generalization of the Metropolis method. If $\psi_{\rm T}$ is an eigenfunction the $E_L(\mathbf{R}_i)$ do not fluctuate. Accuracy of VMC depends crucially on the quality of $\psi_{\rm T}(\mathbf{R})$. Diffusion MC does better by projecting onto ground state.

Three ingredients for accurate Variational Monte Carlo

1. A method for sampling an arbitrary wave function Metropolis-Hastings.
2. A functional form for the wave function that is capable of describing the correct physics/chemistry.
3. An efficient method for optimizing the parameters in the wave functions.

Metropolis-Hastings Monte Carlo

Metropolis, Rosenbluth[2], Teller[2], JCP, **21** 1087 (1953)

W.K. Hastings, Biometrika, **57** (1970)

Metropolis method originally used to sample the Boltzmann distribution. This is still one of its more common uses.

General method for sampling **any known** discrete or continuous density. (Other quantum Monte Carlo methods, e.g., diffusion MC, enable one to sample densities that are not explicitly known but are the eigenstates of known matrices or integral kernels.)

Metropolis-Hastings has serial correlations. Hence, direct sampling methods preferable, but rarely possible for complicated densities in many dimensions.

Metropolis-Hastings Monte Carlo (cont)

A *Markov chain* is specified by two ingredients:
1) an initial state
2) a transition matrix $M(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i})$ (probability of transition $\mathbf{R}_\mathrm{i} \to \mathbf{R}_\mathrm{f}$.)

$$M(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i}) \geq 0, \quad \sum_{\mathbf{R}_\mathrm{f}} M(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i}) = 1. \qquad \textit{Column-stochastic matrix}$$

To sample $\rho(\mathbf{R})$, start from an arbitrary $\mathbf{R}_\mathrm{i}$ and evolve the system by repeated application of M that satisfies the *stationarity condition* (flux out of state $\mathbf{R}_\mathrm{i}$ equals flux into $\mathbf{R}_\mathrm{i}$):

$$\sum_{\mathbf{R}_\mathrm{f}} M(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i})\, \rho(\mathbf{R}_\mathrm{i}) = \sum_{\mathbf{R}_\mathrm{f}} M(\mathbf{R}_\mathrm{i}|\mathbf{R}_\mathrm{f})\, \rho(\mathbf{R}_\mathrm{f}) = \rho(\mathbf{R}_\mathrm{i}) \quad \forall\, \mathbf{R}_\mathrm{i}$$

i.e., $\rho(\mathbf{R})$ is a right eigenvector of M with eigenvalue 1.
Stationarity $\Rightarrow$ if we start with ρ, will continue to sample ρ.
Want more than that: *any* initial density should evolve to ρ.

$$\lim_{n\to\infty} M^n(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i})\, \delta(\mathbf{R}_\mathrm{i}) = \rho(\mathbf{R}_\mathrm{f}), \quad \forall\, \mathbf{R}_\mathrm{i}.$$

i.e., ρ should be the *dominant* right eigenvector.

Metropolis-Hastings Monte Carlo (cont)

Want that any initial density should evolve to ρ.

$$\lim_{n\to\infty} M^n(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i})\delta(\mathbf{R}_\mathrm{i}) = \rho(\mathbf{R}_\mathrm{f}), \quad \forall\, \mathbf{R}_\mathrm{i}.$$

ρ should be the *dominant* right eigenvector. Additional conditions needed to guarantee this.

A nonnegative matrix M is said to be *primitive* if $\exists$ n such that M^n has all elements positive.

(Special case of) Perron-Frobenius Theorem: A column-stochastic primitive matrix has a unique dominant eigenvalue of 1, with a positive right eigenvector and a left eigenvector with all components equal to 1 (by definition of column-stochastic matrix).

In a finite space, necessary and sufficient conditions are that Markov matrix M is primitive. (Same ideas in continuous space (matrix $\to$ integral kernel) but statements and proofs trickier.)

In practice, length of Monte Carlo should be long enough that there be a significant probability of the system making several transitions between the neighborhoods of any pair of representative states that make a significant contribution to the average. This ensures that states are visited with the correct probability with only small statistical fluctuations. For example in a double-well system many transitions between the 2 wells should occur, but we can choose our proposal matrix to achieve this even if barrier between wells is high.

Metropolis-Hastings Monte Carlo (cont)

Construction of M

Need a prescription to construct M, such that ρ is its stationary state. Impose *detailed balance* condition

$$M(\mathbf{R}_f|\mathbf{R}_i)\ \rho(\mathbf{R}_i) = M(\mathbf{R}_i|\mathbf{R}_f)\ \rho(\mathbf{R}_f)$$

Detailed balance more stringent than stationarity condition.
Detailed balance is not necessary but provides way to construct M.
Write elements of M as product of elements of a proposal matrix T and an acceptance Matrix A,

$$M(\mathbf{R}_f|\mathbf{R}_i) = A(\mathbf{R}_f|\mathbf{R}_i)\ T(\mathbf{R}_f|\mathbf{R}_i)$$

$M(\mathbf{R}_f|\mathbf{R}_i)$ and $T(\mathbf{R}_f|\mathbf{R}_i)$ are stochastic matrices, but $A(\mathbf{R}_f|\mathbf{R}_i)$ is not.
Detailed balance is now:

$$A(\mathbf{R}_f|\mathbf{R}_i)\ T(\mathbf{R}_f|\mathbf{R}_i)\ \rho(\mathbf{R}_i) = A(\mathbf{R}_i|\mathbf{R}_f)\ T(\mathbf{R}_i|\mathbf{R}_f)\ \rho(\mathbf{R}_f)$$

$$\text{or} \qquad \frac{A(\mathbf{R}_f|\mathbf{R}_i)}{A(\mathbf{R}_i|\mathbf{R}_f)} = \frac{T(\mathbf{R}_i|\mathbf{R}_f)\ \rho(\mathbf{R}_f)}{T(\mathbf{R}_f|\mathbf{R}_i)\ \rho(\mathbf{R}_i)}\ .$$

Metropolis-Hastings Monte Carlo (cont)

Choice of Acceptance Matrix A

$$\frac{A(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i})}{A(\mathbf{R}_\mathrm{i}|\mathbf{R}_\mathrm{f})} = \frac{T(\mathbf{R}_\mathrm{i}|\mathbf{R}_\mathrm{f})\ \rho(\mathbf{R}_\mathrm{f})}{T(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i})\ \rho(\mathbf{R}_\mathrm{i})}\ .$$

Infinity of choices for A. Any function

$$F\left(\frac{T(\mathbf{R}_\mathrm{i}|\mathbf{R}_\mathrm{f})\ \rho(\mathbf{R}_\mathrm{f})}{T(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i})\ \rho(\mathbf{R}_\mathrm{i})}\right) = A(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i})$$

for which $F(x)/F(1/x) = x$ and $0 \leq F(x) \leq 1$ will do.
Choice of Metropolis *et al.* $F(x) = \min\{1, x\}$, maximizes the acceptance:

$$A(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i}) = \min\left\{1,\ \frac{T(\mathbf{R}_\mathrm{i}|\mathbf{R}_\mathrm{f})\ \rho(\mathbf{R}_\mathrm{f})}{T(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i})\ \rho(\mathbf{R}_\mathrm{i})}\right\}.$$

Other less good choices for $A(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i})$ have been made, *e.g.* $F(x) = \frac{x}{1+x}$

$$A(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i}) = \frac{T(\mathbf{R}_\mathrm{i}|\mathbf{R}_\mathrm{f})\ \rho(\mathbf{R}_\mathrm{f})}{T(\mathbf{R}_\mathrm{i}|\mathbf{R}_\mathrm{f})\ \rho(\mathbf{R}_\mathrm{f}) + T(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i})\ \rho(\mathbf{R}_\mathrm{i})}.$$

Metropolis: $T(\mathbf{R}_\mathrm{i}|\mathbf{R}_\mathrm{f}) = T(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i})$, Hastings: $T(\mathbf{R}_\mathrm{i}|\mathbf{R}_\mathrm{f}) \neq T(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i})$

Metropolis-Hastings Monte Carlo (cont)

Choice of Proposal Matrix T

So, the optimal choice for the acceptance matrix $A(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i})$ is simple and known.

However, there is considerable scope for using one's ingenuity to come up with good proposal matrices, $T(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i})$, that allow one to make large moves with large acceptances, in order to make the autocorrelation time small.

Choice of Proposal Matrix T in Metropolis-Hastings (cont)

CJU, PRL **71**, 408 (1993)

$$A(\mathbf{R}_f|\mathbf{R}_i) = \min\left\{1, \frac{T(\mathbf{R}_i|\mathbf{R}_f)\ \rho(\mathbf{R}_f)}{T(\mathbf{R}_f|\mathbf{R}_i)\ \rho(\mathbf{R}_i)}\right\}$$

Use freedom in T to make $\quad \dfrac{T(\mathbf{R}_i|\mathbf{R}_f)\ \rho(\mathbf{R}_f)}{T(\mathbf{R}_f|\mathbf{R}_i)\ \rho(\mathbf{R}_i)} \approx 1.$

$T(\mathbf{R}_f|\mathbf{R}_i) \propto \rho(\mathbf{R}_f)$ optimal if $T(\mathbf{R}_f|\mathbf{R}_i)$ can be sampled over all space – usually not the case. And if it is, then one would not use Metropolis-Hastings in the first place.

Otherwise, let $\quad T(\mathbf{R}_f|\mathbf{R}_i) = \dfrac{S(\mathbf{R}_f|\mathbf{R}_i)}{\int d\mathbf{R}_f\, S(\mathbf{R}_f|\mathbf{R}_i)} \approx \dfrac{S(\mathbf{R}_f|\mathbf{R}_i)}{S(\mathbf{R}_i|\mathbf{R}_i)\Omega(\mathbf{R}_i)}$

$S(\mathbf{R}_f|\mathbf{R}_i)$ is non-zero only in domain $D(\mathbf{R}_i)$ of volume $\Omega(\mathbf{R}_i)$ around $\mathbf{R}_i$).

$$\frac{A(\mathbf{R}_f,\mathbf{R}_i)}{A(\mathbf{R}_i,\mathbf{R}_f)} = \frac{T(\mathbf{R}_i|\mathbf{R}_f)}{T(\mathbf{R}_f|\mathbf{R}_i)}\frac{\rho(\mathbf{R}_f)}{\rho(\mathbf{R}_i)} \approx \frac{\Omega(\mathbf{R}_i)}{\Omega(\mathbf{R}_f)}\frac{S(\mathbf{R}_i|\mathbf{R}_i)}{S(\mathbf{R}_f|\mathbf{R}_f)}\frac{S(\mathbf{R}_i|\mathbf{R}_f)}{S(\mathbf{R}_f|\mathbf{R}_i)}\frac{\rho(\mathbf{R}_f)}{\rho(\mathbf{R}_i)}$$

from which it is apparent that the choice

$$S(\mathbf{R}_f|\mathbf{R}_i) \overset{\propto}{\approx} \sqrt{\rho(\mathbf{R}_f)/\Omega(\mathbf{R}_f)} \quad \text{yields} \quad A(\mathbf{R}_f,\mathbf{R}_i)/A(\mathbf{R}_i,\mathbf{R}_f) \approx 1.$$

Choice of Proposal Matrix T in Metropolis-Hastings (cont)

To be more precise, if the log-derivatives of $T(\mathbf{R}_f|\mathbf{R}_i)$ equal those of $\sqrt{\rho(\mathbf{R}_f)/\Omega(\mathbf{R}_f)}$ at $\mathbf{R}_f = \mathbf{R}_i$, the average acceptance goes as $1 - \mathcal{O}(\Delta^m)$, where Δ is the linear dimension of $D(\mathbf{R}_i)$.

In general, $m=2$, but if $D(\mathbf{R}_i)$ is inversion symmetric with $\mathbf{R}_i$ at its center, then $m=3$.

Considerable improvement compared to using a symmetric $S(\mathbf{R}_f|\mathbf{R}_i)$ or choosing $S(\mathbf{R}_f|\mathbf{R}_i) \overset{\propto}{\sim} \rho(\mathbf{R}_f)$ for either of which m=1.

Another possible choice, motivated by (DMC) is

$$T(\mathbf{R}_f|\mathbf{R}_i) = \frac{1}{(2\pi\tau)^{3/2}} \exp\left[\frac{-(\mathbf{R}_f - \mathbf{R}_i - \mathbf{V}(\mathbf{R}_i)\tau)^2}{2\tau}\right], \quad \mathbf{V}(\mathbf{R}_i) = \frac{\nabla\Psi(\mathbf{R}_i)}{\Psi(\mathbf{R}_i)}$$

Advantage: allows Metropolis Monte Carlo and diffusion Monte Carlo programs to share almost all the code.
$m = 2$ for this choice of T, so such an algorithm is more efficient than one with a symmetric $S(\mathbf{R}_f|\mathbf{R}_i)$ or one for which $S(\mathbf{R}_f|\mathbf{R}_i) \overset{\propto}{\sim} \rho(\mathbf{R}_f)$, but less efficient than one for which $S(\mathbf{R}_f|\mathbf{R}_i) \overset{\propto}{\sim} \sqrt{\rho(\mathbf{R}_f)/\Omega(\mathbf{R}_f)}$.

These arguments are rigorous only in the small-step limit and are applicable only to functions with sufficiently many derivatives within $D(\mathbf{R}_i)$. In practice these ideas yield large reduction in the autocorrelation time provided that we employ a coordinate system such that ρ has continuous derivatives within $D(\mathbf{R}_i)$.

Choice of Proposal Matrix T in Metropolis-Hastings (cont)

Another possible choice, motivated by (DMC) is

$$T(\mathbf{R}_f|\mathbf{R}_i) = \frac{1}{(2\pi\tau)^{3/2}} \exp\left[\frac{-(\mathbf{R}_f - \mathbf{R}_i - \mathbf{V}(\mathbf{R}_i)\tau)^2}{2\tau}\right], \quad \mathbf{V}(\mathbf{R}_i) = \frac{\nabla\Psi(\mathbf{R}_i)}{\Psi(\mathbf{R}_i)}$$

Advantage: allows Metropolis Monte Carlo and diffusion Monte Carlo programs to share almost all the code.

Metropolis-Hastings Monte Carlo (cont)

Some Observations about Metropolis-Hastings Method

1. To sample states with relative density ρ it is not necessary to know the normalization of ρ. Metropolis automatically samples $\rho(\mathbf{R}_\mathrm{i})/\int d\mathbf{R}_\mathrm{f}\ \rho(\mathbf{R}_\mathrm{f})$. So, it is useful for calculating quantities of the form

$$\frac{\int d\mathbf{R}_\mathrm{i}\ e(\mathbf{R}_\mathrm{i})\ \rho(\mathbf{R}_\mathrm{i})}{\int d\mathbf{R}_\mathrm{f}\ \rho(\mathbf{R}_\mathrm{f})}$$

   which is the form encountered in quantum mechanics and statistical mechanics, but not for doing importance sampling to calculate $\int d\mathbf{R}_\mathrm{i}\ f(\mathbf{R}_\mathrm{i})$, unless one has a $g(\mathbf{R}_\mathrm{i})$ that not only mimics $f(\mathbf{R}_\mathrm{i})$ but whose integral is known.

2. The variance of the estimate for the expectation value $\langle X \rangle$ is given by

$$\frac{1}{N/T_\mathrm{corr}-1}\left(\frac{\sum X(\mathbf{R}_\mathrm{i})^2}{N}-\left(\frac{\sum X(\mathbf{R}_\mathrm{i})}{N}\right)^2\right).$$

   That is, the effective number of configurations N_eff is smaller than N by a factor of T_corr, which we define to be the autocorrelation time.
   (T_corr is related to integrated autocorrelation time, $T_\mathrm{corr} = 1 + 2t_\mathrm{corr}$.)

Metropolis-Hastings Monte Carlo (cont)

Some Observations about Metropolis-Hastings Method

3. The rate of convergence to the desired density and the autocorrelation time of estimates of observables is governed by the sub-dominant eigenvalues of M. In practice reduce $T_{\rm corr}$ by inventing large moves that have large acceptance probabilities.
4. Folklore: when one can choose from a range of proposal matrices, the optimal one has an average acceptance ratio close to $1/2$.
   Reasonable choice in absence of any information, but in fact the optimal choice may have an average acceptance that is anywhere between zero and one.
   I have found instances where the optimum is as small as 0.2 or as large as 0.9.
   A much better criterion is to maximize the rate at which the system diffuses through configuration space $\langle A(\mathbf{R}_{\rm f}|\mathbf{R}_{\rm i})(\mathbf{R}_{\rm f}-\mathbf{R}_{\rm i})^2\rangle$.
   The real measure of goodness is of course to minimize the autocorrelation time for the observables of interest.

Metropolis-Hastings Monte Carlo (cont)
Some Observations about Metropolis-Hastings Method

5. One often employs elementary transition matrices M_i that are non-ergodic to construct a compound transition matrices M that is ergodic.
   Two sorts of combinations are often useful:
   1. $M = \prod_{i=1}^{n} M_i$. Sequential updating. e.g. Ising spins on lattice sites or the electrons in electronic structure.
   2. $M = \sum_{i=1}^{n} c_i M_i$, $c_i \geq 0$ (independent of the current state) and $\sum_{i=1}^{n} c_i = 1$. Choose the transitions M_i randomly with probabilities c_i.

      Sometimes one needs to make moves in a space that has both discrete and continuous components (e.g., spin or particle permutations and space). The individual Markov matrices may make transitions in only one of these spaces and therefore is not ergodic but the compound transition matrix is ergodic.

Metropolis-Hastings Monte Carlo (cont)
Some Observations about Metropolis-Hastings Method

7. The *Gibbs sampler* or *heat bath* algorithm can be considered to be a special case of the (generalized) Metropolis method. $T(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i}) \propto \rho(\mathbf{R}_\mathrm{f})$ for only a small set of accessible states in a domain $D(\mathbf{R}_\mathrm{i})$ in the neighborhood of $\mathbf{R}_\mathrm{i}$:

$$T(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i}) = \begin{cases} \rho(\mathbf{R}_\mathrm{f})/\sum\rho(\mathbf{R}_\mathrm{f}) & \text{if } \mathbf{R}_\mathrm{f} \,\epsilon\, D(\mathbf{R}_\mathrm{i}) \\ 0 & \text{otherwise} \end{cases}$$

If the sum over the accessible states from $\mathbf{R}_\mathrm{i}$ and $\mathbf{R}_\mathrm{f}$ is the same, the acceptance is unity.

$$\frac{A(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i})}{A(\mathbf{R}_\mathrm{i}|\mathbf{R}_\mathrm{f})} = \frac{T(\mathbf{R}_\mathrm{i}|\mathbf{R}_\mathrm{f})\ \rho(\mathbf{R}_\mathrm{f})}{T(\mathbf{R}_\mathrm{f}|\mathbf{R}_\mathrm{i})\ \rho(\mathbf{R}_\mathrm{i})} = 1\ .$$

The heat bath algorithm frequently used for lattice models where the normalization constant can be easily computed, e.g. Potts model.

One may ask, how are all states accessed if accessible states from $\mathbf{R}_\mathrm{i}$ and $\mathbf{R}_\mathrm{f}$ are the same? 1^{st} it is only the sum that must be the same. 2^{nd} it is OK because one may be cycling through the spins.

Optimizing the Markov matrix (discrete space)

Note that it was necessary to compute all the ρ_i in the thermalized space to do heat bath and this can be expensive if this space is not small. If one decides to incur this expense, then one can in fact do better than heat-bath by making all except one of the diagonal moves be zero as follows.

do $j = 1, n$

$S_< = \sum_{k=1}^{j-1} M(k,j)$

$S_> = \sum_{k=j+1}^{n} \pi(k)$

if $(j \neq n)$ then

$M(j+1:n\ ,\ j) = \pi(j+1:n)\ (1 - S_<)/S_>$

$M(j\ ,\ j+1:n) = M(j+1:n\ ,\ j)\ \pi(j)\ /\ \pi(j+1:n)$

$M(0,0) = 0$

else

$M(n,n) = 1 - S_<$

endif

enddo

Whereas the heat-bath matrix has one eigenvalue 1 and the rest zero, this matrix has one eigenvalue one and the rest negative. So, there are negative correlations (a good thing!) in this small space.

Optimizing the Markov matrix (discrete space)

1. Order states in order of increasing probability.
2. Starting with first column, make diagonal element zero and rescale the subdiagonal elements so that column adds up to 1.
3. After fixing each column, the corresponding row is determined by detailed balance.
4. For the last column there are no sub diagonal elements so the diagonal element is set by the condition that the column adds up to 1.

The new matrix has only one nonzero diagonal element – the last one. This procedure can be performed on any Markov matrix and the result will be different, but starting from the heat-bath matrix is probably a good choice.
Example: Suppose heat-bath matrix is

$$M = \begin{pmatrix} 0.1667 & 0.1667 & 0.1667 \\ 0.3333 & 0.3333 & 0.3333 \\ 0.5000 & 0.5000 & 0.5000 \end{pmatrix}$$

which of course has eigenvalues 1, 0, 0.
The optimized matrix is

$$M = \begin{pmatrix} 0.0000 & 0.2000 & 0.2000 \\ 0.4000 & 0.0000 & 0.5333 \\ 0.6000 & 0.8000 & 0.2667 \end{pmatrix}$$

which has eigenvalues 1., -0.2, -0.533333
Cute but probably not very useful since space must be small to construct this M.

Metropolis-Hastings Monte Carlo (cont)

Autocorrelation time

N Monte Carlo steps $=$ N_b blocks $\times$ N_s steps/block
Assume N_b is large enough that the block averages are nearly independent.

$\bar{E}$ $=$ average of $E_{\rm L}$ over the N MC steps
σ $=$ rms fluctuations of individual $E_{\rm L}$
σ_b $=$ rms fluctuations of block averages of $E_{\rm L}$

Effectively, $N/T_{\rm corr}$ independent measurements of $E_{\rm L}$

Define $T_{\rm corr}$ as

$$\mathrm{err}(\bar{E}) = \frac{\sigma}{\sqrt{N_b \times N_s}} \sqrt{T_{\rm corr}} = \frac{\sigma_b}{\sqrt{N_b}}$$

$$\Rightarrow \boxed{T_{\rm corr} = N_s \left(\frac{\sigma_b}{\sigma}\right)^2}$$

Choose $N_s \gg T_{\rm corr}$, say, 100 $T_{\rm corr}$.
If $N_s \approx 10\,T_{\rm corr}$, $T_{\rm corr}$ underest. $\approx 10\%$.

Metropolis-Hastings Monte Carlo (cont)

Autocorrelation time

N Monte Carlo steps $= N_b$ blocks $\times$ N_s steps/block

Assume N_b is large enough that the block averages are nearly independent.

$\bar{E}$ = average of $E_{\rm L}$ over the N MC steps

σ = rms fluctuations of individual $E_{\rm L}$

σ_b = rms fluctuations of block averages of $E_{\rm L}$

Effectively, $N/T_{\rm corr}$ independent measurements of $E_{\rm L}$

Define $T_{\rm corr}$ as

$$\mathrm{err}(\bar{E}) = \frac{\sigma}{\sqrt{\frac{N}{T_{\rm corr}} - 1}} = \frac{\sigma_b}{\sqrt{N_b - 1}}$$

$$\Rightarrow \boxed{T_{\rm corr} = \frac{N - T_{\rm corr}}{N_b - 1}\left(\frac{\sigma_b}{\sigma}\right)^2 \approx N_s\left(\frac{\sigma_b}{\sigma}\right)^2}$$

Choose $N_s \gg T_{\rm corr}$, say, 100 $T_{\rm corr}$.

If $N_s \approx 10\,T_{\rm corr}$, $T_{\rm corr}$ underest. $\approx 10\%$.

Blocking Analysis for error of mean of autocorrelated variables

Compute recursively and plot

$$\frac{1}{N_b(N_b-1)}\sum_{i=1}^{N_b}(m_i-\bar{E})^2$$

for various blocking levels, $N_s = 1,\ 2,\ 2^2,\ 2^3,\ ...,\ N/2$

If the variables were uncorrelated to begin with then these estimates of the error would be equal aside from statistical fluctuations.

If they are autocorrelated, the estimated error will grow and the flatten out when the block means become uncorrelated, which assumes that $N \gg T_{\rm corr}$.

Assuming that block means are independent Gaussian variables (they are not at the lower blocking levels), the estimated uncertainty of the error is

$$\frac{error\ estim}{\sqrt{2(N_b-1)}}$$

using the fact that the probability density function of the sum of squares of $N_b - 1$ normal standard deviates has variance $2(N_b - 1)$.

A reasonable choice of blocking level is the highest one for which the increase in the estimate for the error is larger than the increase in the estimate for the uncertainty in the error. It is possible to get a somewhat better estimate by predicting the shape of the curve and extrapolating when say $N < 1000\,T_{\rm corr}$.

Blocking Analysis for error of mean of autocorrelated variables

In variational Monte Carlo, $T_{\rm corr}$ is usually very small if one makes an intelligent choice for the proposal matrix. With the algorithm we typically use $T_{\rm corr} < 2$ even for systems with say 100 electrons!

However, in some of the projector Monte Carlo methods (e.g. FCIQMC and SQMC) that we discuss next, $T_{\rm corr}$ can be much larger, even for much smaller systems. Further, in these methods one needs to use a large population of walkers, so it becomes expensive to have a large number of Monte Carlo steps. In the next viewgraph, a blocking analysis for a run with $T_{\rm corr} \approx 1000$ and $N = 2^{23}$ is shown.

Blocking Analysis for error of mean of autocorrelated variables

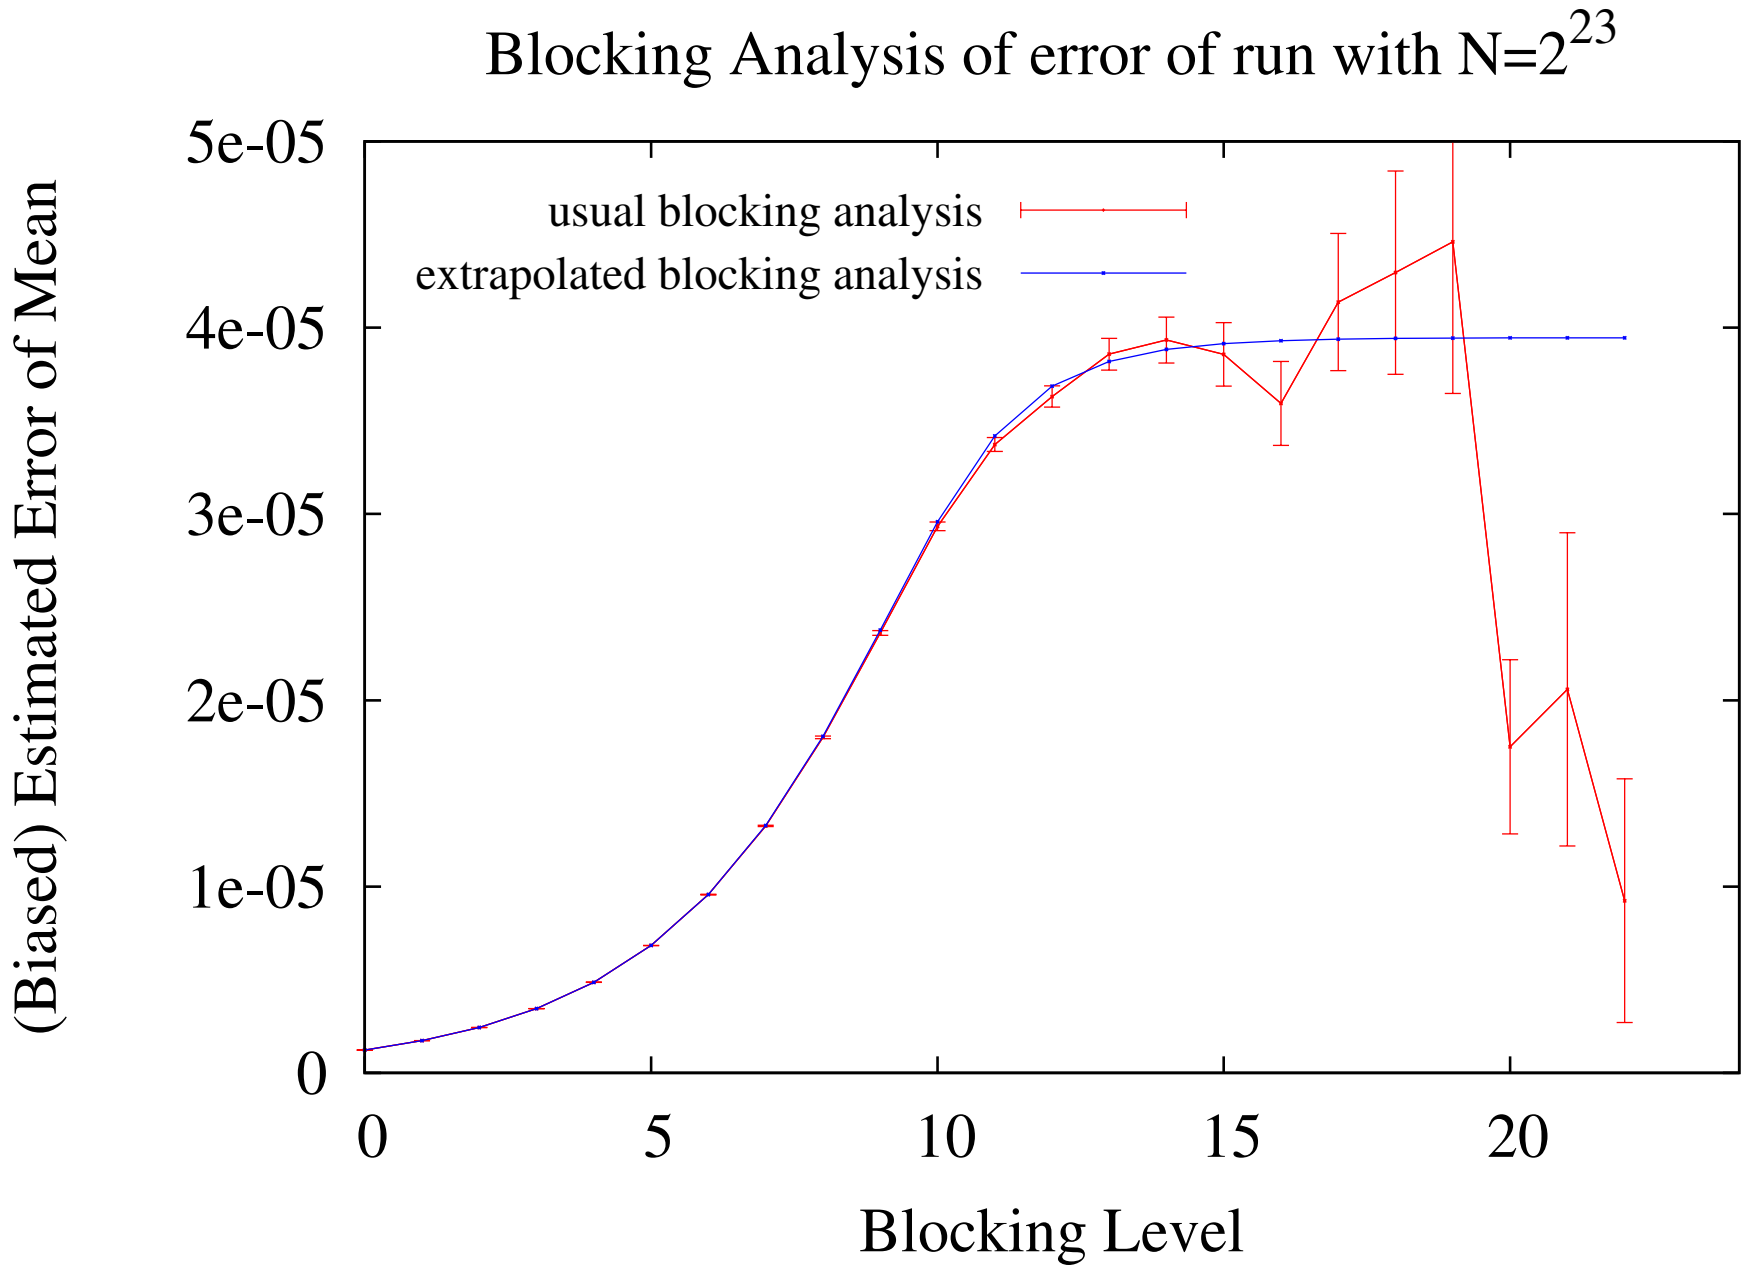

Lecture 3
Optimization of Many-body Wavefunctions

Cyrus J. Umrigar

Functional form of Trial Wave Function

One of the great advantages of QMC is that one has a great deal of freedom in the functional form of the trial wavefunction. This is in contrast to other methods where one is largely restricted to using linear combinations of determinants, and, furthermore the orbitals in these determinants have to be expanded in basis functions, such as gaussians or planewaves, that are amenable to analytic integration.

In QMC one has can utilize one's intuition about the physics or chemistry of the problem to come up with good functional forms for the wavefunction. These functional forms may have several parameters, whose values are not know a priori, so powerful methods for optimizing these parameters have been developed.

Some innovative functional forms that have been used are:

1. Antisymmetrized geminal power times Jastrow Sorella, Casula
2. Pfaffian times Jastrow Schmidt, Mitas and coworkers
3. Inhomogeneous backflow times Jastrow Needs and coworkers

Next we discuss what is probably the most commonly used functional form – multideterminant expansion times Jastrow.

Functional form of Trial Wave Function

$$\Psi_T = \left(\sum_n d_n \mathrm{D}_n^{\uparrow} \mathrm{D}_n^{\downarrow} \right) \times \mathcal{J}(r_i, r_j, r_{ij})$$

- Determinants: $\sum_n d_n \mathrm{D}_n^{\uparrow} \mathrm{D}_n^{\downarrow}$

$\mathrm{D}^{\uparrow}$ and $\mathrm{D}^{\downarrow}$ are determinants of single-particle orbitals ϕ for up ($\uparrow$) and down ($\downarrow$) spin electrons respectively.

The single-particle orbitals ϕ are given by:

$$\phi(\mathbf{r}_i) = \sum_{\alpha k} c_{k_\alpha} \; N_{k_\alpha} \, r_{i\alpha}^{n_{k_\alpha}-1} \, e^{-\zeta_{k_\alpha} r_{i\alpha}} \, \mathrm{Y}_{l_{k_\alpha} m_{k_\alpha}}(\hat{\mathrm{r}}_{i\alpha})$$

- Jastrow: $\mathcal{J}(r_i, r_j, r_{ij}) = \prod_{\alpha i} \exp(A_{\alpha i}) \prod_{ij} \exp(B_{ij}) \prod_{\alpha ij} \exp(C_{\alpha ij})$

  $A_{\alpha i} \Rightarrow$ electron-ion correlation

  $B_{ij} \Rightarrow$ electron-electron correlation

  $C_{\alpha ij} \Rightarrow$ electron-electron-ion correlation

d_n, c_{k_α}, ζ_{k_α} and parms in $\mathcal{J}$ are optimized.

$\sim N_{\mathrm{atomtype}}$ of $\mathcal{J}$ parms.

$\sim N_{\mathrm{atomtype}}$ of ζ_{k_α} parms.

$\sim N_{\mathrm{atom}}^2$ of c_{k_α} parms.

$\sim e^{N_{\mathrm{atom}}}$ of d_n parms.

Power of QMC:

$\mathcal{J}$ parms. do work of d_n parms.

Optimization of many-body wavefunctions

Standard methods do not work for the wavefunctions we are interested in. For example, they work for linear combinations of determinants but not for linear combinations of determinants multiplied by a Jastrow factor.

But is it a worthwhile expenditure of effort to optimize wavefunctions?

Almost all errors reduced by optimizing trial wavefunctions

1. Statistical error (both the rms fluctuations of $E_{\rm L}$ and the autocorrelation time)
2. $E_{\rm VMC}$
3. Fixed-node error in $E_{\rm DMC}$ (nodes move during optimization). Fixed node errors can be LARGE. For C_2, FN error for 1-det wavefn is 1.3 eV for total energy and 0.7 eV for well-depth. However, optimized multidet. wavefn has FN error that is better than chemical accuracy (1 kcal/mole $=$ 0.043 eV/molecule).
4. Time-step error in DMC
5. Population control error in DMC
6. Pseudopotential locality error in DMC when using nonlocal pseudopotentials
7. Error of observables that do not commute with the Hamiltonian (mixed estimators, $\langle \Psi_0 | \hat{A} | \Psi_T \rangle$ not exact even for nodeless ψ_0, $\psi_{\rm T}$) if one does not use forward/side walking.

Choices to be made when optimizing trial wavefunctions

1. What precisely do we want to optimize – the objective function or measure of goodness?
2. What method do we use to do the optimization? If more than one method is applied to the same objective function, they will of course give the same wavefunction, but the efficiency with which we arrive at the solution may be much different.
3. When we test to see if the proposed new parameters are better than the old ones, do we test on a fixed sample of MC points or draw new MC points each time?

Measures of goodness of variational wave functions

$$\min \; E_{\rm VMC} = \frac{\langle \psi_{\rm T}|H|\psi_{\rm T}\rangle}{\langle \psi_{\rm T}|\psi_{\rm T}\rangle} = \langle E_{\rm L}\rangle_{|\psi_{\rm T}|^2}$$

$$\min \; \sigma^2_{\rm VMC} = \frac{\langle \psi_{\rm T}|(H-E_{\rm T})^2|\psi_{\rm T}\rangle}{\langle \psi_{\rm T}|\psi_{\rm T}\rangle} = \langle E_{\rm L}^2(\mathbf{R}_i)\rangle_{|\psi_{\rm T}|^2} - \langle E_{\rm L}(\mathbf{R}_i)\rangle^2_{|\psi_{\rm T}|^2}$$

$$\max \quad \Omega^2 = \frac{|\langle \psi_{\rm FN}|\psi_{\rm T}\rangle|^2}{\langle \psi_{\rm FN}|\psi_{\rm FN}\rangle \langle \psi_{\rm T}|\psi_{\rm T}\rangle} = \frac{\left\langle \frac{\psi_{\rm FN}}{\psi_{\rm T}} \right\rangle^2_{|\psi_{\rm T}|^2}}{\left\langle \left| \frac{\psi_{\rm FN}}{\psi_{\rm T}} \right|^2 \right\rangle_{|\psi_{\rm T}|^2}}$$

$$\min \; E_{\rm DMC} = \frac{\langle \psi_{\rm FN}|H|\psi_{\rm T}\rangle}{\langle \psi_{\rm FN}|\psi_{\rm T}\rangle} = \langle E_{\rm L}\rangle_{|\psi_{\rm FN}\psi_{\rm T}|}$$

For an infinitely flexible wave function all optimizations will yield the exact wavefunction (except that minimizing σ could yield an excited state) but for the imperfect functional forms used in practice they differ.

Progress in optimization of Many-Body Wavefunctions

Naive energy optim. $\rightarrow$ Variance optim. $\rightarrow$ Efficient energy optim.

	–	1988	naive energy optimization, few (~ 3) parameters
1988	–	2001	variance optimization, ~ 100 parameters could be used for more, but, variance does not couple strongly to some parameters
2001	–		efficient energy optimization, ~ 1000's of parameters

Optimization of Many-Body Wavefunctions

A major advantage of quantum Monte Carlo methods is that there is no restriction on the form of $\psi_{\rm T}(\mathbf{R})$. Hence any insight one may have, as regards the nature of the many-body correlations, can be built into $\psi_{\rm T}(\mathbf{R})$ and tested. To exploit this freedom it is necessary to have a method for optimizing arbitrary wavefunctions.

First thought: Minimize the energy on MC sample.

$$\bar{E} = \sum_{i=1}^{N_{\rm conf}} \frac{\mathcal{H}\psi_{\rm T}(\mathbf{R}_i;\{p\})}{\psi_{\rm T}(\mathbf{R}_i;\{p\})}\, w_i, \qquad w_i = \left|\frac{\Psi_{\rm T}(\mathbf{R}_i)}{\Psi^0_{\rm T}(\mathbf{R}_i)}\right|^2 \Bigg/ \sum_{i=1}^{N_{\rm conf}} \left|\frac{\Psi_{\rm T}(\mathbf{R}_i)}{\Psi^0_{\rm T}(\mathbf{R}_i)}\right|^2$$

Second thought: Minimize the variance of the local energy.

$$\sigma^2 = \sum_{i=1}^{N_{\rm conf}} \left(\frac{\mathcal{H}\psi_{\rm T}(\mathbf{R}_i;\{p\})}{\psi_{\rm T}(\mathbf{R}_i;\{p\})} - \bar{E}\right)^2 w_i$$

Third thought: Minimize the energy using MC but not on MC sample.

Variance vs. Energy

$$\sigma^2 = \sum_{i=1}^{N_{\rm conf}} \left(\frac{\mathcal{H}\Psi_{\rm T}(\mathbf{R}_i)}{\Psi_{\rm T}(\mathbf{R}_i)} - \bar{E} \right)^2$$

$$\bar{E} = \sum_{i=1}^{N_{\rm conf}} \frac{\mathcal{H}\Psi_{\rm T}(\mathbf{R}_i)}{\Psi_{\rm T}(\mathbf{R}_i)}$$

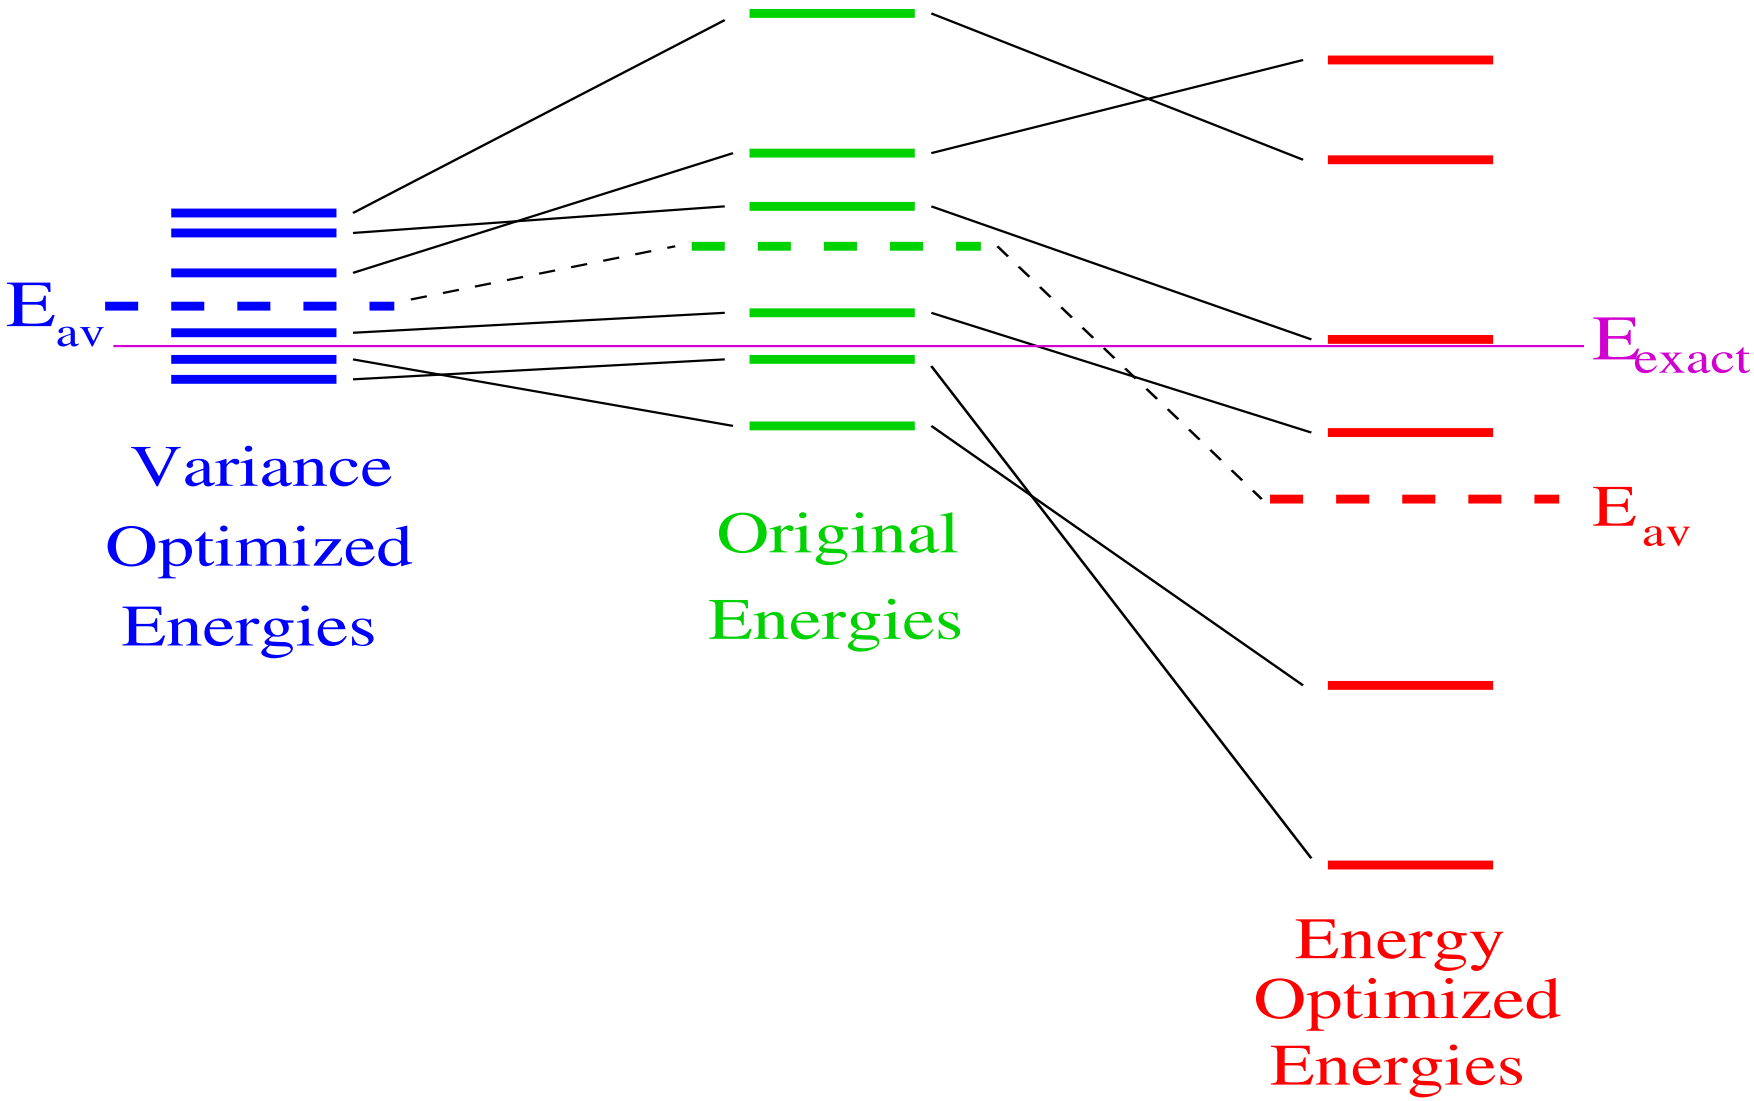

Take-home Message

Energy optimization methods that minimize the energy evaluated on finite sample will yield poor energies on other samples, unless the sample used to do the minimization is very large.

So, efficient energy optimization methods do NOT optimize the energy evaluated on a finite sample, although they do minimize the energy in the limit of an infinite sample.

Advantages of Energy (or Mixed) Optim. vs. Variance Optim.

1. Want lowest energy; fluctuations are of secondary importance. Energy and variance do not always go hand-in-hand enough.
2. Some parameters couple more strongly to energy than variance.
3. Some variance-optimized parameters make wave function too extended.
4. Hellman-Feynman theorem can be used for forces (when combined with variance reduction methods).

Optimization Methods

The optimization methods we use are based on standard methods:

1. Levenberg-Marquardt method
2. Newton method
3. Linear method (though with significant extension to nonlinear parameters)
4. Perturbation theory

However, all of them need additional ingredients to work with stochastic methods, and these ingredients improve the efficiency of the method by several orders of magnitude!

Newton Method

Calculate gradient $\mathbf{g}$ and Hessian $\mathbf{h}$ of objective function and update parameters:

$$\mathbf{p}_{\rm next} = \mathbf{p}_{\rm current} - \mathbf{h}^{-1}\mathbf{g}$$

or more efficiently ($\mathcal{O}(N_p^2)$ vs. $\mathcal{O}(N_p^3)$) find parameter changes, $\delta\mathbf{p}$, by solving linear equations:

$$\mathbf{h}\ \delta\mathbf{p} = -\mathbf{g},$$

Optimization of Jastrow and determinantal parameters encounter different problems.

Jastrow: For the form of the Jastrow we use and the systems we study the eigenvalues of the Hessian span 10-12 orders of magnitude. So using steepest descent is horribly slow and using the Hessian, or a reasonable approximation to it, is essential even if there were no statistical noise.

determinantal: The eigenvalues of the Hessian span only 1-2 orders of magnitude. However, the Hessian has terms involving

$$\frac{\frac{\partial\psi}{\partial p_i}}{\psi}$$

that diverge as $\psi \to 0$. The strongest divergence among various terms cancels.

Energy Minimization via Newton

Lin, Zhang, Rappe, JCP 2000; CJU, Filippi, PRL 2005

$$\bar{E} = \frac{\langle\psi|H|\psi\rangle}{\langle\psi|\psi\rangle} = \langle E_{\rm L}\rangle_{\psi^2}; \qquad E_{\rm L}(\mathbf{R}) = \frac{H\psi(\mathbf{R})}{\psi(\mathbf{R})}$$

Energy gradient components, $\bar{E}_i$:

$$\bar{E}_i = \frac{\langle\psi_i|H\psi\rangle + \langle\psi|H\psi_i\rangle}{\langle\psi|\psi\rangle} - 2\frac{\langle\psi|H|\psi\rangle\,\langle\psi|\psi_i\rangle}{\langle\psi|\psi\rangle^2}$$

$$= \frac{\langle\psi_i|H\psi\rangle + \langle\psi|H\psi_i\rangle}{\langle\psi|\psi\rangle} - 2\frac{\bar{E}\,\langle\psi|\psi_i\rangle}{\langle\psi|\psi\rangle} = 2\frac{\langle\psi_i|H\psi\rangle - \bar{E}\,\langle\psi|\psi_i\rangle}{\langle\psi|\psi\rangle} \quad \text{(by Hermiticity)}$$

$$= \left\langle \frac{\psi_i}{\psi} E_{\rm L} + \frac{H\psi_i}{\psi} - 2\bar{E}\frac{\psi_i}{\psi} \right\rangle_{\psi^2} = 2\left\langle \frac{\psi_i}{\psi}(E_{\rm L} - \bar{E}) \right\rangle_{\psi^2}$$

Is blue or green expression better for MC?

Energy Minimization via Newton

Lin, Zhang, Rappe, JCP 2000; CJU, Filippi, PRL 2005

$$\bar{E} = \frac{\langle\psi|H|\psi\rangle}{\langle\psi|\psi\rangle} = \langle E_{\rm L}\rangle_{\psi^2}; \qquad E_{\rm L}(\mathbf{R}) = \frac{H\psi(\mathbf{R})}{\psi(\mathbf{R})}$$

Energy gradient components, $\bar{E}_i$:

$$\bar{E}_i = \frac{\langle\psi_i|H\psi\rangle + \langle\psi|H\psi_i\rangle}{\langle\psi|\psi\rangle} - 2\frac{\langle\psi|H|\psi\rangle\,\langle\psi|\psi_i\rangle}{\langle\psi|\psi\rangle^2}$$

$$= \frac{\langle\psi_i|H\psi\rangle + \langle\psi|H\psi_i\rangle}{\langle\psi|\psi\rangle} - 2\frac{\bar{E}\,\langle\psi|\psi_i\rangle}{\langle\psi|\psi\rangle} = 2\frac{\langle\psi_i|H\psi\rangle - \bar{E}\,\langle\psi|\psi_i\rangle}{\langle\psi|\psi\rangle} \quad \text{(by Hermiticity)}$$

$$= \left\langle \frac{\psi_i}{\psi} E_{\rm L} + \frac{H\psi_i}{\psi} - 2\bar{E}\frac{\psi_i}{\psi} \right\rangle_{\psi^2} = 2\left\langle \frac{\psi_i}{\psi}(E_{\rm L} - \bar{E}) \right\rangle_{\psi^2}$$

Is blue or green expression better for MC?

Green is better because it is a zero-variance expression in the limit that ψ is the exact ground state (CJU, Filippi, PRL 2005) Moreover it is simpler and faster.

Energy Minimization via Newton

CJU, Filippi, PRL 2005

Energy hessian components, E_{ij}:

$$\bar{E}_i = 2\frac{\langle\psi_i|H\psi\rangle - \bar{E}\,\langle\psi|\psi_i\rangle}{\langle\psi|\psi\rangle} \equiv 2\frac{\langle\psi_i\psi(E_{\rm L}-\bar{E})\rangle}{\langle\psi^2\rangle}$$

$$E_{ij} = 2\left[\frac{\langle(\psi_{ij}\psi+\psi_i\psi_j)(E_{\rm L}-\bar{E})\rangle + \langle\psi_i\psi(E_{{\rm L},j}-\bar{E}_j)\rangle - \bar{E}_i\,\langle\psi\psi_j\rangle}{\langle\psi^2\rangle}\right]$$

$$= 2\left[\left\langle\left(\frac{\psi_{ij}}{\psi}+\frac{\psi_i\psi_j}{\psi^2}\right)(E_{\rm L}-\bar{E})\right\rangle_{\psi^2} - \left\langle\frac{\psi_i}{\psi}\right\rangle_{\psi^2}\bar{E}_j - \left\langle\frac{\psi_j}{\psi}\right\rangle_{\psi^2}\bar{E}_i + \left\langle\frac{\psi_i}{\psi}E_{{\rm L},j}\right\rangle_{\psi^2}\right].$$

What can be done to improve this expression?

Energy Minimization via Newton

CJU, Filippi, PRL 2005

Energy hessian components, E_{ij}:

$$\bar{E}_i = 2\frac{\langle\psi_i|H\psi\rangle - \bar{E}\,\langle\psi|\psi_i\rangle}{\langle\psi|\psi\rangle} \equiv 2\frac{\langle\psi_i\psi(E_{\rm L}-\bar{E})\rangle}{\langle\psi^2\rangle}$$

$$E_{ij} = 2\left[\frac{\langle(\psi_{ij}\psi+\psi_i\psi_j)(E_{\rm L}-\bar{E})\rangle + \langle\psi_i\psi(E_{{\rm L},j}-\bar{E}_j)\rangle - \bar{E}_i\,\langle\psi\psi_j\rangle}{\langle\psi^2\rangle}\right]$$

$$= 2\left[\left\langle\left(\frac{\psi_{ij}}{\psi}+\frac{\psi_i\psi_j}{\psi^2}\right)(E_{\rm L}-\bar{E})\right\rangle_{\psi^2} - \left\langle\frac{\psi_i}{\psi}\right\rangle_{\psi^2}\bar{E}_j - \left\langle\frac{\psi_j}{\psi}\right\rangle_{\psi^2}\bar{E}_i + \left\langle\frac{\psi_i}{\psi}E_{{\rm L},j}\right\rangle_{\psi^2}\right].$$

What can be done to improve this expression?

1) Symmetrize – but this does not reduce fluctuations much

Energy Minimization via Newton

CJU, Filippi, PRL 2005

Energy hessian components, E_{ij}:

$$\bar{E}_i = 2\frac{\langle\psi_i|H\psi\rangle - \bar{E}\,\langle\psi|\psi_i\rangle}{\langle\psi|\psi\rangle} \equiv 2\frac{\langle\psi_i\psi(E_{\rm L}-\bar{E})\rangle}{\langle\psi^2\rangle}$$

$$E_{ij} = 2\left[\frac{\langle(\psi_{ij}\psi+\psi_i\psi_j)(E_{\rm L}-\bar{E})\rangle + \langle\psi_i\psi(E_{{\rm L},j}-\bar{E}_j)\rangle - \bar{E}_i\,\langle\psi\psi_j\rangle}{\langle\psi^2\rangle}\right]$$

$$= 2\left[\left\langle\left(\frac{\psi_{ij}}{\psi}+\frac{\psi_i\psi_j}{\psi^2}\right)(E_{\rm L}-\bar{E})\right\rangle_{\psi^2} - \left\langle\frac{\psi_i}{\psi}\right\rangle_{\psi^2}\bar{E}_j - \left\langle\frac{\psi_j}{\psi}\right\rangle_{\psi^2}\bar{E}_i + \left\langle\frac{\psi_i}{\psi}E_{{\rm L},j}\right\rangle_{\psi^2}\right].$$

What can be done to improve this expression?

1) Symmetrize – but this does not reduce fluctuations much

2) Noting that $\langle E_{{\rm L},j}\rangle_{\psi^2} = \dfrac{\left\langle\psi^2\left(\frac{H\psi}{\psi}\right)_j\right\rangle}{\langle\psi^2\rangle} = \dfrac{\left\langle\psi^2\left(\frac{H\psi_j}{\psi}-\frac{\psi_j}{\psi^2}H\Psi\right)\right\rangle}{\langle\psi^2\rangle} = \dfrac{\langle\psi H\psi_i-\psi_i H\psi\rangle}{\langle\psi^2\rangle} = 0$ by hermiticity of $\hat{H}$, and, that the fluctuations of the covariance $\langle ab\rangle - \langle a\rangle\langle b\rangle$ are smaller than those of the product $\langle ab\rangle$, when $\sqrt{\langle a^2\rangle - \langle a\rangle^2} \ll |\langle a\rangle|$ and $\langle b\rangle = 0$ on ∞ sample but $\langle b\rangle \neq 0$ on finite sample, replace

$$\left\langle\frac{\psi_i}{\psi}E_{{\rm L},j}\right\rangle_{\psi^2} \rightarrow \frac{1}{2}\left(\left\langle\frac{\psi_i}{\psi}E_{{\rm L},j}\right\rangle_{\psi^2} - \left\langle\frac{\psi_i}{\psi}\right\rangle_{\psi^2}\langle E_{{\rm L},j}\rangle_{\psi^2} + \left\langle\frac{\psi_j}{\psi}E_{{\rm L},i}\right\rangle_{\psi^2} - \left\langle\frac{\psi_j}{\psi}\right\rangle_{\psi^2}\langle E_{{\rm L},i}\rangle_{\psi^2}\right)$$

3) Too hard to describe here.

Energy Minimization via Newton

$$\bar{E}_{ij} = 2\left[\left\langle\left(\frac{\psi_{ij}}{\psi}+\frac{\psi_i\psi_j}{\psi^2}\right)(E_{\rm L}-\bar{E})\right\rangle_{\psi^2} - \left\langle\frac{\psi_i}{\psi}\right\rangle_{\psi^2}\bar{E}_j - \left\langle\frac{\psi_j}{\psi}\right\rangle_{\psi^2}\bar{E}_i\right]$$

$$+\left\langle\frac{\psi_i}{\psi}E_{{\rm L},j}\right\rangle_{\psi^2} - \left\langle\frac{\psi_i}{\psi}\right\rangle_{\psi^2}\langle E_{{\rm L},j}\rangle_{\psi^2} + \left\langle\frac{\psi_j}{\psi}E_{{\rm L},i}\right\rangle_{\psi^2} - \left\langle\frac{\psi_j}{\psi}\right\rangle_{\psi^2}\langle E_{{\rm L},i}\rangle_{\psi^2}$$

$$= 2\left[\left\langle\left(\frac{\psi_{ij}}{\psi}-\frac{\psi_i\psi_j}{\psi^2}\right)(E_{\rm L}-\bar{E})\right\rangle_{\psi^2}\right. \quad \text{(0 for } p_i \text{ linear in exponent)}$$

$$\left.+2\left\langle\left(\frac{\psi_i}{\psi}-\left\langle\frac{\psi_i}{\psi}\right\rangle_{\psi^2}\right)\left(\frac{\psi_j}{\psi}-\left\langle\frac{\psi_j}{\psi}\right\rangle_{\psi^2}\right)(E_{\rm L}-\bar{E})\right\rangle_{\psi^2}\right]$$

$$+\left\langle\frac{\psi_i}{\psi}E_{{\rm L},j}\right\rangle_{\psi^2} - \left\langle\frac{\psi_i}{\psi}\right\rangle_{\psi^2}\langle E_{{\rm L},j}\rangle_{\psi^2} + \left\langle\frac{\psi_j}{\psi}E_{{\rm L},i}\right\rangle_{\psi^2} - \left\langle\frac{\psi_j}{\psi}\right\rangle_{\psi^2}\langle E_{{\rm L},i}\rangle_{\psi^2}\,.$$

1) Blue and green terms are zero variance estimators.
2) Red terms are not, but, additional terms =0 for infinite sample and cancel most of the fluctuations for a finite sample.

Energy optimization of determinantal parameters via Newton

Different issues arise in optimizing Jastrow parameters and determinantal parameters:
Jastrow: eigenvalues of Hessian have a range of 11 orders of magnitude!
Determinantal parameters: divergences in elements of Hessian and Hamiltonian matrices.

$$\bar{E}_i \;=\; 2\left\langle \frac{\psi_i}{\psi}(E_{\rm L}-\bar{E})\right\rangle_{\psi^2} \qquad \text{(by Hermiticity)}.$$

Linear divergence – no problem since weighted by $|\psi|^2$.

$$\bar{E}_{ij} = 2\left[\left\langle\left(\frac{\psi_{ij}}{\psi}+\frac{\psi_i\psi_j}{\psi^2}\right)(E_{\rm L}-\bar{E})\right\rangle_{\psi^2} - \left\langle\frac{\psi_i}{\psi}\right\rangle_{\psi^2}\bar{E}_j - \left\langle\frac{\psi_j}{\psi}\right\rangle_{\psi^2}\bar{E}_i\right]$$
$$+\left\langle\frac{\psi_i}{\psi}E_{{\rm L},j}\right\rangle_{\psi^2} - \left\langle\frac{\psi_i}{\psi}\right\rangle_{\psi^2}\langle E_{{\rm L},j}\rangle_{\psi^2} + \left\langle\frac{\psi_j}{\psi}E_{{\rm L},i}\right\rangle_{\psi^2} - \left\langle\frac{\psi_j}{\psi}\right\rangle_{\psi^2}\langle E_{{\rm L},i}\rangle_{\psi^2}$$

Leading divergences (3^{rd}-order) cancel! (They cancel in the linear method too.) 2^{nd} order divergences would give ∞ variance, but in practice do not seem to be problematic. Do 2^{nd} order ones cancel?

Variance minimization via Newton

The parameters p_i in a real-valued trial wave function ψ are varied to minimize the variance of the local energy,

$$\sigma^2 = \frac{\int \mathrm{d}^{3N}R\ \psi^2 (E_\mathrm{L} - \bar{E})^2}{\int \mathrm{d}^{3N}R\ \psi^2} = \left\langle (E_\mathrm{L} - \bar{E})^2 \right\rangle_{\psi^2} .$$

$E_\mathrm{L} = H\psi/\psi$ is the local energy; $\quad \bar{E} = \langle E_\mathrm{L} \rangle$

The derivative wrt p_i, is

$$(\sigma^2)_i = 2\left[\left\langle E_{\mathrm{L},i}(E_\mathrm{L} - \bar{E}) \right\rangle_{\psi^2} + \left\langle \frac{\psi_i}{\psi} E_\mathrm{L}^2 \right\rangle_{\psi^2} - \left\langle \frac{\psi_i}{\psi} \right\rangle_{\psi^2} \left\langle E_\mathrm{L}^2 \right\rangle_{\psi^2} - 2\bar{E} \left\langle \frac{\psi_i}{\psi} (E_\mathrm{L} - \bar{E}) \right\rangle_{\psi^2} \right].$$

Note, the above is a zero-variance expression for the gradient.

If we ignore change in ψ and minimize on fixed MC pts. (no reweighting):

$$(\sigma^2)_i \quad = \quad 2\left\langle E_{\mathrm{L},i}(E_\mathrm{L} - \bar{E}) \right\rangle_{\psi^2} \ = \ 2\left\langle (E_{\mathrm{L},i} - \bar{E}_i)(E_\mathrm{L} - \bar{E}) \right\rangle_{\psi^2}$$

and setting $E_{ij} = E_{\mathrm{L},ij} = 0$ (Levenberg-Marquardt approximation) we get the (+ve definite) approx. to Hessian:

$$(\sigma^2)_{ij} \quad = \quad 2\left\langle (E_{\mathrm{L},i} - \bar{E}_i)(E_{\mathrm{L},j} - \bar{E}_j) \right\rangle_{\psi^2} .$$

This simple hessian works as well or better than the hessian we tried that does not make either approximation.

Linear method for linear parameters

Symmetric or nonsymmetric H?

1) true **H** is symmetric:

$$\text{True} \quad H_{ij} = \int \mathrm{d}^{3N}R\, \psi_i(\mathbf{R})\, \hat{H}\, \psi_j(\mathbf{R}) \quad \text{symmetric}$$

$$\text{MC estim.} \quad H_{ij} = \sum_{n=1}^{N_{\mathrm{MC}}} \frac{\psi_i(\mathbf{R}_n)}{\psi(\mathbf{R}_n)} \left(\frac{\hat{H}\psi_j(\mathbf{R}_n)}{\psi(\mathbf{R}_n)} \right) \quad \text{nonsymmetric}$$

$$\text{MC estim.} \quad H_{ij} = \frac{1}{2} \sum_{n=1}^{N_{\mathrm{MC}}} \left(\frac{\psi_i(\mathbf{R}_n)}{\psi(\mathbf{R}_n)} \frac{\hat{H}\psi_j(\mathbf{R}_n)}{\psi(\mathbf{R}_n)} + \frac{\hat{H}\psi_i(\mathbf{R}_n)}{\psi(\mathbf{R}_n)} \frac{\psi_j(\mathbf{R}_n)}{\psi(\mathbf{R}_n)} \right) \quad \text{symmetric}$$

Linear method for linear parameters

Symmetric or nonsymmetric H?

1) true **H** is symmetric:

$$\text{True} \quad H_{ij} = \int \mathrm{d}^{3N}R\, \psi_i(\mathbf{R})\, \hat{H}\, \psi_j(\mathbf{R}) \quad \text{symmetric}$$

$$\text{MC estim.} \quad H_{ij} = \sum_{n=1}^{N_{\mathrm{MC}}} \frac{\psi_i(\mathbf{R}_n)}{\psi(\mathbf{R}_n)} \left(\frac{\hat{H}\psi_j(\mathbf{R}_n)}{\psi(\mathbf{R}_n)} \right) \quad \text{nonsymmetric}$$

$$\text{MC estim.} \quad H_{ij} = \frac{1}{2} \sum_{n=1}^{N_{\mathrm{MC}}} \left(\frac{\psi_i(\mathbf{R}_n)}{\psi(\mathbf{R}_n)} \frac{\hat{H}\psi_j(\mathbf{R}_n)}{\psi(\mathbf{R}_n)} + \frac{\hat{H}\psi_i(\mathbf{R}_n)}{\psi(\mathbf{R}_n)} \frac{\psi_j(\mathbf{R}_n)}{\psi(\mathbf{R}_n)} \right) \quad \text{symmetric}$$

2) Minimizing the energy evaluated on a finite sample, i.e., minimizing the Rayleigh quotient, $\partial E/\partial p_k = 0$, even with nonsymmetric H evaluated on finite sample, gives generalized eigenvalue equation with symmetric **H**:

$$E = \min_{\mathbf{p}} \frac{\mathbf{p}^{\mathsf{T}}\mathbf{H}\mathbf{p}}{\mathbf{p}^{\mathsf{T}}\mathbf{S}\mathbf{p}} = \min_{\mathbf{p}} \frac{\sum_{ij} p_i H_{ij} p_j}{\sum_{ij} p_i S_{ij} p_j}$$

$$\frac{\partial E}{\partial p_k} = 0 \implies \left(\sum_{ij} p_i S_{ij} p_j \right) \left(\sum_j H_{kj} p_j + \sum_i p_i H_{ik} \right) - \left(\sum_{ij} p_i H_{ij} p_j \right) \left(2 \sum_j S_{kj} p_j \right) = 0$$

$$\frac{(\mathbf{H} + \mathbf{H}^{\mathsf{T}})}{2}\mathbf{p} = E\mathbf{S}\mathbf{p}$$

Nonsymm. H satisfies strong zero-variance principle

M. P. Nightingale and Melik-Alaverdian, PRL, **87**, 043401 (2001).

Nightingale's strong zero-variance principle:
If the states $\psi_i(\mathbf{R})$ are closed under $\hat{H}$ then the values of the optimized parameters using nonsymmetric H_{ij} are independent of the MC sample

Proof: If closed $\exists\ \{p_j\}$ s.t. $$\hat{H}\sum_{j=1}^{N_p} p_j|\psi_j\rangle \;=\; E\sum_{j=1}^{N_p} p_j|\psi_j\rangle$$

$\times\ \langle\psi_i|\mathbf{R}_n\rangle\langle\mathbf{R}_n|/\langle\psi|\mathbf{R}_n\rangle^2$ and sum over $N_{\rm MC}$ pts. (not complete sum over $\mathbf{R}$ states), sampled from $|\psi(\mathbf{R})|^2$:

$$\sum_{j=1}^{N_p} p_j \underbrace{\sum_{n=1}^{N_{\rm MC}} \frac{\langle\psi_i|\mathbf{R}_n\rangle}{\langle\psi|\mathbf{R}_n\rangle}\frac{\langle\mathbf{R}_n|\hat{H}|\psi_j\rangle}{\langle\mathbf{R}_n|\psi\rangle}}_{H_{ij}} \;=\; E\sum_{j=1}^{N_p} p_j \underbrace{\frac{\sum_{n=1}^{N_{\rm MC}}\langle\psi_i|\mathbf{R}_n\rangle\langle\mathbf{R}_n|\psi_j\rangle}{\langle\psi|\mathbf{R}_n\rangle\langle\mathbf{R}_n|\psi\rangle}}_{S_{ij}}$$

$$\sum_{i=1}^{N_p} H_{ij}p_j \;=\; E\sum_{i=1}^{N_p} S_{ij}p_j$$

H is nonsymmetric H of previous slide. Becomes symmetric when $\sum \to \int$.

Nonsymm. H satisfies strong zero-variance principle

On the other hand, if the states $\psi_i(\mathbf{R})$ are not closed under $\hat{H}$ then the values of the optimized parameters using nonsymmetric H_{ij} depend on the MC sample. In that case, the best one can do is to find a linear combination that minimizes $\langle\Delta|\Delta\rangle$:

$$\hat{H}\sum_{j=0}^{N_p} p_j|\psi_j\rangle \quad = \quad E\sum_{j=0}^{N_p} p_j|\psi_j\rangle + |\Delta\rangle$$

Now

$$\int d\mathbf{R}\ \langle\psi_i|\mathbf{R}\rangle\langle\mathbf{R}|\Delta\rangle \quad = \quad \langle\psi_i|\Delta\rangle = 0$$

$$\text{but} \quad \sum_{n=1}^{N_{\rm MC}} \langle\psi_i|\mathbf{R}_n\rangle\langle\mathbf{R}_n|\Delta\rangle \quad \neq \quad 0$$

So, if one does the complete sum over states (the integral) one gets the usual generalized eigenvalue equation but summing over a finite set of MC points yields parameters $\{p\}$ that depend on the sampled points. In practice, the nonsymmetric H reduces fluctuations in parameter values by 1-2 orders of magnitude (reduces computer time by 2-4 orders of magnitude) when one has good wavefunctions.

Convergence of energy with symmetric and nonsymmetric Hamiltonians

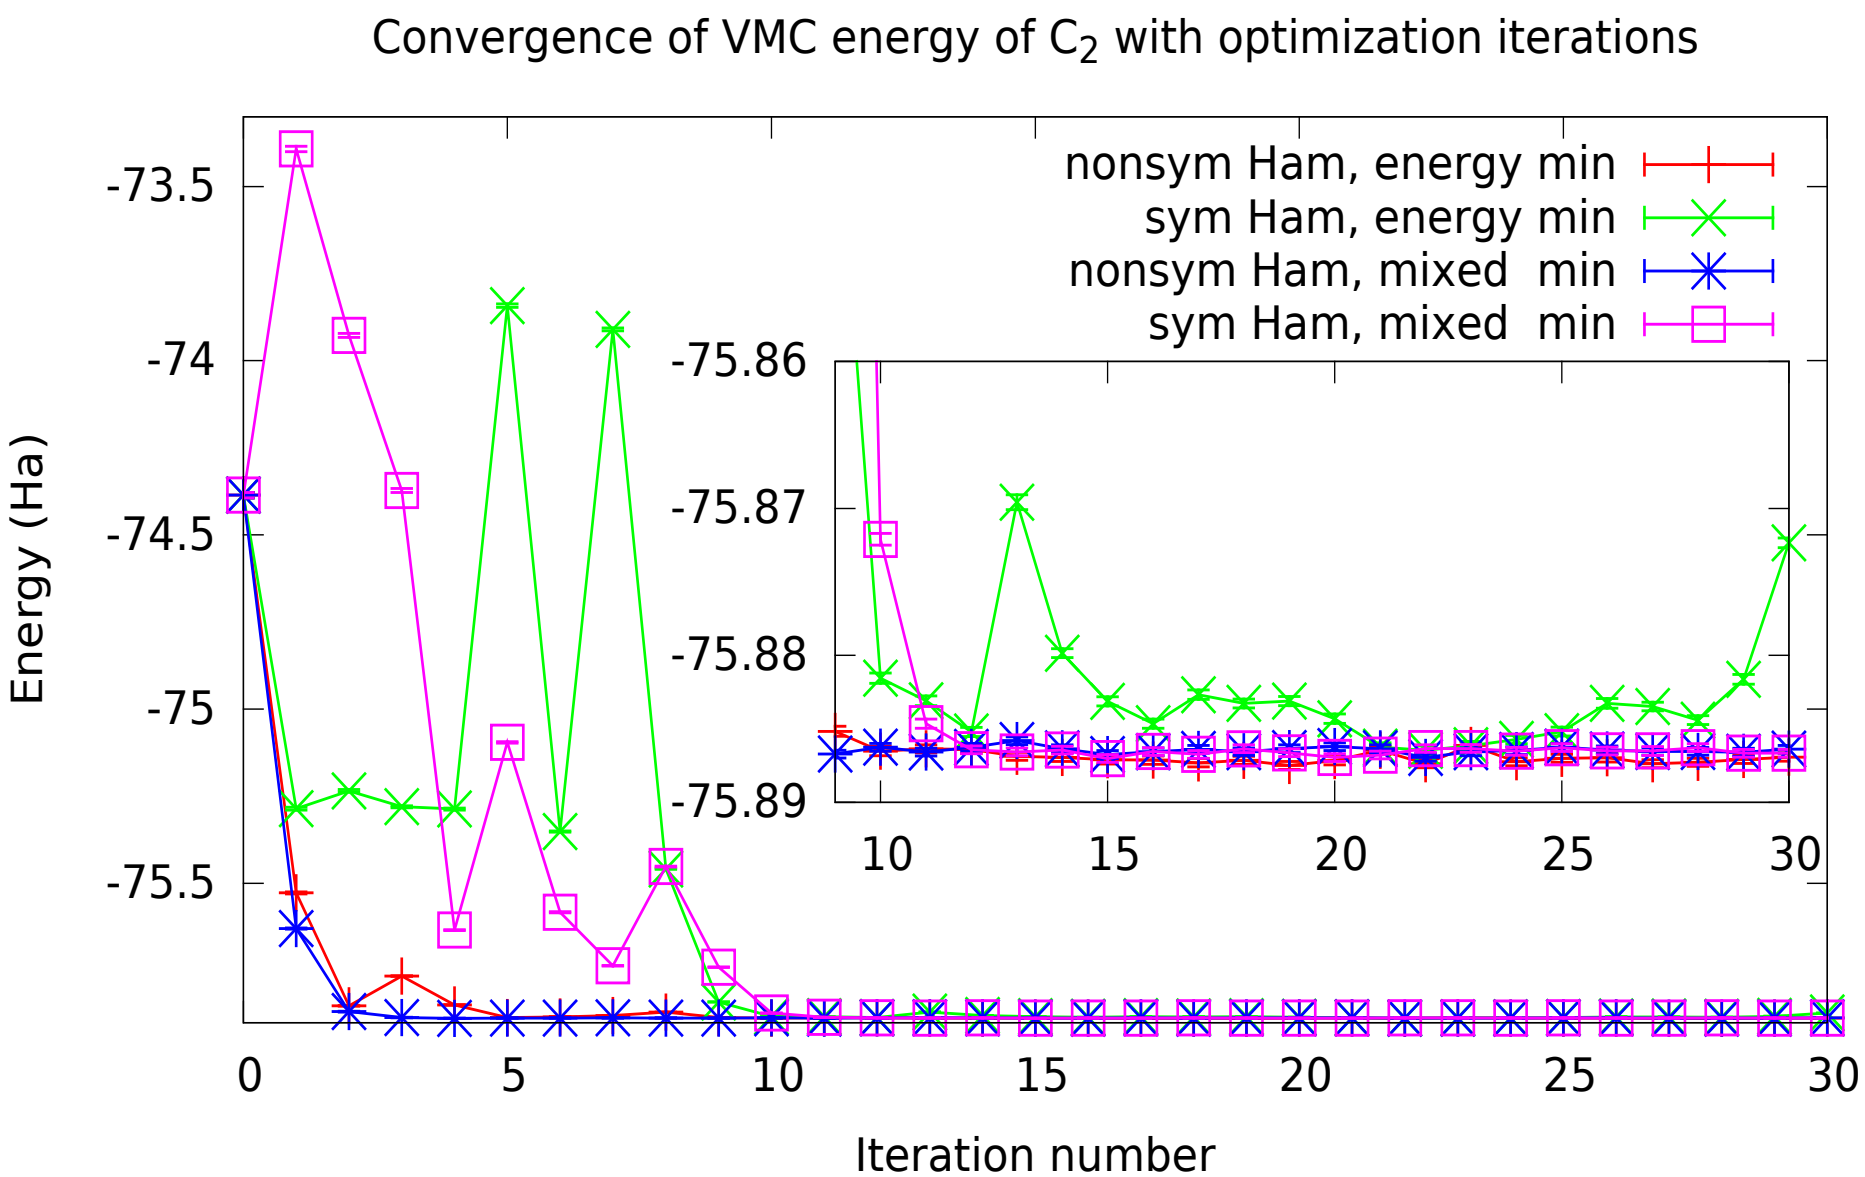

Convergence of energy with symmetric and nonsymmetric Hamiltonians

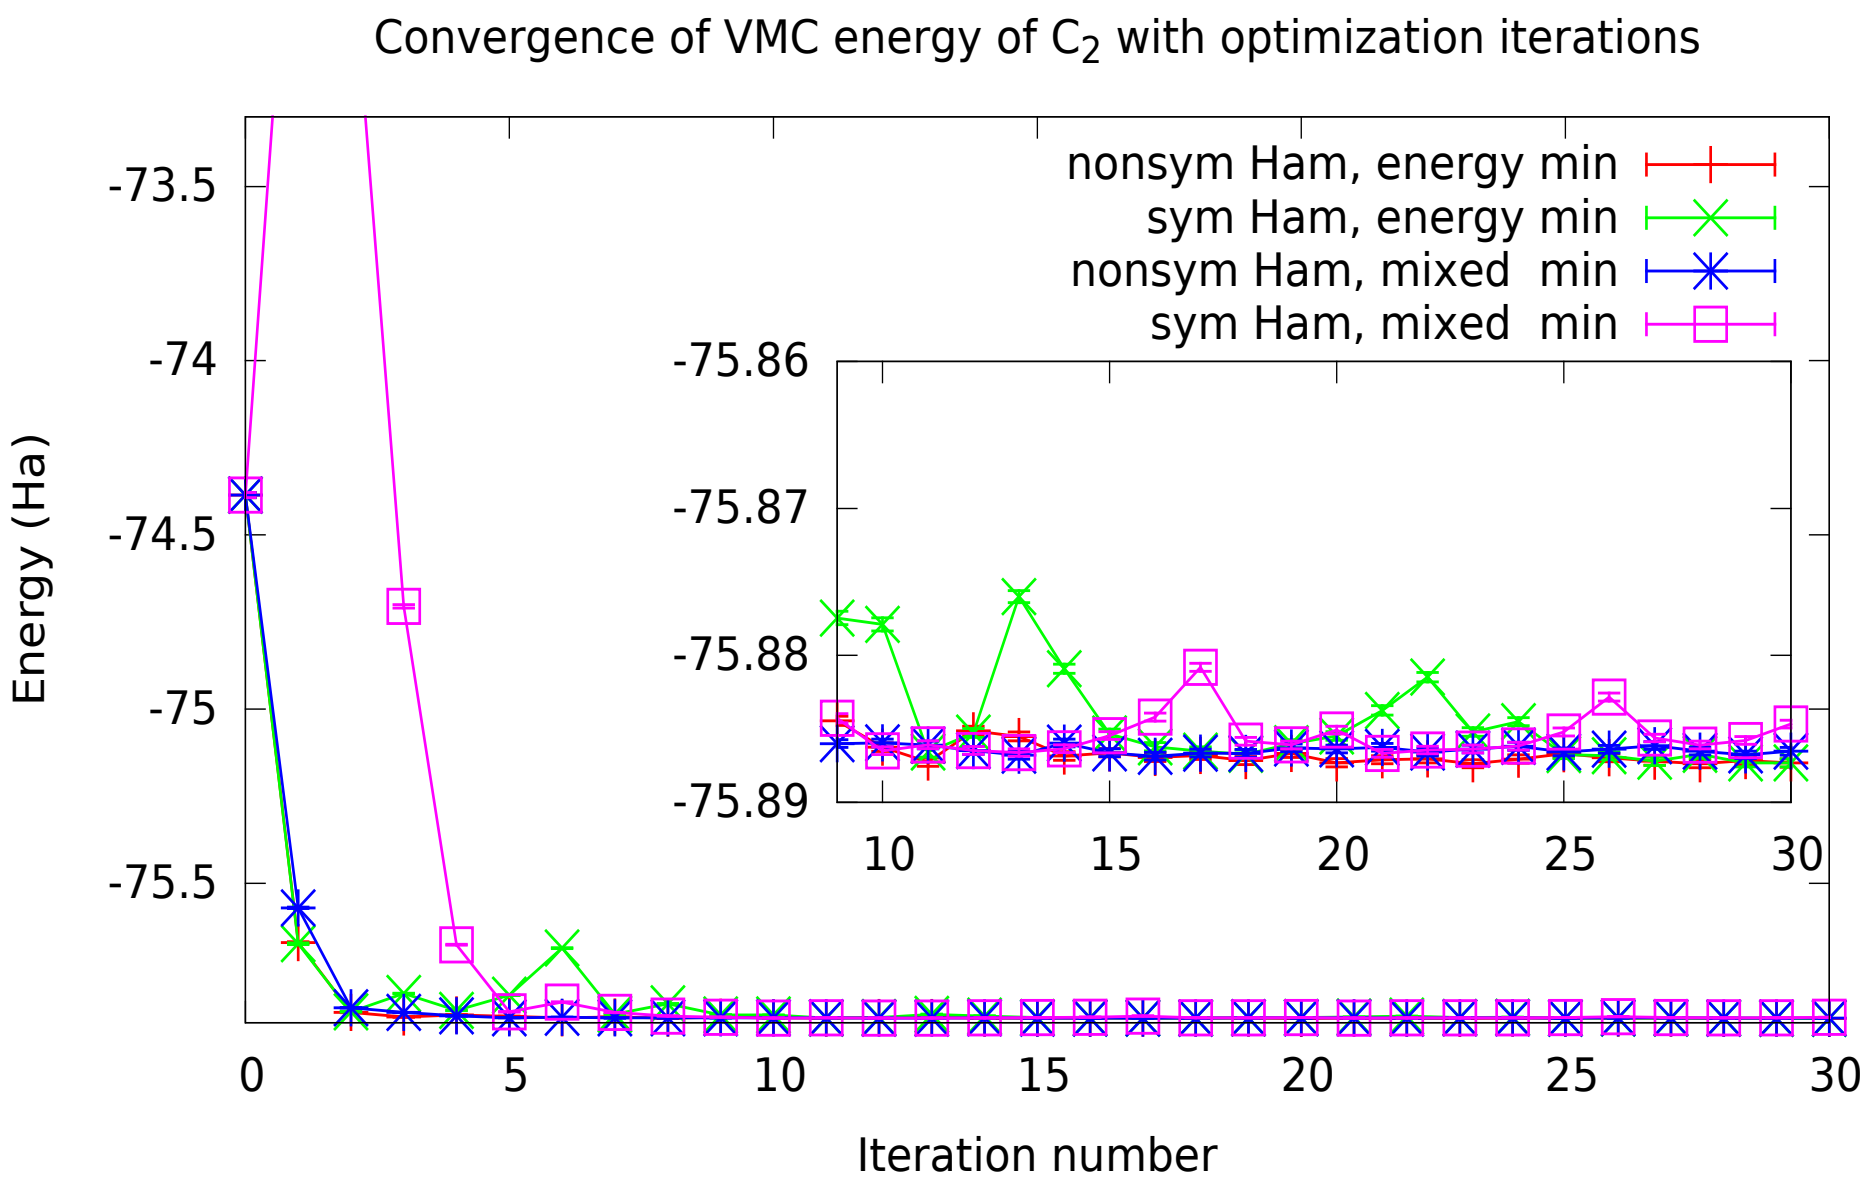

Linear method for nonlinear parameters

Toulouse, CJU, JCP (2007,2008); CJU et al., PRL, **87**, 043401 (2007).

Make linear-order Taylor expansion of Ψ (use $\Psi_i = \partial\Psi/\partial p_i$ as basis):

$$\Psi_{\rm lin} = \Psi_0 + \sum_{i=0}^{N_{\rm parm}} \Delta p_i \, \Psi_i, \qquad \text{(Normalization: } \Delta p_0 = 1\text{)}$$

$\Psi_0 \equiv \Psi(\mathbf{p}_0, \mathbf{R})$ = current wave function
$\Psi_{\rm lin}$ = next linearized wave function
Ψ_i = derivative of Ψ at $\mathbf{p}_0$, wrt i^{th} parameter.
No unique way to obtain new nonlinear parameters.
The simplest procedure: is $p_i^{\rm new} = p_i + \Delta p_i$. Will not work in general. What can one do?
More complicated procedure: fit wave function form to optimal linear combination.
Simpler, yet efficient approach, freedom of norm to make linear approximation better

$$\begin{aligned} \bar{\Psi}(\mathbf{p}, \mathbf{R}) &= N(\mathbf{p}) \, \Psi(\mathbf{p}, \mathbf{R}), \quad N(\mathbf{p}_0) = 1 \\ \bar{\Psi}_i(\mathbf{p}_0, \mathbf{R}) &= \Psi_i(\mathbf{p}_0, \mathbf{R}) + N_i(\mathbf{p}_0)\Psi(\mathbf{p}_0, \mathbf{R}) \end{aligned}$$

Note, $N_i = 0$ for linear parameters by definition.

Change of normalization

Toulouse, CJU, JCP (2007,2008); CJU et al., PRL, **87**, 043401 (2007).

$$\bar{\Psi}(\mathbf{p},\mathbf{R}) = N(\mathbf{p})\,\Psi(\mathbf{p},\mathbf{R}), \quad N(\mathbf{p}_0)=1$$
$$\bar{\Psi}_i = \Psi_i + N_i \Psi_0$$

$$\Psi = \Psi_0 + \sum_{i=1}^{N_{\rm parm}} \delta p_i \Psi_i$$

$$\bar{\Psi} = \Psi_0 + \sum_{i=1}^{N_{\rm parm}} \delta \bar{p}_i \bar{\Psi}_i = \left(1 + \sum_{i=1}^{N_{\rm parm}} N_i \delta \bar{p}_i\right) \Psi_0 . + \sum_{i=1}^{N_{\rm parm}} \delta \bar{p}_i \Psi_i$$

Since Ψ and $\bar{\Psi}$ are the optimal linear combin., they are the same aside from normalization

$$\delta p_i = \frac{\delta \bar{p}_i}{1 + \sum_{i=1}^{N_{\rm parm}} N_i \delta \bar{p}_i} \quad \Longrightarrow \quad \delta \bar{p}_i = \frac{\delta p_i}{1 - \sum_{i=1}^{N_{\rm parm}} N_i \delta p_i}. \qquad (1)$$

One can get $\delta \bar{p}_i$ directly from solving the eigenvalue problem in the renormalized basis or get δp_i from eigenvalue problem in the original basis and use the above transformation. In either case, use $\delta \bar{p}_i$ to update the parameters, $p_i^{\rm new} = p_i + \delta \bar{p}_i$.

The denominator in Eq. 1 can be +ve, -ve or zero! So, predicted parameter changes can change sign depending on normalization!! Of course for the linear parameters it does not matter since all it implies is a change in the normalization of new wavefunction, but for nonlinear parameters the change is crucial! If all parm. linear, $\delta \bar{p}_i = \delta p_i$, since all $N_i = 0$.

General semiorthogonalization

How to choose N_i ?

Toulouse, CJU, JCP (2007,2008); CJU et al., PRL, **87**, 043401 (2007).

$$\begin{aligned} \bar{\Psi}(\mathbf{p},\mathbf{R}) &= N(\mathbf{p})\,\Psi(\mathbf{p},\mathbf{R}), \quad N(\mathbf{p}_0)=1 \\ \bar{\Psi}_i &= \Psi_i + N_i\Psi_0 \end{aligned}$$

Choose N_i such that the derivatives are orthogonal to a linear combination of Ψ_0 and $\Psi_{\rm lin}$.

$$\left\langle \xi\frac{\Psi_0}{|\Psi_0|} + s(1-\xi)\frac{\Psi_{\rm lin}}{|\Psi_{\rm lin}|}\,\middle|\, \Psi_i + N_i\Psi_0 \right\rangle = 0$$

Solving for N_i we get $\left[s=1(-1) \text{ if } \langle\Psi_0|\Psi_{\rm lin}\rangle = 1+\sum_j S_{0j}\Delta p_j > 0 (<0)\right]$,

$$N_i = -\frac{\xi D S_{0i} + s(1-\xi)(S_{0i}+\sum_j S_{ij}\Delta p_j)}{\xi D + s(1-\xi)(1+\sum_j S_{0j}\Delta p_j)}$$

where $$D = \frac{|\Psi_{\rm lin}|}{|\Psi_0|} = \left(1 + 2\sum_j S_{0j}\Delta p_j + \sum_{i,j} S_{ij}\Delta p_i \Delta p_j\right)^{1/2}$$

Semiorthogonalization in the linear method

Comparison of semiorthogonalizations with xi = 1, 0.5, 0 versus no semiorthogonalization

Ψ_0 the initial wavefn.

Ψ_i^ξ derivative of the wavefn. wrt parameter p_i

$\Delta\Psi_i^\xi$ the change in the wavefn.

Ψ_{lin}^ξ the linear wavefn.

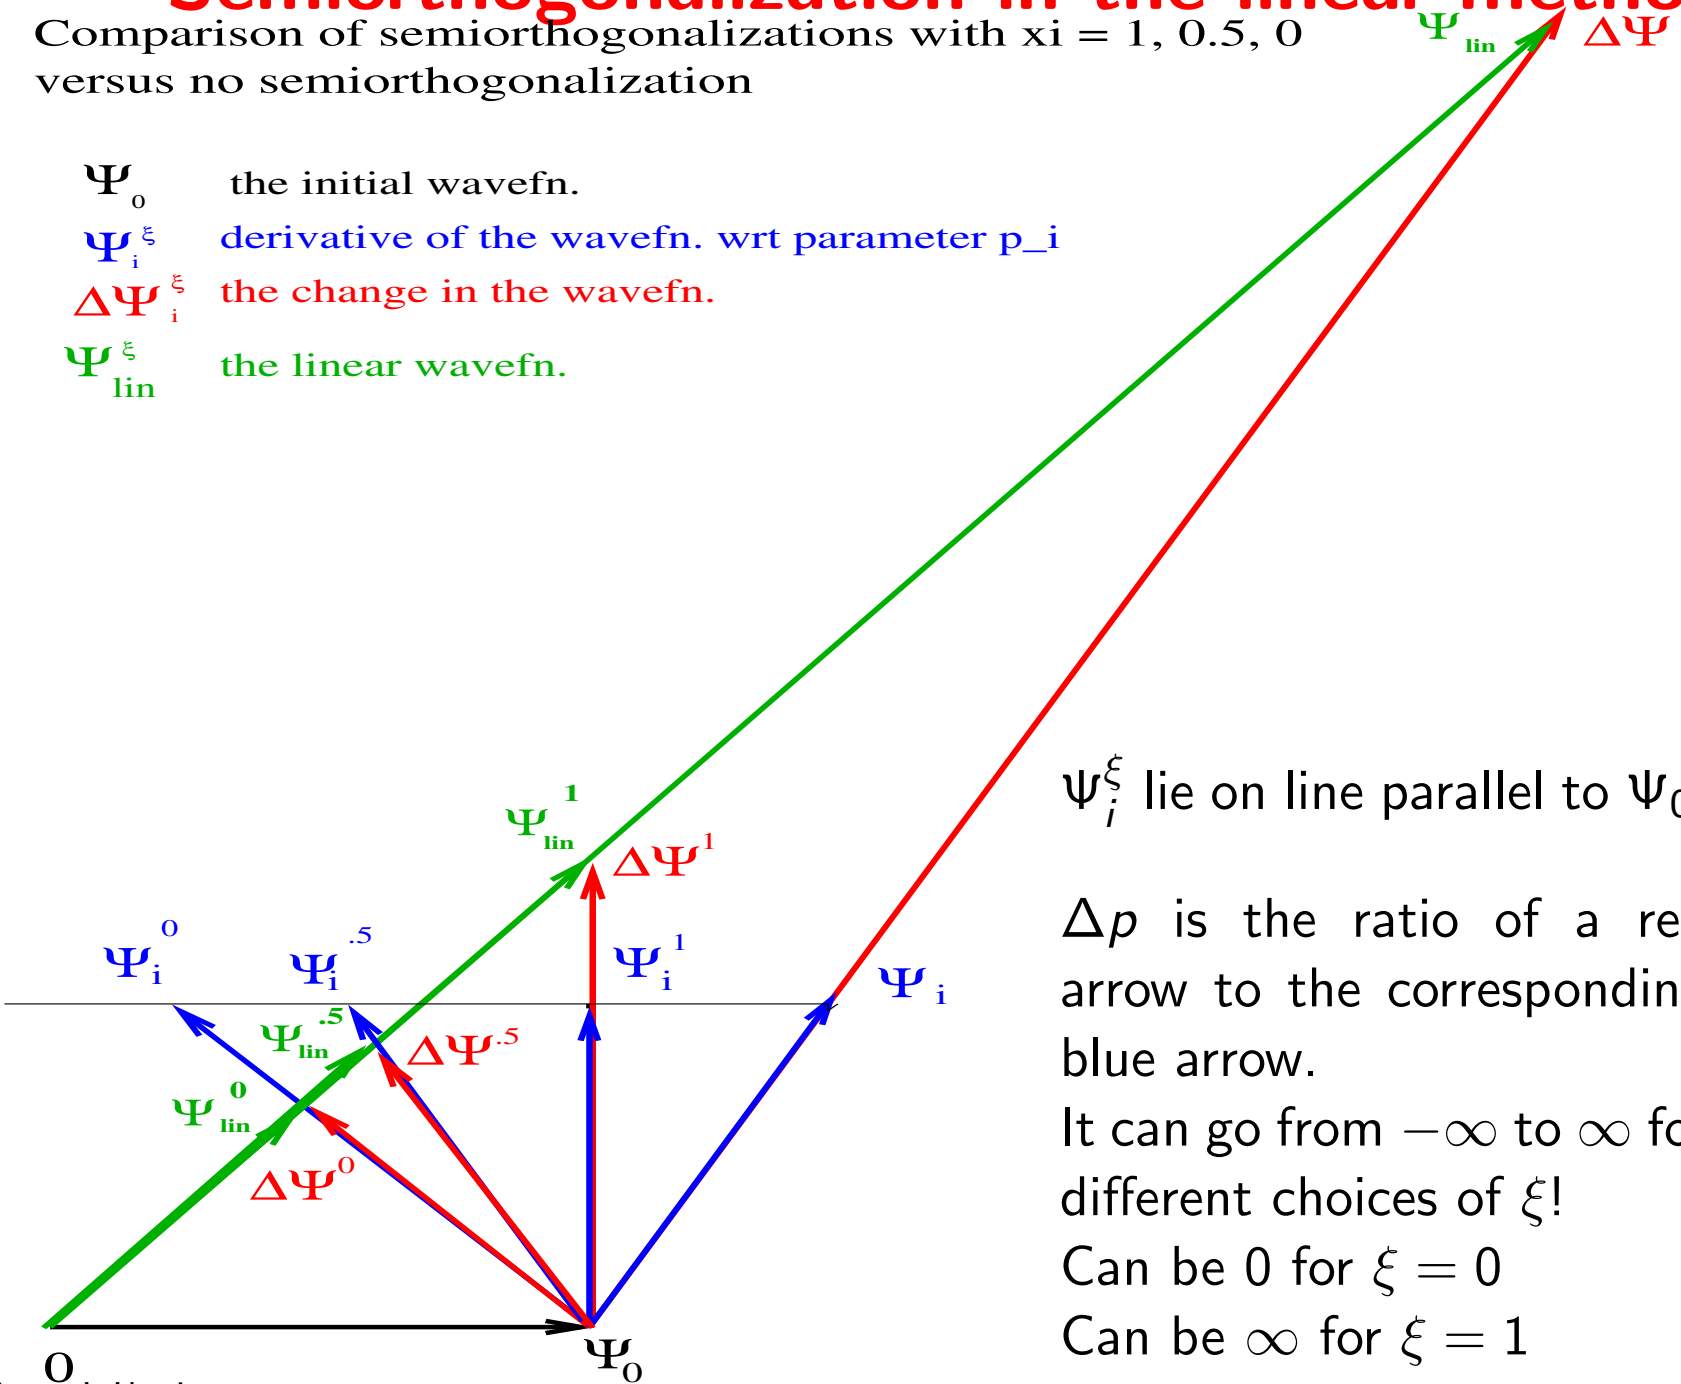

Ψ_i^ξ lie on line parallel to Ψ_0.

Δp is the ratio of a red arrow to the corresponding blue arrow.

It can go from $-\infty$ to ∞ for different choices of ξ!

Can be 0 for $\xi = 0$

Can be ∞ for $\xi = 1$

Semiorthogonalization in the linear method

Ψ_0 is the initial wave function, Ψ_i^ζ is the derivative of the wave function wrt parameter p_i for ζ. If superscript ζ is omitted that denotes that no semiorthogonalization is done. Then

$$\Psi_{\rm lin} \;=\; \Psi_0 + \sum_{i=1}^{N_{\rm parm}} \Delta\Psi_i^\zeta = \Psi_0 + \sum_{i=1}^{N_{\rm parm}} \Delta p_i^\zeta \Psi_i^\zeta, \quad \Delta p_i^\zeta = \frac{\Delta\Psi_i^\zeta}{\Psi_i^\zeta}$$

Note that $||\Delta\Psi^\zeta||$ is smallest for $\zeta = 1$ and that $||\Psi_{\rm lin}^{0.5}|| = ||\Psi_0||$.
Also note that when there is just 1 parameter (can be generalized to > 1):

1. In the limit that $\Psi_{\rm lin} \parallel \Psi_i$, $\Delta p_i = \pm\infty$
2. In the limit that $\Psi_{\rm lin} \perp \Psi_0$, $\Delta p_i^1 = \pm\infty$ because $\Delta\Psi^1 = \infty$, and, $\Delta p_i^0 = 0$ because $\Psi_i^0 = \infty$
3. $\Delta p_i^{0.5}$ is always finite

Note that Δp_i^ζ decreases as ζ decreases from 1 to 0. In Fig. 1, Δp_i is > 1 for $\zeta = 1$, and, < 1 for $\zeta = 0.5, 0$.
Also note that in Fig. 1 if we rotate $\Psi_{\rm lin}$ such that $\frac{\nabla\Psi\cdot\Psi_0}{||\nabla\Psi||||\Psi_0||} > \frac{\Psi_{\rm lin}\cdot\Psi_0}{||\Psi_{\rm lin}||||\Psi_0||}$ then Δp_i has the opposite sign as Δp_i^ζ!

Variance Minimization via Linear method

Toulouse, CJU, J. Chem. Phys., **128**, 174101 (2008)

Can one use the linear method to optimize the variance?

Variance Minimization via Linear method

Toulouse, CJU, J. Chem. Phys., **128**, 174101 (2008)

Can one use the linear method to optimize the variance?
Suppose we have some quadratic model of the energy variance to minimize

$$V_{\rm min} \quad = \quad \min_{\Delta \mathbf{p}} \left\{ V_0 + \mathbf{g}_V^{\rm T} \cdot \Delta \mathbf{p} + \frac{1}{2} \Delta \mathbf{p}^{\rm T} \cdot \mathbf{h}_V \cdot \Delta \mathbf{p} \right\}, \qquad (2)$$

where $V_0 = \langle \overline{\Psi}_0 | (\hat{H} - E_0)^2 | \overline{\Psi}_0 \rangle$ is the energy variance of the current wave function $|\overline{\Psi}_0\rangle$, $\mathbf{g}_V$ is the gradient of the energy variance with components $g_{V,i} = 2\langle \overline{\Psi}_i | (\hat{H} - E_0)^2 | \overline{\Psi}_0 \rangle$ and $\mathbf{h}_V$ is some approximation to the Hessian matrix of the energy variance. Then, one could instead minimize the following rational quadratic model (*augmented hessian method*)

$$V_{\rm min} \quad = \quad \min_{\Delta \mathbf{p}} \frac{\begin{pmatrix} 1 & \Delta \mathbf{p}^{\rm T} \end{pmatrix} \begin{pmatrix} V_0 & \mathbf{g}_V^{\rm T}/2 \\ \mathbf{g}_V/2 & \mathbf{h}_V/2 + V_0 \overline{\mathbf{S}} \end{pmatrix} \begin{pmatrix} 1 \\ \Delta \mathbf{p} \end{pmatrix}}{\begin{pmatrix} 1 & \Delta \mathbf{p}^{\rm T} \end{pmatrix} \begin{pmatrix} 1 & \mathbf{0}^{\rm T} \\ \mathbf{0} & \overline{\mathbf{S}} \end{pmatrix} \begin{pmatrix} 1 \\ \Delta \mathbf{p} \end{pmatrix}},$$

which agrees with the quadratic model in Eq. (2) up to second order in $\Delta \mathbf{p}$, and which leads to the following generalized eigenvalue equation

$$\begin{pmatrix} V_0 & \mathbf{g}_V^{\rm T}/2 \\ \mathbf{g}_V/2 & \mathbf{h}_V/2 + V_0 \overline{\mathbf{S}} \end{pmatrix} \begin{pmatrix} 1 \\ \Delta \mathbf{p} \end{pmatrix} = V_{\rm min} \begin{pmatrix} 1 & \mathbf{0}^{\rm T} \\ \mathbf{0} & \overline{\mathbf{S}} \end{pmatrix} \begin{pmatrix} 1 \\ \Delta \mathbf{p} \end{pmatrix}.$$

Hence, we can use linear method to optimize a linear combination of energy and variance!

Connection between Linear and Newton methods

Toulouse, CJU, J. Chem. Phys., **128**, 174101 (2008)

In semiorthogonal basis with $\xi = 1$, linear eqs. are:

$$\begin{pmatrix} E_0 & \mathbf{g}^T/2 \\ \mathbf{g}/2 & \overline{\mathbf{H}} \end{pmatrix} \begin{pmatrix} 1 \\ \Delta\mathbf{p} \end{pmatrix} = E_{\rm lin} \begin{pmatrix} 1 & \mathbf{0}^T \\ \mathbf{0} & \overline{\mathbf{S}} \end{pmatrix} \begin{pmatrix} 1 \\ \Delta\mathbf{p} \end{pmatrix}, \tag{3}$$

Defining, $\Delta E = E_{\rm lin} - E_0 \le 0$, the 1^{st} and 2^{nd} eqs. are:

$$2\Delta E = \mathbf{g}^T \cdot \Delta\mathbf{p}, \qquad 1^{st} \text{ eq.} \tag{4}$$

$$\frac{\mathbf{g}}{2} + \overline{\mathbf{H}}\Delta\mathbf{p} = E_{\rm lin}\overline{\mathbf{S}}\Delta\mathbf{p} \qquad 2^{nd} \text{ eq.} \tag{5}$$

$$\text{i.e.,} \quad 2\left(\overline{\mathbf{H}} - E_{\rm lin}\overline{\mathbf{S}}\right)\Delta\mathbf{p} = -\mathbf{g}, \tag{6}$$

This can be viewed as the Newton method with an approximate hessian, $\mathbf{h} = 2(\overline{\mathbf{H}} - E_{\rm lin}\overline{\mathbf{S}})$ which is nonnegative definite. (It has all nonnegative eigenvalues since we are subtracting out the lowest eigenvalue.) This also means that the linear method can be stabilized in much the same way as the Newton method.

Note that $2(\overline{\mathbf{H}} - E_{\rm lin}\overline{\mathbf{S}}) = 2(\overline{\mathbf{H}} - E_0\mathbf{S} - \Delta E\,\overline{\mathbf{S}})$ and $2(\overline{\mathbf{H}} - E_0\overline{\mathbf{S}})$ is the approximate hessian of Sorella's stochastic reconfiguration with approximate hessian (SRH) method (which converges more slowly that our linear and Newton methods). The present method provides an automatic stabilization of the SRH method by a positive definite matrix $-\Delta E\,\overline{\mathbf{S}}$ making the hessian nonnegative definite.

Perturbation Theory in an arbitrary nonorthogonal basis

Toulouse, CJU, J. Chem. Phys., **126**, 084102 (2007)

.

Given a Hamiltonian $\hat{H}$ and an arbitrary nonorthogonal basis, $\{|\Psi_i\rangle\}$, use perturbation theory to get approximate eigenstates of $\hat{H}$.

Define dual basis: $\langle\bar{\Psi}_i|\Psi_j\rangle = \delta_{ij}$ and zeroth order Hamiltonian, $\hat{H}^{(0)}$:

$$\langle\bar{\Psi}_i| \quad = \quad \sum_{j=0}^{N_{\rm opt}} (\mathbf{S}^{-1})_{ij}\langle\Psi_j|, \qquad \hat{H}^{(0)} = \sum_{i=0}^{N_{\rm opt}} E_i|\Psi_i\rangle\langle\bar{\Psi}_i|$$

First order perturbation correction is

$$|\Psi^{(1)}\rangle \quad = \quad -\sum_{i=1}^{N_{\rm opt}} |\Psi_i\rangle \sum_{j=1}^{N_{\rm opt}} (\mathbf{S}^{-1})_{ij} \frac{\langle\Psi_j|\hat{H}|\Psi_0\rangle}{E_i - E_0}$$

Want $\hat{H}$ and $\hat{H}^{(0)}$ close – choose E_i so $\hat{H}$ and $\hat{H}^{(0)}$ have same diagonals

$$E_i \quad = \quad \frac{\langle\Psi_i|\hat{H}|\Psi_i\rangle}{\langle\Psi_i|\Psi_i\rangle}$$

If the E_i are evaluated without the Jastrow factor then this is the same as the perturbative eff. fluct. pot. (EFP) method of Schautz and Filippi and Scemama and Filippi.

Stabilization

If far from the minimum, or, $N_{\rm MC}$, is small, then the Hessian, $\bar{E}_{ij}$, need not be positive definite (whereas variance-minimization Levenberg-Marquardt $\bar{E}_{ij}$ is positive definite).

Even for positive definite $\bar{E}_{ij}$, the new parameter values may make the wave function worse if quadratic approximation is not good.

Add $a_{\rm diag}$ to the diagonal elements of the Hessian. This shifts the eigenvalues by the added constant. As $a_{\rm diag}$ is increased, the proposed parameter changes become smaller and rotate from the Newtonian direction to the steepest descent direction, but in practice $a_{\rm diag}$ is tiny.

The linear method and the perturbative method can be approximately recast into the Newton method. Consequently we can use the same idea for the linear and perturbative methods too.

Stabilization with Correlated Sampling

Each method has a parameter $a_{\rm diag}$ that automatically adjusts to make the method totally stable:

1. Do a MC run to compute the gradient and the Hessian (or overlap and Hamiltonian).
2. Using the above gradient and Hessian (or overlap and Hamiltonian), use 3 different values of $a_{\rm diag}$ to predict 3 different sets of updated parameters.
3. Do a short correlated sampling run for the 3 different wave functions to compute the energy differences for the 3 wave functions more accurately than the energies themselves.
4. Fit a parabola through the 3 energies to find the optimal $a_{\rm diag}$.
5. Use this optimal $a_{\rm diag}$ to predict a new wave function, using the gradient and Hessian computed in step 1.
6. Loop back

Comparison of Newton, linear and perturbative methods

Programming effort and cost per iteration:

1. Newton method requires $\psi,\ \ \psi_i,\ \ \psi_{ij},\ \ \hat{H}\psi,\ \ \hat{H}\psi_i$.
2. Linear method requires $\psi,\ \ \psi_i,\ \ \hat{H}\psi,\ \ \hat{H}\psi_i$.
3. Perturbative method requires $\psi,\ \ \psi_i,\ \ \hat{H}\psi,\ \ \hat{H}\psi_i$.
   Perturbative method with approx. denom., and, SR require $\psi,\ \ \psi_i,\ \ \hat{H}\psi$.

Convergence with number of iterations:

1. Newton and linear methods converge in 2-10 iterations for all parameters (CSF, orbital and Jastrow), but sometimes orbitals and exponents can take much longer.
2. Perturbative method converges in 2-10 iterations for CSF and orbital parameters but is very slow for Jastrow because eigenvalues of Hessian for Jastrow span 10-12 orders of magnitude. (Perturbative method can be viewed as Newton with crude Jastrow.)

Things to note

Eigenvalues of $\bar{E}_{ij}$ for Jastrow parameters typically span 10-12 orders of magnitude. So steepest descent would be horribly slow to converge!

Linear and Newton methods can be used for all parameters, even basis-set exponents.

Take Home Message:

Any method that attempts to minimize the energy, by minimizing the energy evaluated on a set of MC points, will require a very large sample and be highly inefficient.
Each of the 3 methods presented above avoids doing this.

Optimization of linear combination of energy and variance

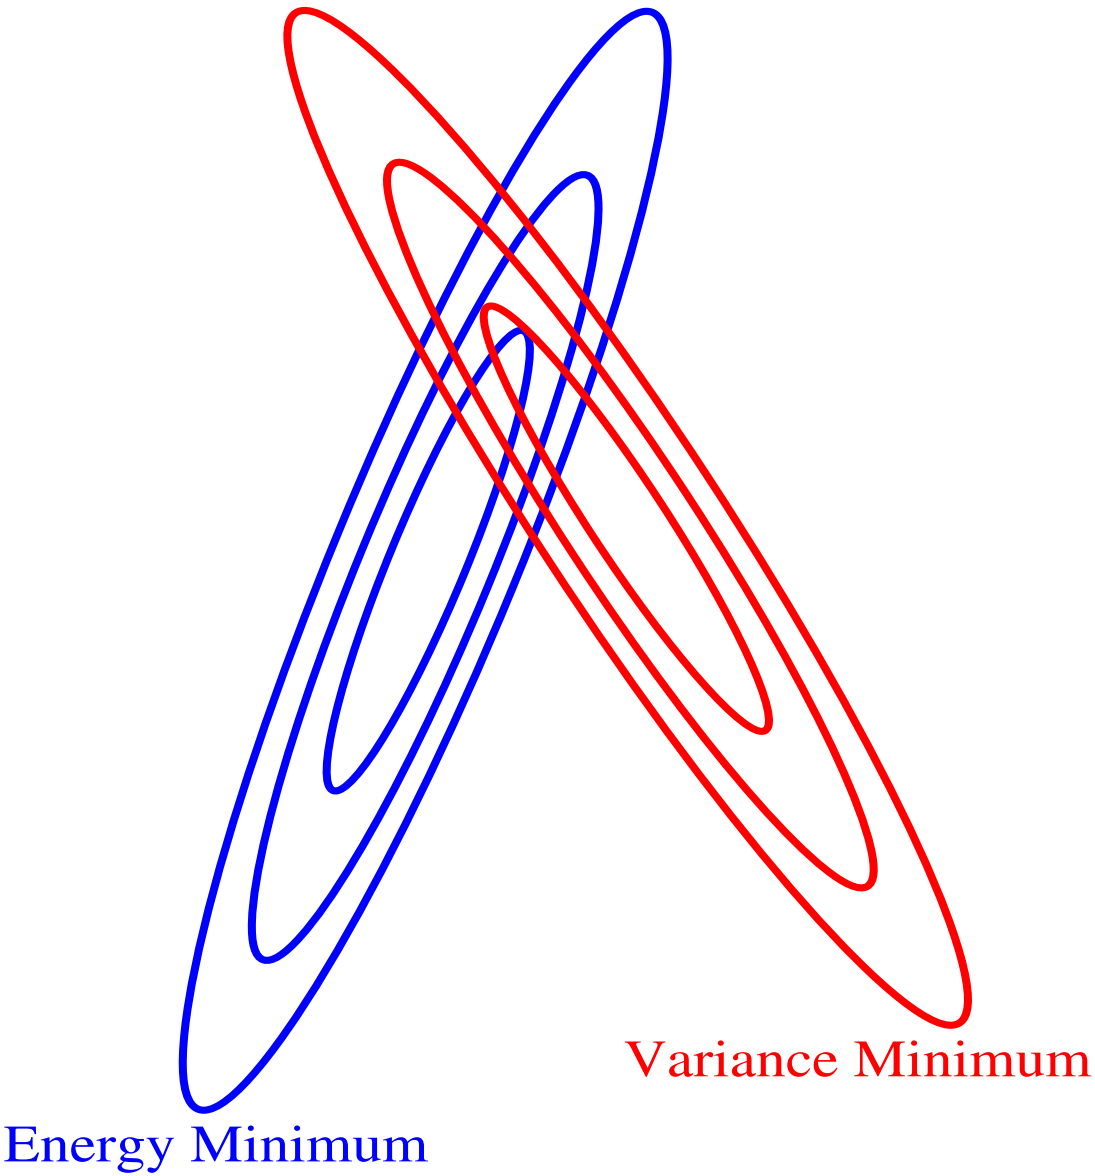

1. Can reduce the variance, without sacrificing appreciably the energy, by minimizing a linear combination, particularly since the ratio of hard to soft directions is 11 orders of magnitude.
2. Easy to do – obvious for Newton. Not obvious, but easy to do for linear method as shown above.
3. Measure of efficiency of the wave function is $\sigma^2 T_{\rm corr}$.

Lectures 4 and 5

Projector MC (PMC)
Path-integral MC (PIMC)
Reptation MC

Cyrus J. Umrigar

The problem

We wish to find the lowest energy eigenstate(s) of a (Hamiltonian) matrix.

If the number of basis states is sufficiently small that one can store a vector (say $< 10^{10}$), then one can use a deterministic iterative method, such as the power method or the Lanczos method.

Quantum Monte Carlo: If the space is larger than this, even infinite, one can use a stochastic implementation of the power method. At any instant in time only a random sample of the vector is stored in computer memory, and the solution is given by the time-average.

Definitions

Given a complete or incomplete basis: $\{|\phi_i\rangle\}$, either discrete or continuous

$$\text{Exact} \quad |\Psi_0\rangle = \sum_i e_i|\phi_i\rangle, \quad \text{where,} \quad e_i = \langle\phi_i|\Psi_0\rangle$$

$$\text{Trial} \quad |\Psi_T\rangle = \sum_i t_i|\phi_i\rangle, \quad \text{where,} \quad t_i = \langle\phi_i|\Psi_T\rangle$$

$$\text{Guiding} \quad |\Psi_G\rangle = \sum_i g_i|\phi_i\rangle, \quad \text{where,} \quad g_i = \langle\phi_i|\Psi_G\rangle$$

(If basis incomplete then "exact" means "exact in that basis".)

Ψ_T used to calculate variational and mixed estimators of operators $\hat{A}$, i.e., $\langle\Psi_{\rm T}|\hat{A}|\Psi_{\rm T}\rangle/\langle\Psi_{\rm T}|\Psi_{\rm T}\rangle$, $\langle\Psi_{\rm T}|\hat{A}|\Psi_0\rangle/\langle\Psi_{\rm T}|\Psi_0\rangle$

Ψ_G used to alter the probability density sampled, i.e., $\Psi_{\rm G}^2$ in VMC, $\Psi_{\rm G}\Psi_0$ in PMC. Affects only the statistical error of VMC and PMC methods.

$\Psi_{\rm G}$ must be such that $g_i \neq 0$ if $e_i \neq 0$. If $\Psi_{\rm T}$ also satisfies this condition then $\Psi_{\rm G}$ can be chosen to be $\Psi_{\rm T}$. To simplify expressions, we use $\Psi_{\rm G} = \Psi_{\rm T}$ or $\Psi_{\rm G} = 1$ in what follows.

Variational MC

$$E_V = \frac{\langle \Psi_{\rm T}|\hat{H}|\Psi_{\rm T}\rangle}{\langle \Psi_{\rm T}|\Psi_{\rm T}\rangle} = \frac{\sum_{ij}^{N_{\rm st}} \langle \Psi_{\rm T}|\phi_i\rangle \ \langle \phi_i|\hat{H}|\phi_j\rangle \ \langle \phi_j|\Psi_{\rm T}\rangle}{\sum_{i}^{N_{\rm st}} \langle \Psi_{\rm T}|\phi_i\rangle \ \langle \phi_i|\Psi_{\rm T}\rangle}$$

$$= \frac{\sum_{ij}^{N_{\rm st}} t_i H_{ij} t_j}{\sum_{i}^{N_{\rm st}} t_i^2} = \frac{\sum_{i}^{N_{\rm st}} \ g_i^2 \ \frac{t_i^2}{g_i^2} \ \frac{\sum_j^{N_{\rm st}} H_{ij} t_j}{t_i}}{\sum_{i}^{N_{\rm st}} \ g_i^2 \frac{t_i^2}{g_i^2}}$$

$$= \frac{\sum_{i}^{N_{\rm st}} \ g_i^2 \ \frac{t_i^2}{g_i^2} \ E_{\rm L}(i)}{\sum_{i}^{N_{\rm st}} \ g_i^2 \ \frac{t_i^2}{g_i^2}} = \frac{\left[\sum_{i}^{N_{\rm MC}} \ \frac{t_i^2}{g_i^2} \ E_{\rm L}(i)\right]_{\Psi_{\rm G}^2}}{\left[\sum_{i}^{N_{\rm MC}} \ \frac{t_i^2}{g_i^2}\right]_{\Psi_{\rm G}^2}} \equiv \frac{\left\langle \frac{t_i^2}{g_i^2} \ E_{\rm L}(i)\right\rangle_{\Psi_{\rm G}^2}}{\left\langle \frac{t_i^2}{g_i^2}\right\rangle_{\Psi_{\rm G}^2}}$$

Sample probability density function $\frac{g_i^2}{\sum_k^{N_{\rm st}} g_k^2}$ using Metropolis-Hastings.

Value depends only on $\Psi_{\rm T}$. Statistical error depend on $\Psi_{\rm T}$ and $\Psi_{\rm G}$.

Energy bias and statistical error vanish as $\Psi_{\rm T} \to \Psi_0$.

For fixed Ψ_T, $\Psi_G = \Psi_T$ does not minimize statistical fluctuations!

Projector MC

Pure and Mixed estimators for energy are equal: $E_0 = \dfrac{\langle \Psi_0 | \hat{H} | \Psi_0 \rangle}{\langle \Psi_0 | \Psi_0 \rangle} = \dfrac{\langle \Psi_0 | \hat{H} | \Psi_{\rm T} \rangle}{\langle \Psi_0 | \Psi_{\rm T} \rangle}$

Projector: $|\Psi_0\rangle = \hat{P}(\infty)\,|\Psi_{\rm T}\rangle = \lim_{n\to\infty} \hat{P}^n(\tau)\,|\Psi_{\rm T}\rangle$

$$E_0 = \frac{\langle \Psi_0 | \hat{H} | \Psi_{\rm T} \rangle}{\langle \Psi_0 | \Psi_{\rm T} \rangle} = \frac{\sum_{ij}^{N_{\rm st}} \langle \Psi_0 | \phi_i \rangle \, \langle \phi_i | \hat{H} | \phi_j \rangle \, \langle \phi_j | \Psi_{\rm T} \rangle}{\sum_{i}^{N_{\rm st}} \langle \Psi_0 | \phi_i \rangle \, \langle \phi_i | \Psi_{\rm T} \rangle}$$

$$= \frac{\sum_{ij}^{N_{\rm st}} e_i H_{ij} t_j}{\sum_{i}^{N_{\rm st}} e_i t_i} = \frac{\sum_{i}^{N_{\rm st}} e_i g_i \, \frac{t_i}{g_i} \, \frac{\sum_j^{N_{\rm st}} H_{ij} t_j}{t_i}}{\sum_{i}^{N_{\rm st}} e_i g_i \, \frac{t_i}{g_i}}$$

$$= \frac{\sum_{i}^{N_{\rm st}} e_i g_i \, \frac{t_i}{g_i} \, E_{\rm L}(i)}{\sum_{i}^{N_{\rm st}} e_i g_i \, \frac{t_i}{g_i}} = \frac{\left[\sum_{i}^{N_{\rm MC}} \frac{t_i}{g_i} \, E_{\rm L}(i)\right]_{\Psi_{\rm G}\Psi_0}}{\left[\sum_{i}^{N_{\rm MC}} \frac{t_i}{g_i}\right]_{\Psi_{\rm G}\Psi_0}} \equiv \frac{\left\langle \frac{t_i}{g_i} \, E_{\rm L}(i) \right\rangle_{\Psi_{\rm G}\Psi_0}}{\left\langle \frac{t_i}{g_i} \right\rangle_{\Psi_{\rm G}\Psi_0}}$$

Sample $e_i g_i / \sum_k^{N_{\rm st}} e_k g_k$ using projector.

For exact PMC, value indep. of $\Psi_{\rm T}$, $\Psi_{\rm G}$, statistical error depends on $\Psi_{\rm T}$, $\Psi_{\rm G}$.
(For FN-PMC, value and statistical error depends on $\Psi_{\rm T}, \Psi_{\rm G}$.)
(For FN-DMC, value depends on $\Psi_{\rm G}$, statistical error on $\Psi_{\rm T}, \Psi_{\rm G}$.)
Statistical error vanishes as $\Psi_{\rm T} \to \Psi_0$, since $E_{{\rm L},i} = E_0$, indep. of i.
For fixed Ψ_T, $\Psi_G = \Psi_T$ does not minimize statistical fluctuations!

Variational and Projector MC

$$E_V = \frac{\left[\sum_i^{N_{\rm MC}} \frac{t_i^2}{g_i} E_{\rm L}(i)\right]_{\Psi_{\rm G}}}{\left[\sum_i^{N_{\rm MC}} \frac{t_i^2}{g_i}\right]_{\Psi_{\rm G}}} = \frac{\left[\sum_i^{N_{\rm MC}} \frac{t_i^2}{g_i^2} E_{\rm L}(i)\right]_{\Psi_{\rm G}^2}}{\left[\sum_i^{N_{\rm MC}} \frac{t_i^2}{g_i^2}\right]_{\Psi_{\rm G}^2}}$$ (Value depends on $\Psi_{\rm T}$, error on $\Psi_{\rm T}, \Psi_{\rm G}$)

$$E_0 = \underbrace{\frac{\left[\sum_i^{N_{\rm MC}} t_i E_{\rm L}(i)\right]_{\Psi_0}}{\left[\sum_i^{N_{\rm MC}} t_i\right]_{\Psi_0}}}_{\text{"no imp. sampl."}} = \underbrace{\frac{\left[\sum_i^{N_{\rm MC}} \frac{t_i}{g_i} E_{\rm L}(i)\right]_{\Psi_{\rm G}\Psi_0}}{\left[\sum_i^{N_{\rm MC}} \frac{t_i}{g_i}\right]_{\Psi_{\rm G}\Psi_0}}}_{\text{"imp. sampl."}}$$ (Value exact†, error depends on $\Psi_{\rm T}, \Psi_{\rm G}$.)

$$E_{\rm L}(i) = \frac{\sum_j^{N_{\rm st}} H_{ij} t_j}{t_i}$$ (Note: $E_{\rm L}(i)$ is ∞ if $t_i = 0$, but, $t_i E_{\rm L}(i)$ is finite.)

In both VMC and PMC compute weighted average of the *configuration value of* $\hat{H}$ aka *local energy,* $E_{\rm L}(i)$, but from points sampled from different distributions.

This is practical for systems that are large enough to be interesting if

1. $t_i = \langle \phi_i | \Psi_{\rm T} \rangle$ can be evaluated in polynomial time, say N^3
2. the sum in $E_{\rm L}(i)$ can be done quickly, i.e., $\hat{H}$ is sparse (if space discrete) or semi-diagonal (if space continuous).

† In practice, usually necessary to make fixed-node approximation.

Importance Sampling in Projector Monte Carlo

If

$$\sum_j P_{ij} e_j \;\; = \;\; e_i$$

the similarity transformed matrix with elements $\tilde{P}_{ij} = \frac{g_i P_{ij}}{g_j}$ has eigenstate with elements $g_i e_i$:

$$\sum_j \tilde{P}_{ij} (g_j e_j) \;\; = \;\; \sum_j \left(\frac{g_i P_{ij}}{g_j} \right) (g_j e_j) \;\; = \;\; g_i e_i$$

Projector MC

Projector: $|\Psi_0\rangle = \hat{P}(\infty)\,|\Psi_{\rm T}\rangle = \lim_{n\to\infty} \hat{P}^n(\tau)\,|\Psi_{\rm T}\rangle$

Projector is any function of the Hamiltonian that maps the ground state eigenvalue of $\hat{H}$ to 1, and the highest eigenvalue of $\hat{H}$ to an absolute that is < 1 (preferably close to 0).

Exponential projector: $\hat{P} = e^{\tau(E_T\hat{\mathbf{1}}-\hat{H})}$

Linear projector: $\hat{P} = \hat{\mathbf{1}} + \tau(E_T\hat{\mathbf{1}} - \hat{H})$

Projector Monte Carlo Methods

The amplitudes of Ψ_0 in the chosen basis are obtained by using a "Projector", $\hat{P}$, that is a function of the Hamiltonian, $\hat{H}$, and has Ψ_0 as its dominant state.

Various Projector Monte Carlo Methods differ in:
a) form of the projector, and,
b) space in which the walk is done (single-particle basis and quantization).
(1^{st}-quantized $\equiv$ unsymmetrized basis, 2^{nd}-quantized $\equiv$ antisymmetrized basis.)

Method	Projector	SP Basis	Quantiz
Diffusion Monte Carlo	$e^{\tau(E_T\hat{\mathbf{1}}-\hat{H})}$	$\mathbf{r}$	1^{st}
GFMC (Kalos, Ceperley, Schmidt)	$\frac{1}{\hat{\mathbf{1}}-\tau(E_T\hat{\mathbf{1}}-\hat{H})}$	$\mathbf{r}$	1^{st}
LRDMC (Sorella, Casula)	$\hat{\mathbf{1}}+\tau(E_T\hat{\mathbf{1}}-\hat{H})$	$\mathbf{r}_i$	1^{st}
PMC/FCIQMC/SQMC	$\hat{\mathbf{1}}+\tau(E_T\hat{\mathbf{1}}-\hat{H})$	$\phi_i^{\rm orthog}$	2^{nd}
phaseless AFQMC (Zhang, Krakauer)	$e^{\tau(E_T\hat{\mathbf{1}}-\hat{H})}$	$\phi_i^{\rm nonorthog}$	2^{nd}

$1+\tau(E_T\hat{\mathbf{1}}-\hat{H})$ and $\frac{1}{\hat{\mathbf{1}}-\tau(E_T\hat{\mathbf{1}}-\hat{H})}$ can be used only if the spectrum of $\hat{H}$ is bounded.

Sign Problem in DMC

$\hat{P}(\tau) = e^{\tau(E_T\hat{\mathbf{1}}-\hat{H})}$

Walk is done in the space of the $3N$ coordinates of the N electrons.

$\langle \mathbf{R}|\hat{P}(\tau)|\mathbf{R}'\rangle \approx \frac{e^{\frac{-(\mathbf{R}-\mathbf{R}')^2}{2\tau}+\left(E_T-\frac{\mathcal{V}(\mathbf{R})+\mathcal{V}(\mathbf{R}')}{2}\right)\tau}}{(2\pi\tau)^{3N/2}}$ is nonnegative.

Problem: However, since the Bosonic energy is always lower than the Fermionic energy, the projected state is the Bosonic ground state.

Fixed-node approximation

All except a few calculations (release-node, Ceperley) are done using FN approximation. Instead of doing a free projection, impose the boundary condition that the projected state has the same nodes as the trial state $\Psi_{\rm T}(\mathbf{R})$.

This gives an upper bound to the energy and becomes exact in the limit that $\Psi_{\rm T}$ has the same nodes as Ψ_0.

Sign Problem in 1^{st} Quantization and R space

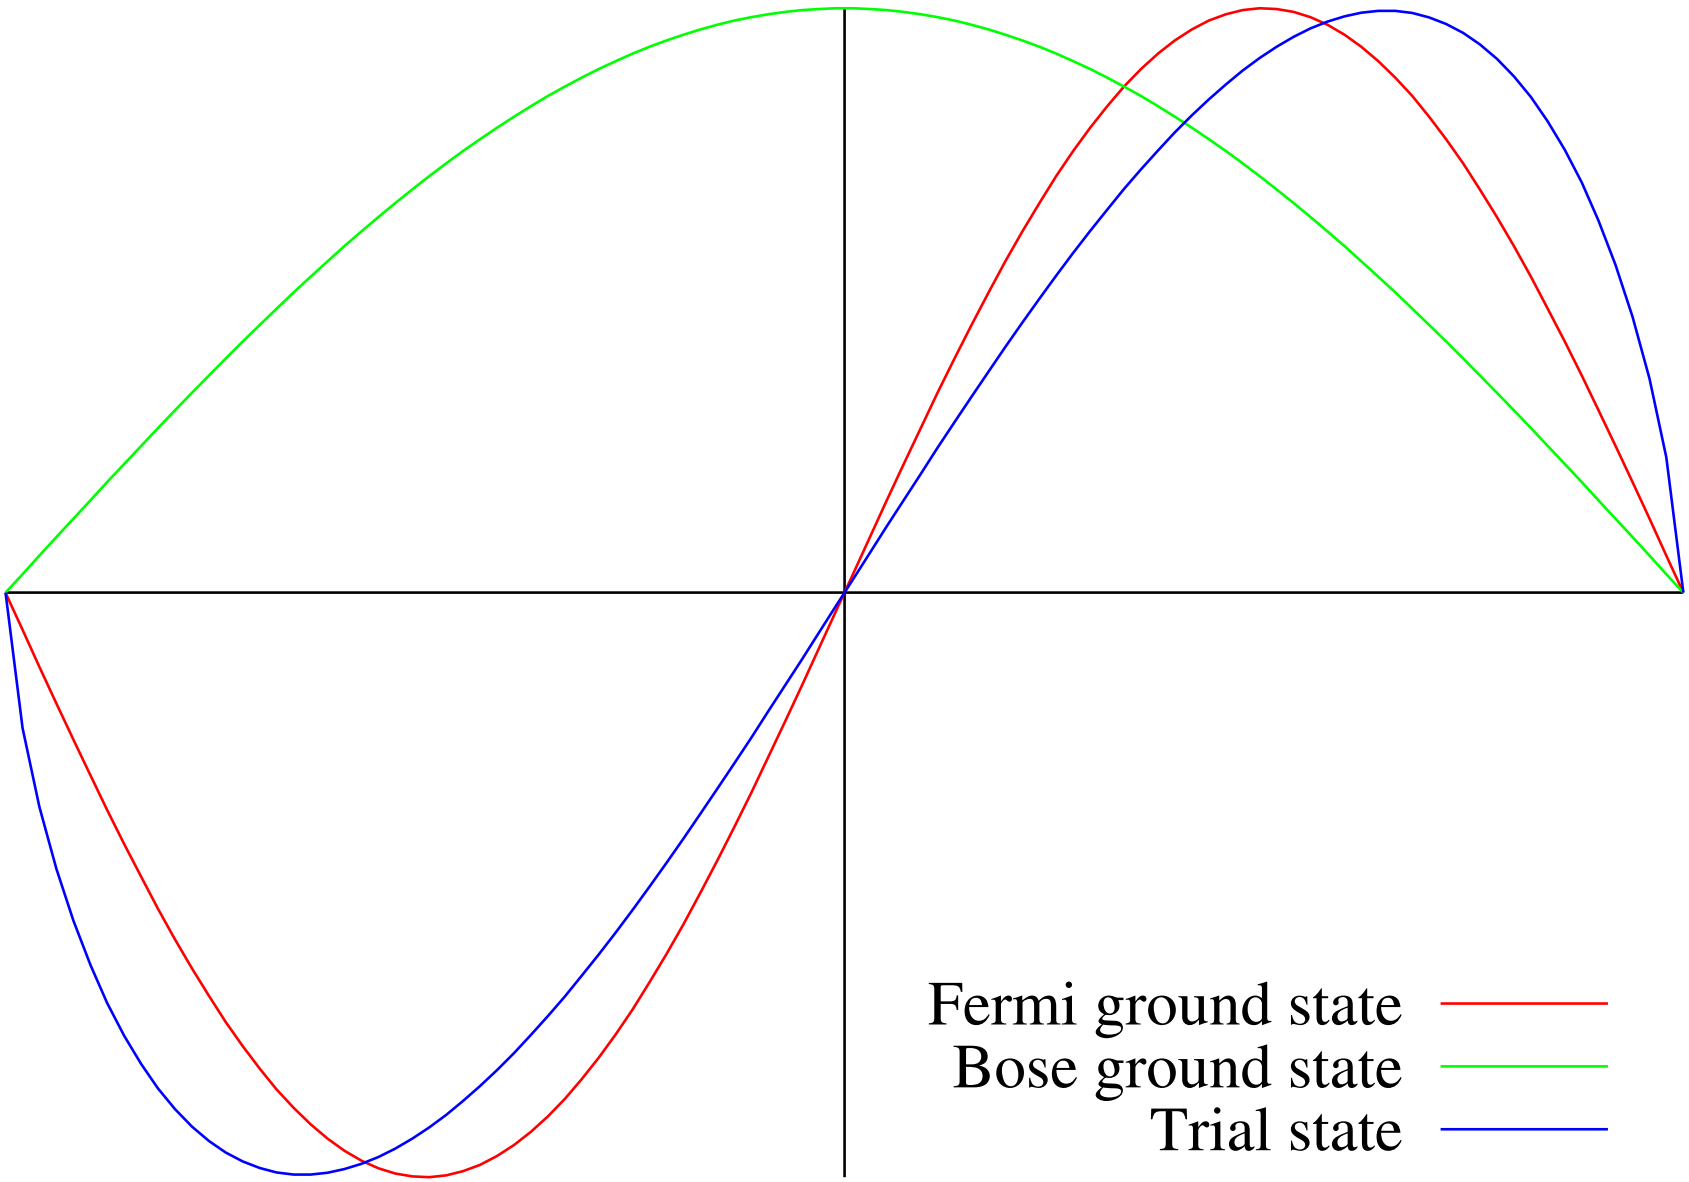

Sign Problem in 1^{st} Quantization and R space

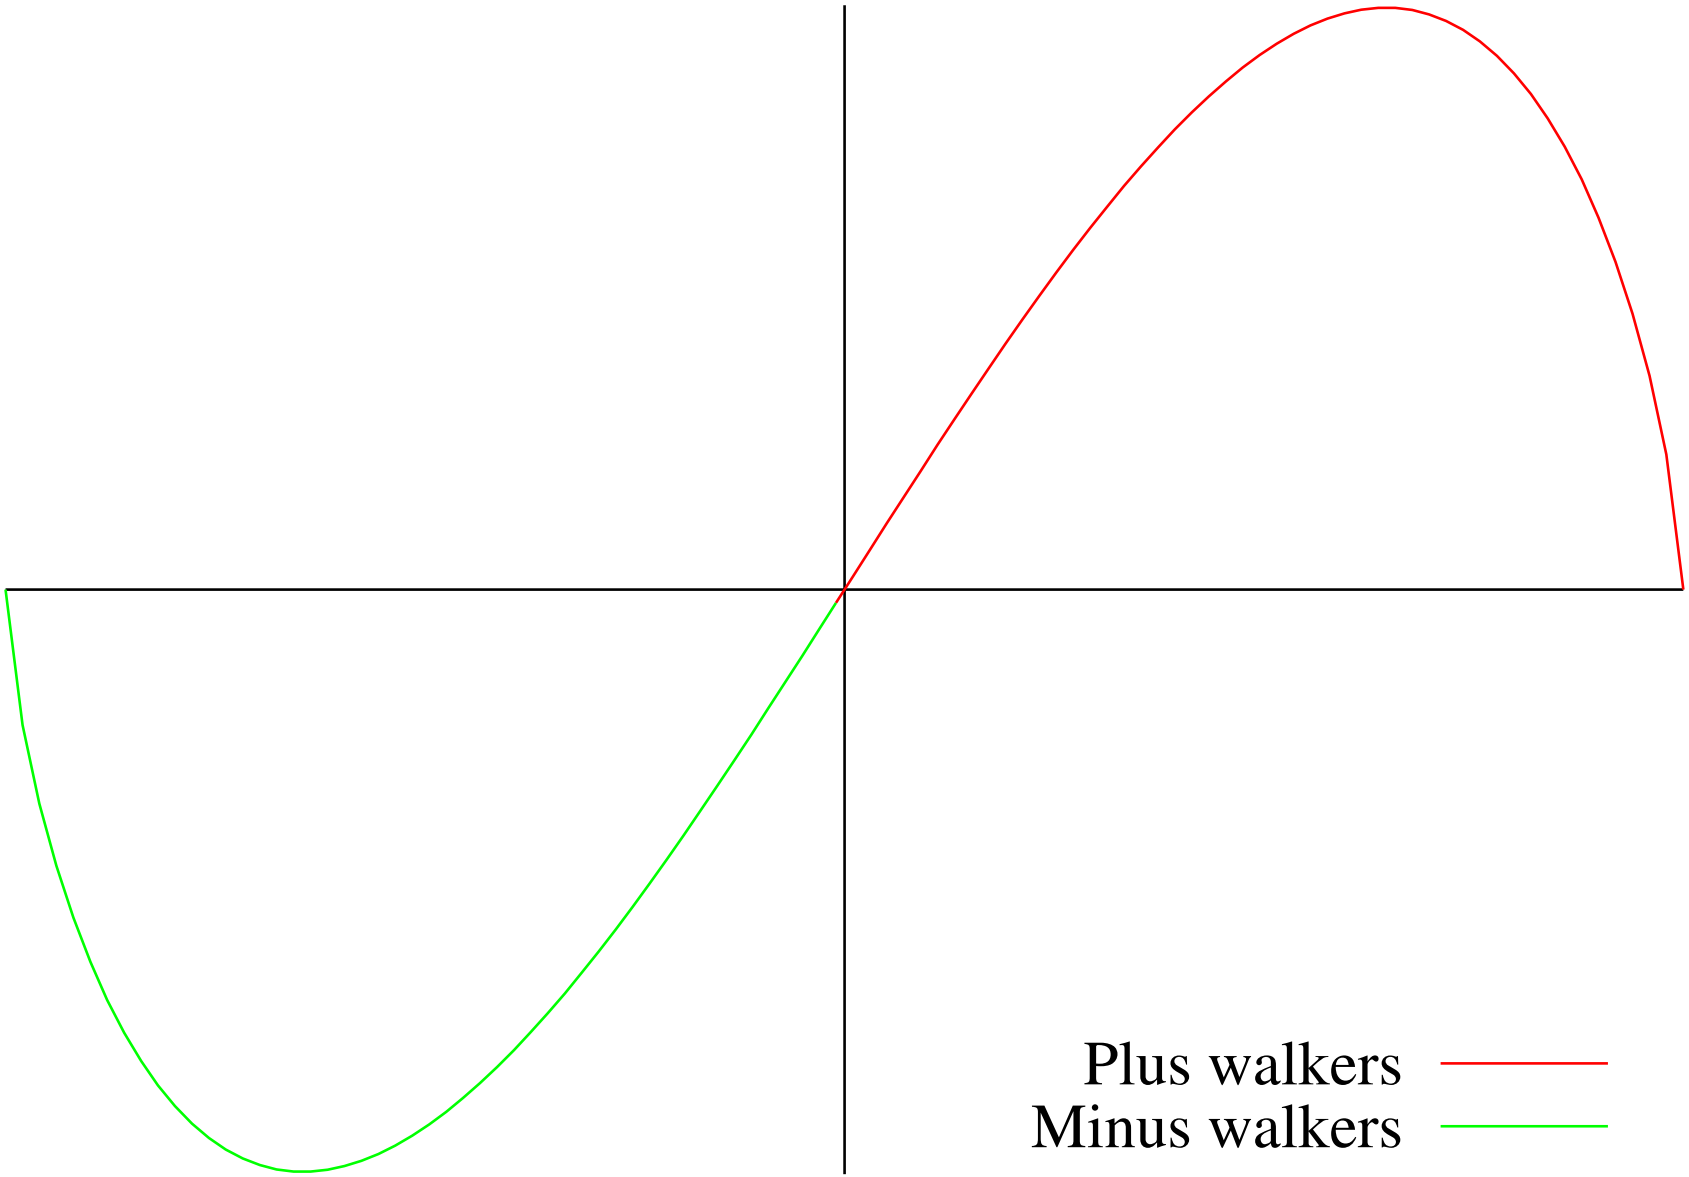

Sign Problem in 1^{st} Quantization and R space

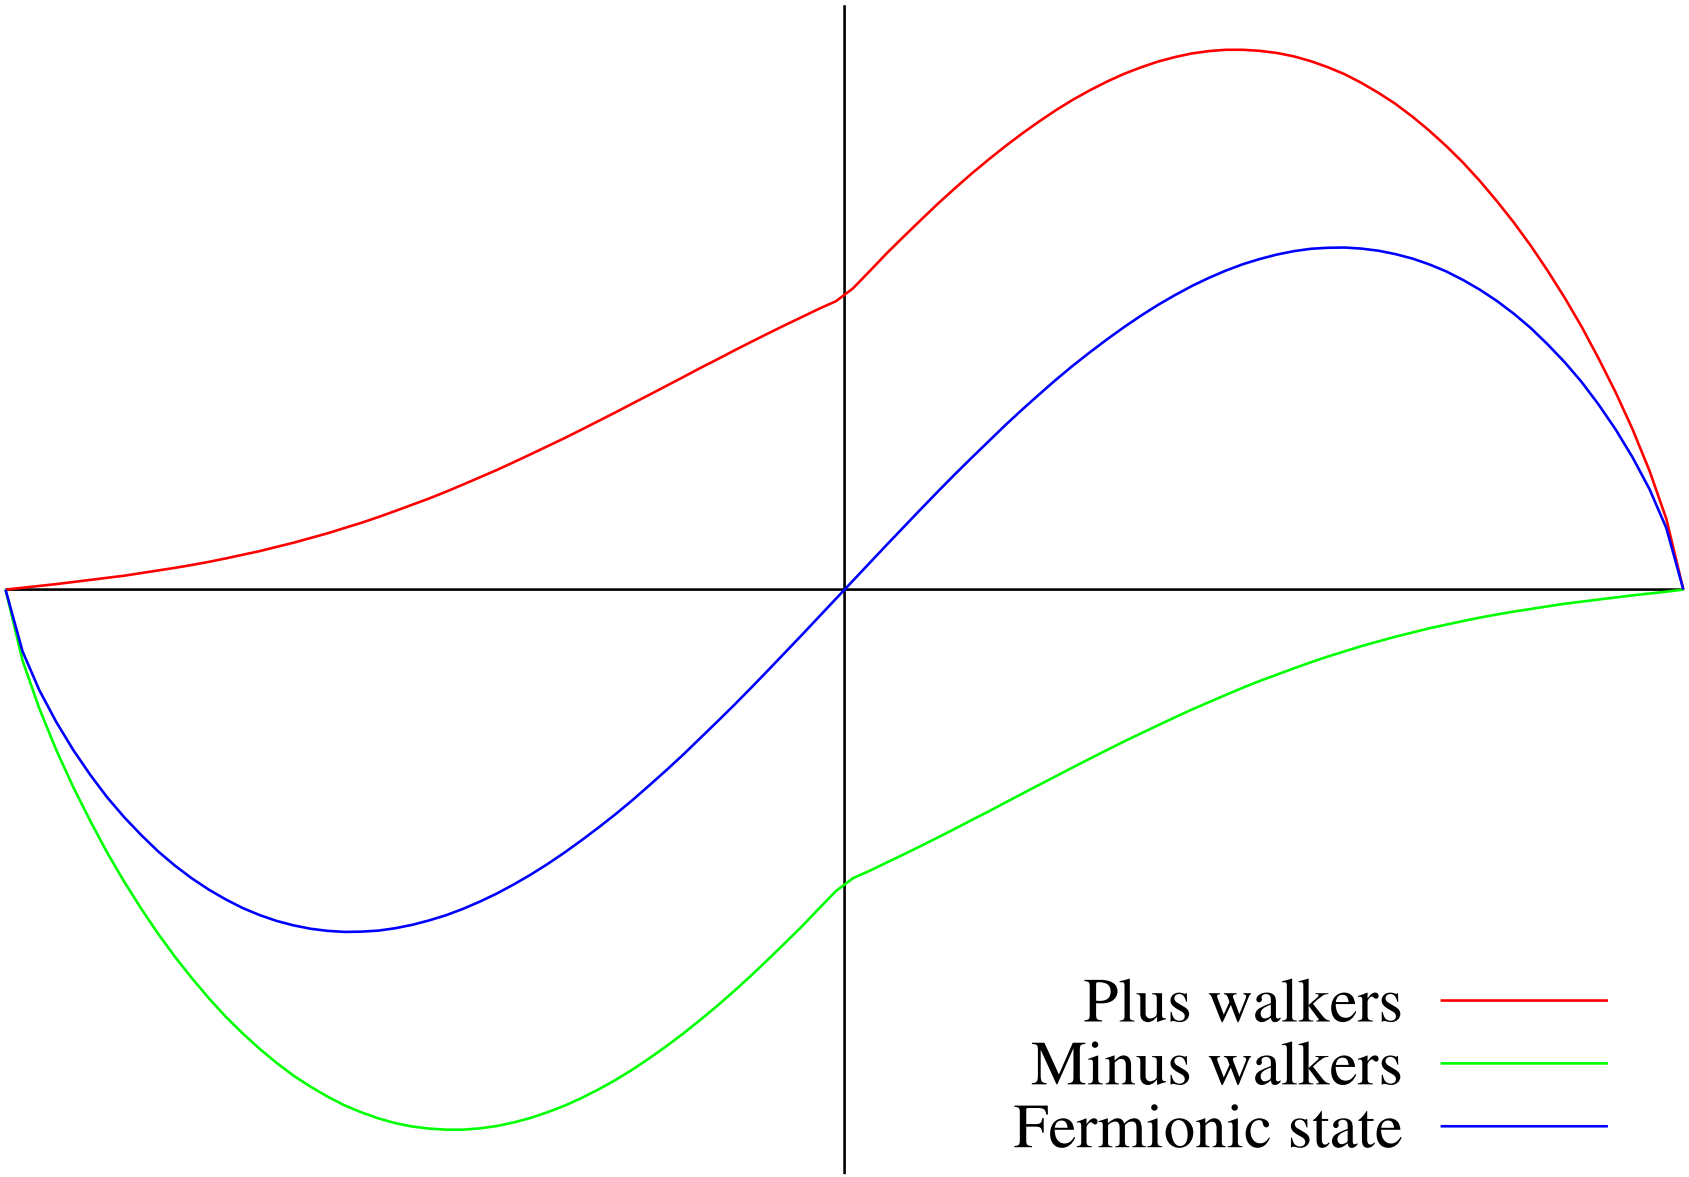

Sign Problem in 1^{st} Quantization and R space

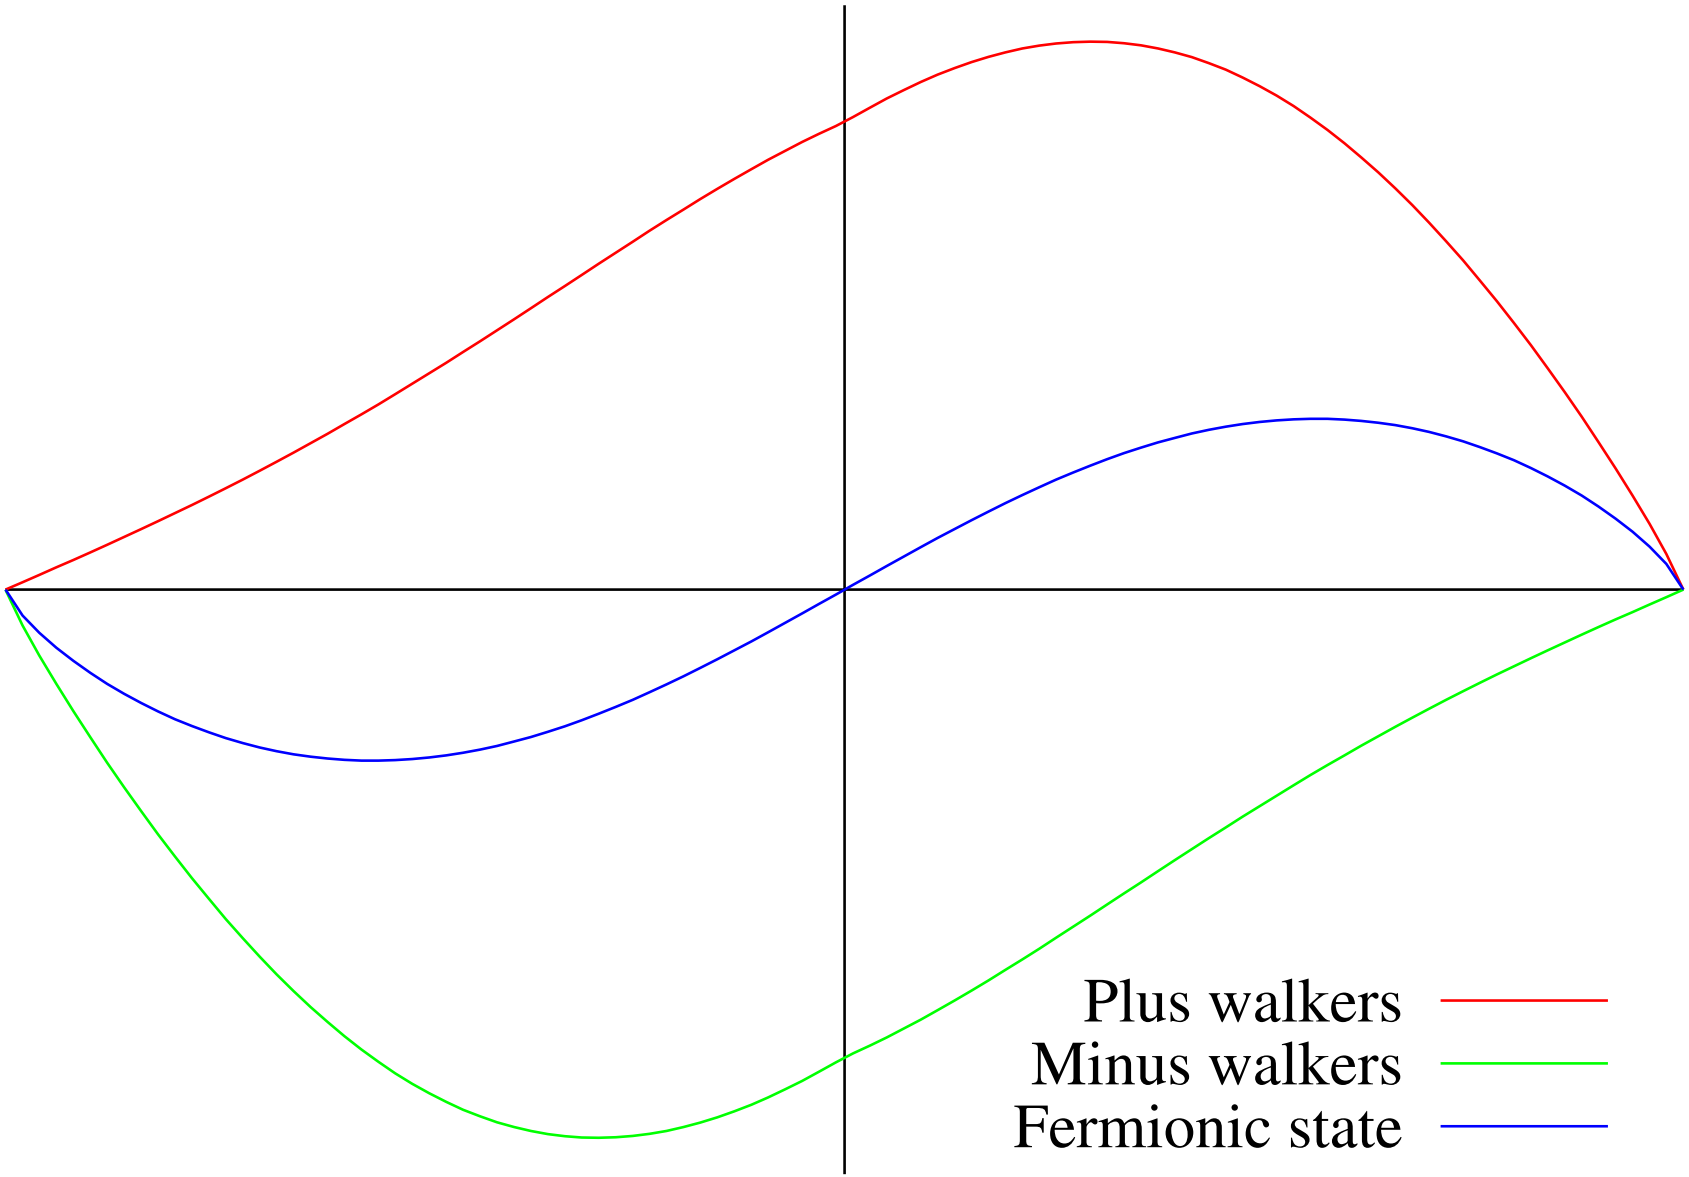

Sign Problem in 1^{st} Quantization and R space

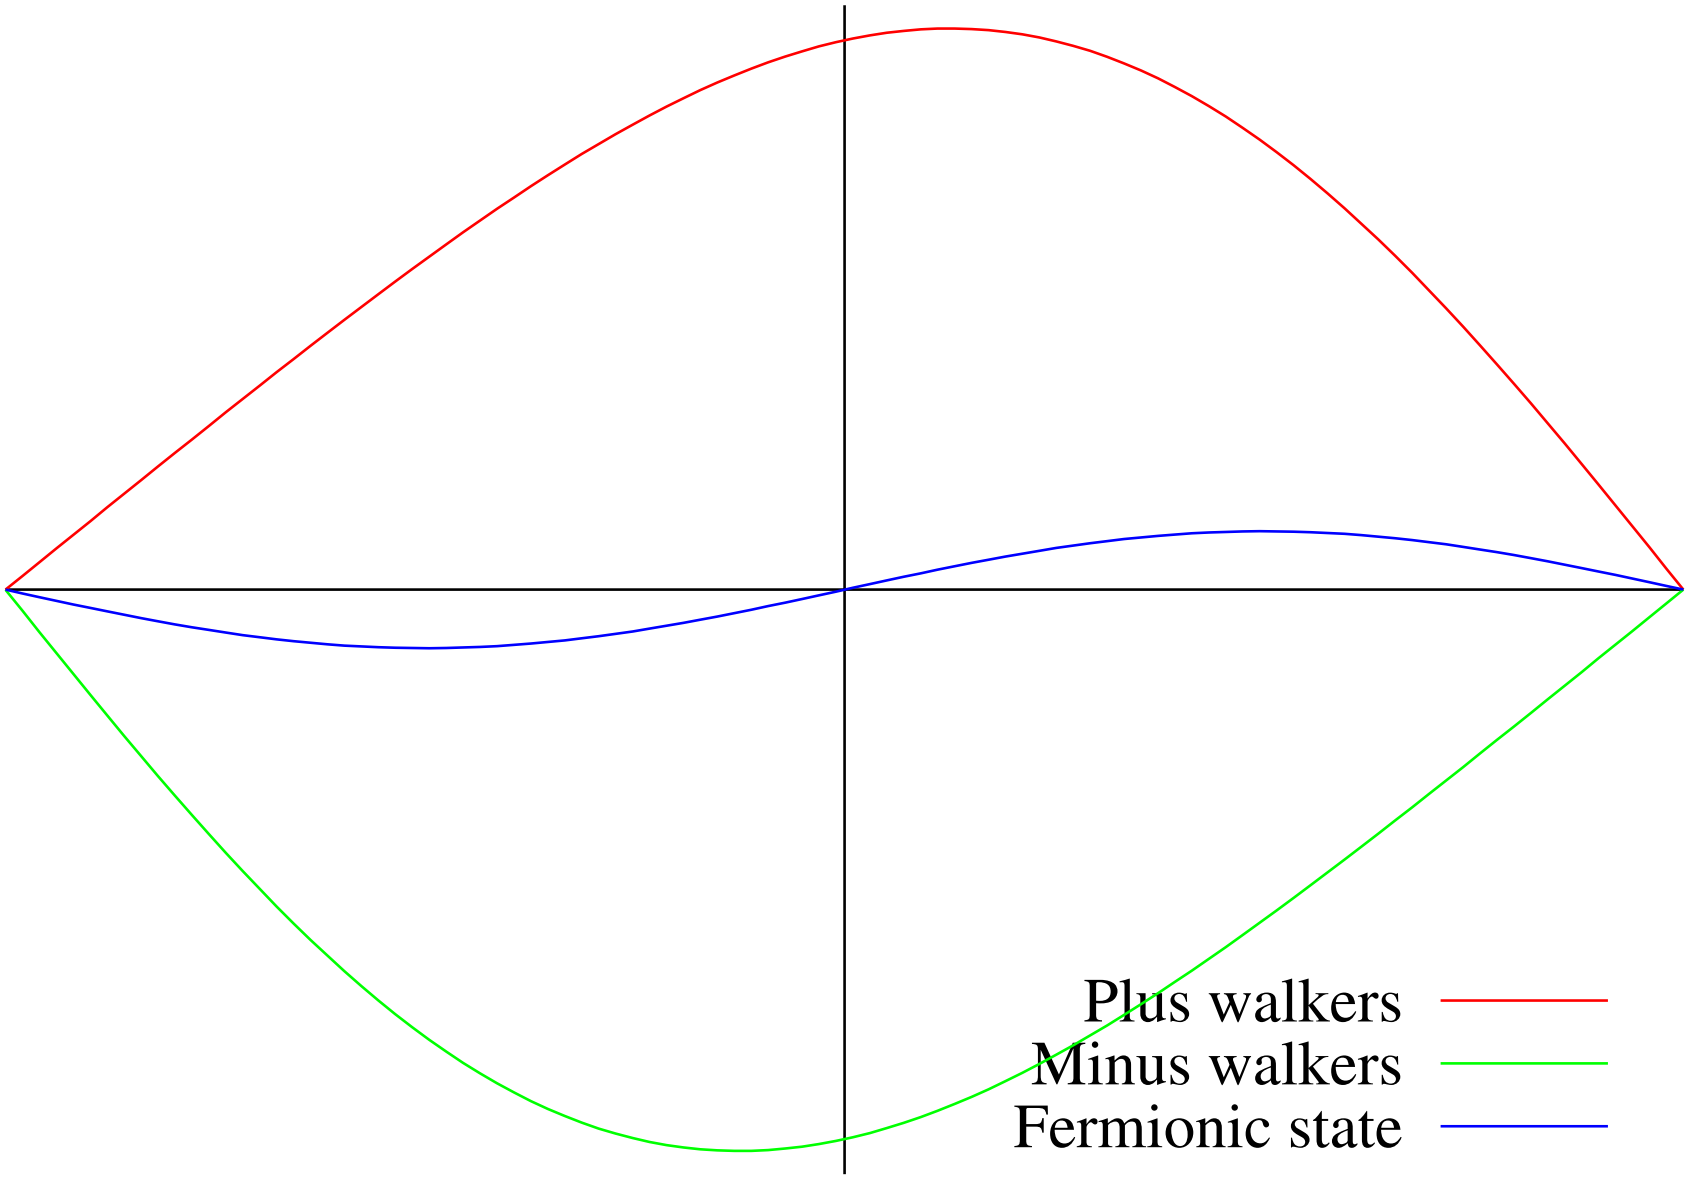

Sign Problem in 1^{st} Quantization and R space

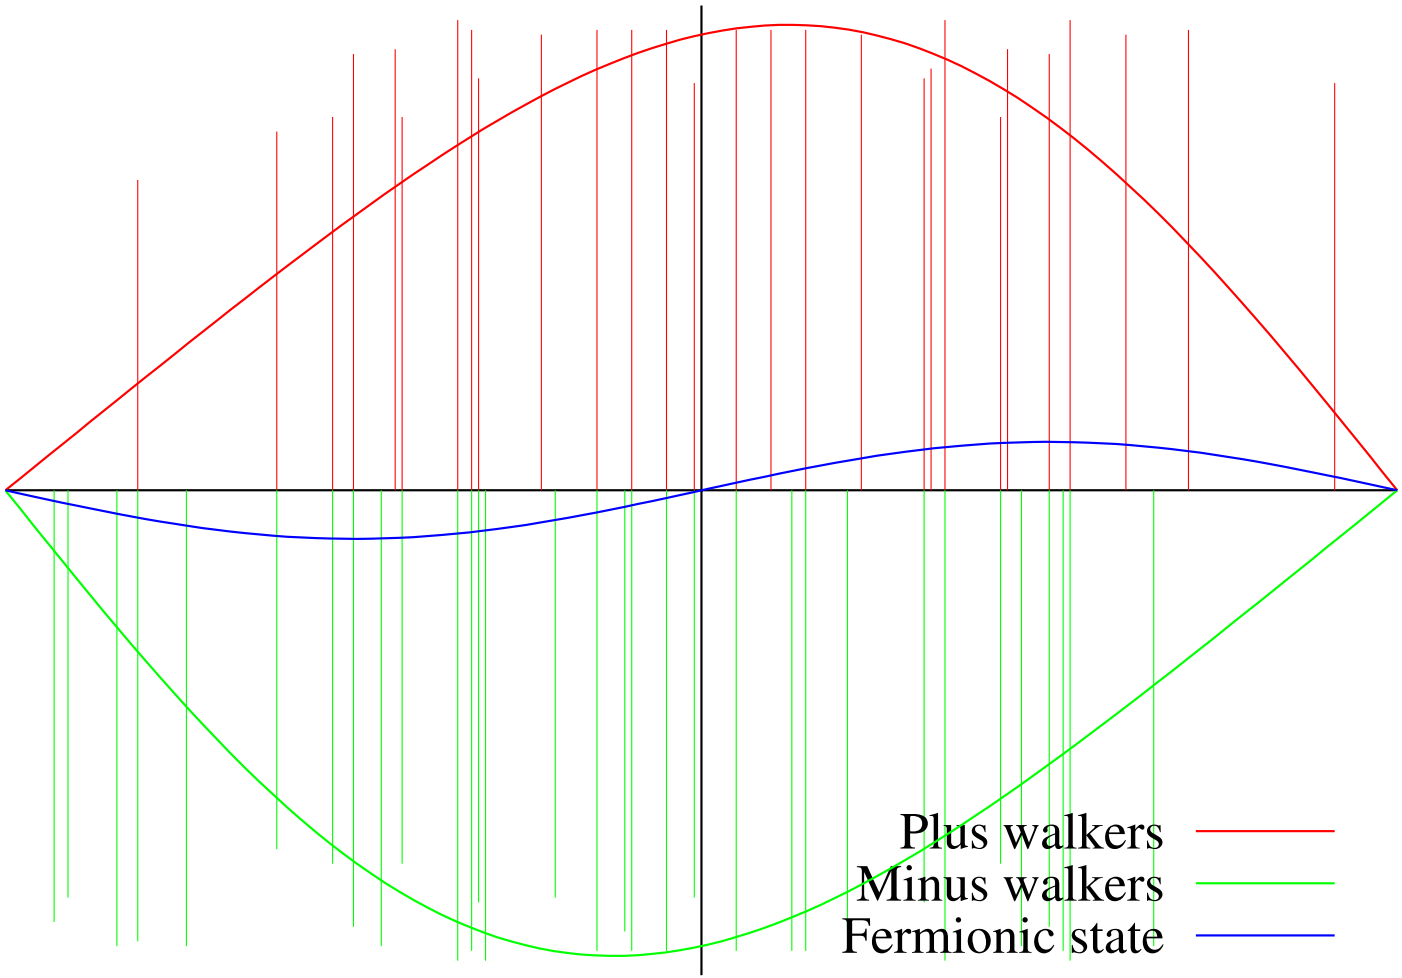

Problem: In large space walkers rarely meet and cancel, so tiny signal/noise! Further, if there are many cancellations, eventually there will be exclusively walkers of one sign only and a purely Bosonic distribution.

Sign Problem in 2^{nd} quantization

It would appear from the above discussion that one could eliminate the sign problem simply by using an antisymmetrized basis. In that case there are no Bosonic states or states of any other symmetry than Fermionic, so there is no possibility of getting noise from non-Fermionic states.

Wrong!

Sign Problem in 2^{nd} quantization

Walk is done in the space of determinants.

Since Bosonic and other symmetry states are eliminated, there is some hope of having a stable signal to noise, but there is still a sign problem.

Problem: Paths leading from state i to state j can contribute with opposite sign. Further, Ψ and $-\Psi$ are equally good.

The projector in the chosen 2^{nd}-quantized basis does not have a sign problem if:
The columns of the projector have the same sign structure aside from an overall sign, e.g.

$$P\Psi = \begin{bmatrix} + & - & + & + \\ - & + & - & - \\ + & - & + & + \\ + & - & + & + \end{bmatrix} \begin{bmatrix} + \\ - \\ + \\ + \end{bmatrix} = \begin{bmatrix} + \\ - \\ + \\ + \end{bmatrix}$$

or equivalently:
It is possible to find a set of sign changes of the basis functions such that all elements of the projector are nonnegative.

The sign problem is an issue only because of the stochastic nature of the algorithm.
Walkers of different signs can be spawned onto a given state in different MC generations.

Sign Problem in orbital space and 2^{nd} Quantization

FCIQMC (Booth, Thom, Alavi, JCP (2009)

When walk is done in space of determinants of HF orbitals, it is practical to have a population that is sufficiently large that cancellations can result in a finite signal to noise ratio. Once a critical population size is reached the probability of sign flips of the population rapidly become very small.

Initiator approximation (Cleland, Booth, Alavi, JCP (2010)

The required population size can be greatly reduced by allowing only determinants occupied by more than a certain number of walkers to spawn progeny on unoccupied determinants.

Becomes exact in the limit of infinite population size.

In subsequent papers they published FCIQMC calculations on various molecules, the homogeneous electron gas, and, real solids. Largest system has as many as 10^{108} states. (Note, however, that what matters is not the number of states, but, the number of states that have significant occupation.)

Sign Problem in FCIQMC/SQMC

Spencer, Blunt, Foulkes, J. Chem. Phys. (2012)
Kolodrubetz, Spencer, Clark, Foulkes, J. Chem. Phys. (2013)

1. The instability gap is given by the difference in the dominant eigenvalues of the projector, and, those of the projector with all off-diagonal elements replaced by their absolute values.

2. More than 1 Hartree product in a given initial determinant may connect via P (or H) to a given Hartree product in a final determinant. The instability gap is smaller in 2^{nd} quantization than in 1^{st} quantization if there are internal cancellations within these contributions, otherwise it is the same as in 1^{st} quantization.
   For example, it is the same in lattice real-space Coulomb systems, real- and momentum-space Hubbard models, but, is different for orbital-space Coulomb systems.

Sign Problem in FCIQMC/SQMC

These papers did not point out that even when the instability gap is the same, there are two important advantages of 2^{nd} quantization:

1. Since the Hilbert space is $N!$ times smaller in 2^{nd} quantization, cancellation are much more effective.
2. In first quantization, one of the two Bosonic populations will dominate and the signal to noise will go to zero even in the limit of an infinite population, unless additional steps are taken to prevent that.

Using a large population and cancellations, it is possible to get a finite signal to noise ratio in 2^{nd} quantization but not in 1^{st} quantization (unless some further constraints are imposed).

Original attempts at using cancellation to control sign problem (in continuum problems): Mal Kalos and coworkers (David Arnow (1982), Shiwei Zhang, Francesco Pederiva, ...)

History of Projector Quantum Monte Carlo

1962 Kalos invents Green's function MC for continuum problems

1974 Kalos, Leveque, Verlet invent another Green's function MC for continuum problems

1983 Blankenbecler and Sugar invented Projector MC for 1-D lattice model (no sign problem)

1984 Ceperley, Alder invent release-node method (converges before noise blows up for some systems)

1989 Trivedi and Ceperley, Heisenberg antiferromagnet (no sign problem)

1994 Bemmel, et al., invent fixed-node approximation for lattice Fermions

many Kalos discusses cancellations of $+ve$ and $-ve$ walkers in continuum for controlling sign problem, but gets it to work for toy problems only

2009 Booth, Thom, Alavi, get cancellations to work in finite basis, but number of walkers very large, exponential cost

2010 Cleland, Booth, Alavi, introduce systematically improvable initiator approximation that greatly reduces population size, exponential cost, reduced exponent

Semistochastic Quantum Monte Carlo (SQMC)

Frank Petruzielo, Adam Holmes, Hitesh Changlani, Peter Nightingale, CJU, PRL 2012

SQMC is hybrid of Exact Diagonalization and QMC

Exact diagonalization has no statistical error or sign problem but is limited to a small number of states ($\sim 10^{10}$ on a single core).

QMC has statistical errors and a sign problem but can employ a much larger number of states.

SQMC combines to some extent the advantages of the above by doing a deterministic projection in a small set of important states and stochastic projection in the rest of the space. It has a much smaller statistical error than stochastic projection and can employ a large number of states.

More generally Semistochastic Projection is an efficient way to find the dominant eigenvalue and corresponding expectation values of any large sparse matrix that has much of its spectral weight on a manageable number of states.

Semistochastic Projection

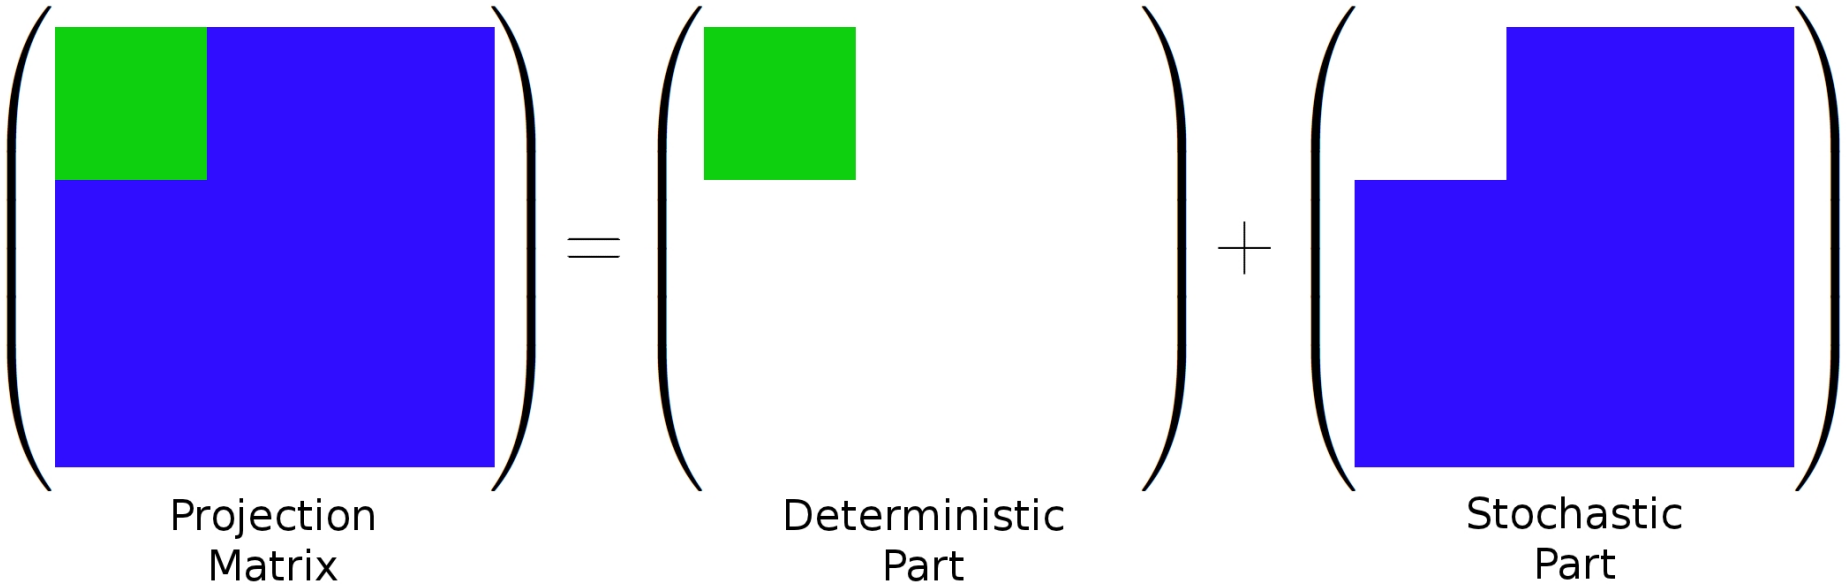

The part of the projection with both indices in the deterministic part is done deterministically. The part of the projection with either index in the stochastic part is done stochastically.

$$P = P^{\mathcal{D}} + P^{\mathcal{S}}$$

$$P^{\mathcal{D}}_{ij} = \begin{cases} P_{ij}, & i,j \in \mathcal{D} \\ 0, & \text{otherwise} \end{cases}$$

$$P^{\mathcal{S}} = P - P^{\mathcal{D}}$$

Diagonal elements in $P^{\mathcal{S}}$

The contribution to the total walker weight on $|\phi_j\rangle$, with $j \in \mathcal{S}$, is

$$P_{jj} w_j(t) \quad = \quad [1 + \tau(E_T - H_{jj})]\, w_j(t)$$

Off-diagonal elements in $P^{\mathcal{S}}$

Weight w_i is divided amongst $n_i = \max(\lfloor w_i \rceil, 1)$ walkers of wt. w_i/n_i.
For each walker on $|\phi_i\rangle$, a move to $|\phi_j\rangle \neq |\phi_i\rangle$ is proposed with probability $T_{ji} > 0$, $(\sum_j T_{ji} = 1)$, where T is the proposal matrix.

The magnitude of the contribution to the walker weight on $|\phi_j\rangle$ from a single walker on $|\phi_i\rangle$ is

$$\begin{cases} 0, & i,j \in \mathcal{D} \\ \dfrac{P_{ji}}{T_{ji}} \dfrac{w_i(t)}{n_i(t)} = -\tau \dfrac{H_{ji}}{T_{ji}} \dfrac{w_i(t)}{n_i(t)} & \text{otherwise} \end{cases}$$

Elements in $P^{\mathcal{D}}$

The contribution to the weight on $|\phi_j\rangle$, with $j \in \mathcal{D}$, is

$$\sum_{i \in \mathcal{D}} P^{\mathcal{D}}_{ji} w_i(t).$$

$P^{\mathcal{D}}$ is stored and applied as a sparse matrix

Semistochastic Projection

Walkers have a label (bit string of orbital occupation numbers) and signed real weights.

Project Do deterministic and stochastic projection

Sort Walker labels are sorted.

Merge Walkers on the same determinant are merged

Initiator The initiator criterion is used to discard some walkers.

Join Because we use real weights, there are many walkers with small weights. Join stochastic space small wt. walkers on different determinants (in unbiased way).

Update Energy Use stored E_{L} components to update energy estimator. So E_{L} never needs to be computed during body of run.

The only additional steps are the deterministic projection and the "join" step.

SQMC

Precompute:

Before MC part of the calculation do following:

1. Choose the deterministic space $\mathcal{D}$ and precompute matrix elements of projector, P, between all pairs of deterministic determinants.
2. Choose the trial wave function, $\Psi_{\rm T}$, and precompute the local energy components of all determinants connected to those in $\Psi_{\rm T}$.

Efficiency gain for C_2 ($3-\zeta$ basis) from semistochastic projection and Ψ_T

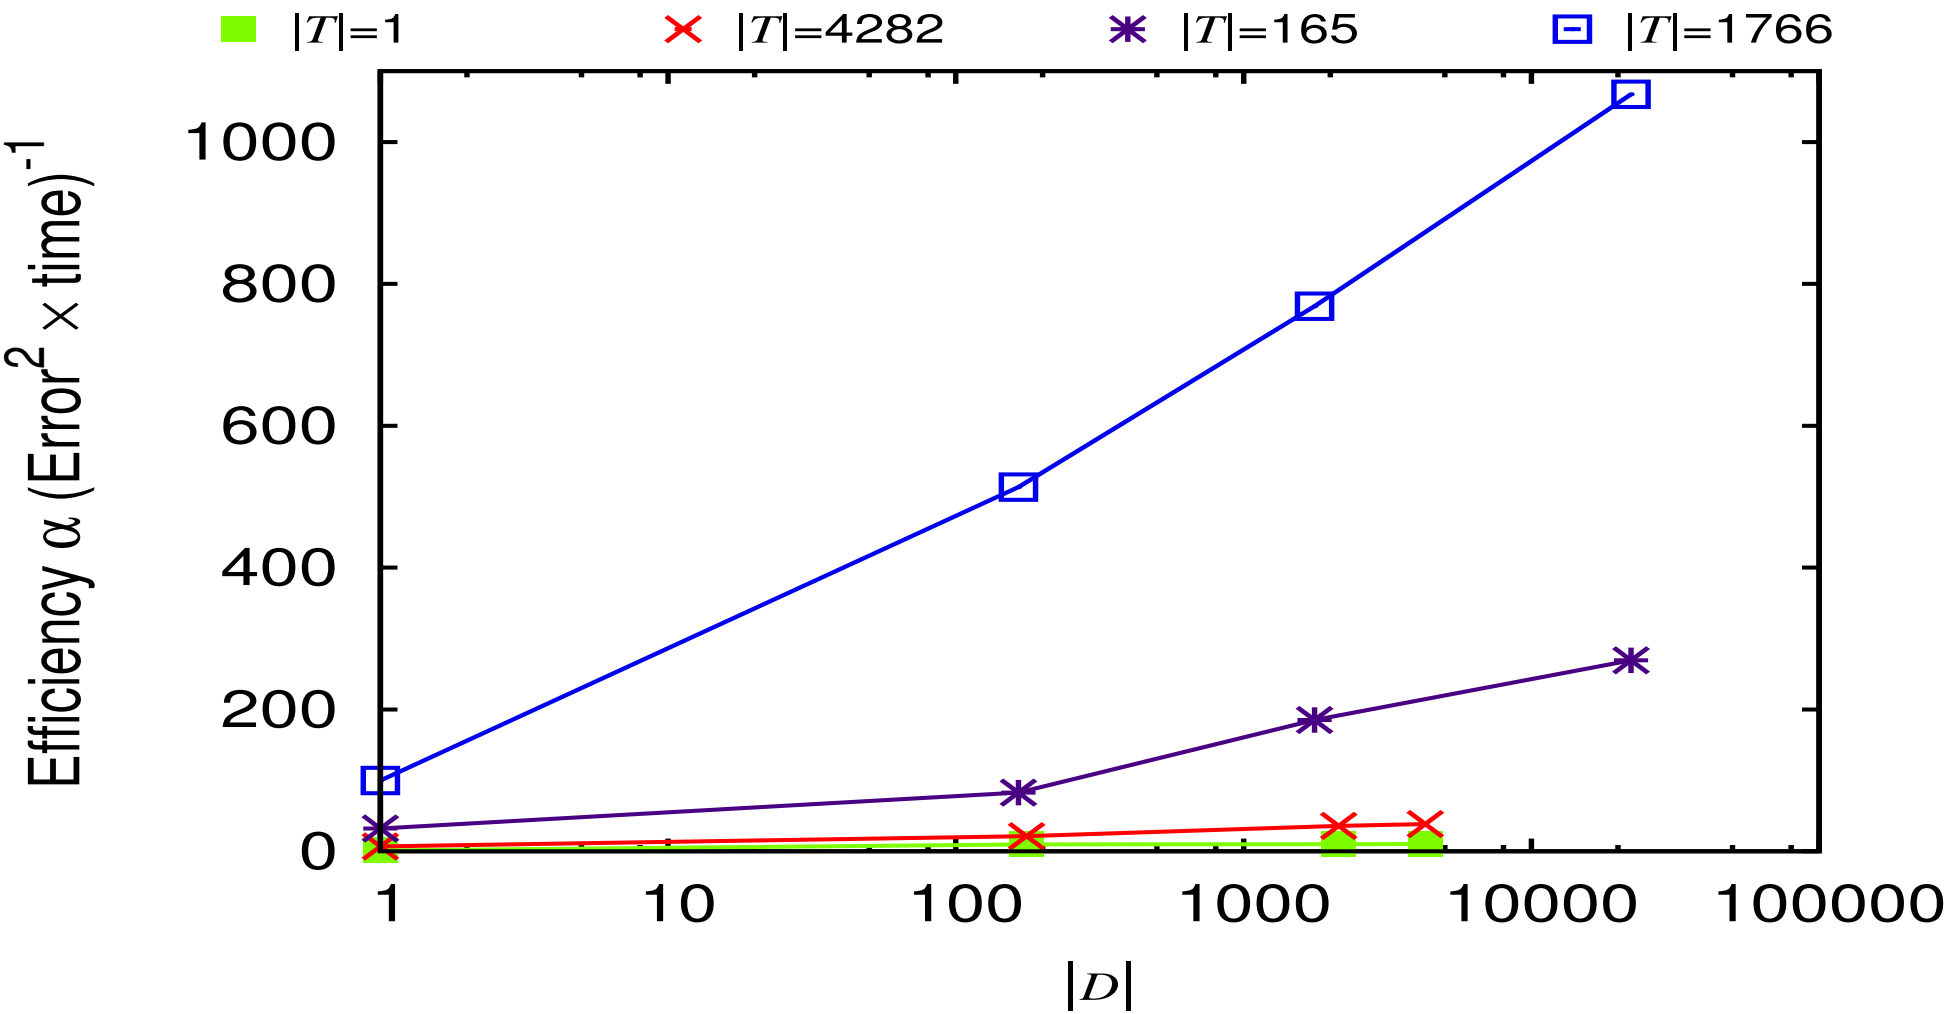

Wavefns. with 165 or 1766 dets. containing some 4^{th}-order excit. are much more efficient than wavefn. with 4282 dets. containing only upto 2^{nd}-order excit.

Discrete-space fixed-node approximation

Even with the semistochastic approach, calculating the exact energy when there is a sign problem is prohibitely expensive if a very large number of states contribute significantly. In that case, one can do approximate calculations using the fixed-node method.

Discrete-space fixed-node approximation

ten Haaf, Bemmel, Leeuwen, van Saarloos, Ceperley, PRB (1995)

Starting from importance sampled $\tilde{P}_{ji} = g_j P_{ji}/g_i$, define $\tilde{P}^{\rm FN}_{ji}$ and corresponding $\tilde{H}^{\rm FN}_{ji}$,

$$\tilde{P}^{\rm FN}_{ji} = \begin{cases} \tilde{P}_{ji}, & j \neq i, \tilde{P}_{ji} \geq 0, \\ 0, & j \neq i, \tilde{P}_{ji} < 0, \\ P_{ii} + \sum_k \min(\tilde{P}_{ki}, 0) & j = i, \end{cases}$$

Note that $P^{\rm FN}_{{\rm L},i} = \sum_j \tilde{P}^{\rm FN}_{ji} = \sum_j \tilde{P}_{ji} = P_{{\rm L},i}$, so $\hat{H}$ and $\hat{H}^{\rm FN}$ have the same local energies.

If there is no sign problem, then all the elements of $\tilde{P}$ are nonnegative. Changing the sign of a basis state changes the sign of the corresponding row and column of P (aside from the diagonal element) but $\tilde{P}$ is unchanged because the corresponding element of $\Psi_{\rm G}$ also changes.

Algorithm 1: ten Haaf et al.

1. Propose move from i to $j \neq i$ with probability $T_{ji} = \frac{\max(\tilde{P}_{ji},0)}{\sum'_k \max(\tilde{P}_{ki},0)}$
2. Give walker on chosen state weight, $w_i \left(\sum'_k \max(\tilde{P}_{ki}, 0) \right)$
3. Give walker on state i weight, $w_i \left(P_{ii} + \sum'_k \min(\tilde{P}_{ki}, 0) \right)$

where the prime on the sum indicates that $j = i$ is omitted from the sum. The diagonal could become negative, especially on high-lying states (both because P_{ii} and g_i are small), so use small τ and possibly place lower bound on $|g_i|$ to ensure that it does not.

Discrete-space fixed-node approximation

Algorithm 2: Kolodrubetz, Clark, PRB 2012

1. Propose move from i to $j \neq i$ with any probability T_{ji}
2. Give walker on chosen state weight, $w_i \left(\frac{\max(\tilde{P}_{ji},0)}{T_{ji}} \right)$
3. Give walker on state i weight, $w_i \left(P_{ii} + \frac{\min(\tilde{P}_{ji},0)}{T_{ji}} \right)$

Advantage: It is no longer necessary to perform the sum $\sum_k' \max(\tilde{P}_{ki}, 0)$, which can be very expensive if the number of connections to state i is large.

Disadvantage: If one makes a simple choice such as $T_{ji} = 1/N_{\rm connections}$ (uniform moves) then it becomes necessary to use even smaller τ than in the previous algorithm, because the walker is no longer moving preferentially to states with large $|\tilde{P}_{ji}|$ and so, when it does move there, it makes a larger contribution, which is problematic if the contribution is negative.

Solution: Instead of using the linear projector, $\hat{\mathbf{1}} + \tau(E_T\hat{\mathbf{1}} - \hat{H})$ devise an exact (i.e. no Trotter time-step error) algorithm to sample the exponential projector $e^{\tau(E_T\hat{\mathbf{1}}-\hat{H})}$. In certain areas of QMC, such algorithms are called "continuous time methods". We will not discuss this further in these lectures.

Desirable properties of fixed-node approximations

For the moment, assume $\Psi_{\rm G} = \Psi_{\rm T}$.

Property	discrete-space 2^{nd}-quant. FN	continuous-real-space 1^{st}-quant. FN-DMC
Exact distribution in $\Psi_{\rm G} \to \Psi_0$ limit	Y	Y
Exact energy when sign of $\Psi_{\rm G} = \Psi_{\rm T}$ correct $\forall$ states	N[a]	Y
FN energy lower than variational energy $E_{\rm FN} \leq E_V$	Y	Y
$E_{\rm FN} \geq E_0$	Y	Y
Projector elements do not become negative for large τ	N	Y
No time-step error	Y	N

[a] Because $\tilde{P}$ may have sign-flip elements even when sign of $\Psi_{\rm G} = \Psi_{\rm T}$ correct on all states.

Now we prove the 4 properties with Y in the discrete-space column.
The last of these is obvious.

Desirable properties of fixed-node approximations

Gives exact distribution in $\Psi_G \rightarrow \Psi_0$ limit

The FN approx. makes 2 changes to the weight on a state i.

1) Enhances the weight on state i because it acquires no negative contributions from connected states.

2) Reduces the weight on state i because negative weight that would have gone to other states is now put back on state i.

The expected values of these 2 contributions are:

$$-\sum_j{}' \tilde{P}_{ij} g_j^2 \quad = \quad -\sum_j{}' g_i P_{ij} g_j$$

and

$$\sum_j{}' \tilde{P}_{ji} g_i^2 \quad = \quad \sum_j{}' g_i P_{ij} g_j$$

where the prime on the sum indicates that we sum over negatively connected states only. Since P is symmetrical, these 2 contributions cancel and the equilibrium distribution is unchanged by the fixed node approx. in the $\Psi_G = \Psi_0$ limit.

Desirable properties of fixed-node approximations

$\underline{E_{\rm FN} \leq E_V}$

$$E_{\rm FN} \;=\; \langle \Psi_{\rm FN} | H_{\rm FN} | \Psi_{\rm FN} \rangle \leq \langle \Psi_{\rm T} | H_{\rm FN} | \Psi_{\rm T} \rangle = \langle \Psi_{\rm T} | H | \Psi_{\rm T} \rangle = E_V$$

Desirable properties of fixed-node approximations

$\underline{E_{\rm FN} \geq E_0}$

Consider an arbitrary wavefunction, $|\Psi\rangle = \sum_i \; c_i |\phi_i\rangle$

$$\begin{aligned}
\langle\Psi|H_{FN} - H|\Psi\rangle &= \sum_i c_i \sum_j{}' \left(H_{ji} \frac{g_j}{g_i} c_i - H_{ji} c_j \right) && \text{prime means only sum over sign-flip terms} \\
&= \sum_{\{i,j\}}{}' H_{ji} \left[c_i^2 \frac{g_j}{g_i} + c_j^2 \frac{g_j}{g_i} - 2 c_i c_j \right] && \text{summing over pairs} \\
&= \sum_{\{i,j\}}{}' |H_{ji}| \left[c_i^2 \left| \frac{g_j}{g_i} \right| + c_j^2 \left| \frac{g_i}{g_j} \right| - 2 c_i c_j \mathrm{sgn}(H_{ji}) \right] && \text{since only sign-flip terms appear} \\
&= \sum_{\{i,j\}}{}' |H_{ji}| \left[c_i \sqrt{\left| \frac{g_j}{g_i} \right|} - \mathrm{sgn}(H_{ji}) c_j \sqrt{\left| \frac{g_i}{g_j} \right|} \right]^2 \\
&\geq 0
\end{aligned}$$

So

$$E_{\rm FN} \;=\; \langle\Psi_{FN}|H_{FN}|\Psi_{FN}\rangle \;\;\geq\;\; \langle\Psi_{FN}|H|\Psi_{FN}\rangle \;\geq\; \langle\Psi_0|H|\Psi_0\rangle \;=\; E_0$$

Comparison of DMC with FCIQMC/SQMC

DMC (walk in electron coordinate space)	FCIQMC/SQMC (walk in determin. space)
Severe Fermion sign problem due to growth of Bosonic component relative to Fermionic.	Less severe Fermion sign problem due to opposite sign walkers being spawned on the same determinant
Fixed-node approximation needed for stable algorithm. Exact if $\Psi_{\rm T}$ nodes exact.	Walker cancellation plus initiator approximation needed for stable algorithm. Exact in ∞-population limit.
Infinite basis.	Finite basis. (Same basis set dependence as in other quantum chemistry methods.
Computational cost is low-order polynomial in N	Computational cost is exponential in N but with much smaller exponent than full CI
Energy is variational	Energy not variational but DM variant is
Need to use pseudopotentials for large Z.	Can easily do frozen-core

Diffusion Monte Carlo

i.e., $\hat{P}(\tau) = \exp(\tau(E_T - \hat{H}))$, $|\phi_i\rangle = |\mathbf{R}\rangle$, walkers are 1^{st}-quantized

The imaginary-time evolved state $|\Psi(t)\rangle$ is

$$|\Psi(t)\rangle = e^{-\hat{H}t}|\Psi(0)\rangle = \sum_i e^{-\hat{H}t}|\Psi_i\rangle\langle\Psi_i|\Psi(0)\rangle = \sum_i e^{-E_i t}|\Psi_i\rangle\langle\Psi_i|\Psi(0)\rangle$$

$$= \sum_i c_i(t)\,|\Psi_i\rangle,$$

$c_i(t) = c_i(0)e^{-E_i t}$, $c_i(0) = \langle\Psi_i|\Psi(0)\rangle$. Expon. decay of excited states

$$E_{\rm mix} = \lim_{t\to\infty} \frac{\langle\Psi_{\rm T}|\hat{H}e^{-\hat{H}t}|\Psi_{\rm T}\rangle}{\langle\Psi_{\rm T}|e^{-\hat{H}t}|\Psi_{\rm T}\rangle} = \frac{\langle\Psi_{\rm T}|\hat{H}|\Psi_0\rangle}{\langle\Psi_{\rm T}|\Psi_0\rangle} = E_0\frac{\langle\Psi_{\rm T}|\Psi_0\rangle}{\langle\Psi_{\rm T}|\Psi_0\rangle} = E_0$$

The mixed estimator for the energy is exact (for Bosonic ground states) but expectation values of operators that do not commute with the energy are not.

Diffusion Monte Carlo

In DMC the MC walk is in real space, so introducing resolution of identity in terms of the complete set of position eigenstates,

$$\begin{aligned} \langle \mathbf{R}'|\Psi(t)\rangle &= \int d\mathbf{R}\, \langle \mathbf{R}'|e^{-\hat{H}t}|\mathbf{R}\rangle\, \langle \mathbf{R}|\Psi(0)\rangle \\ \Psi(\mathbf{R}',t) &= \int d\mathbf{R}\, G(\mathbf{R}',\mathbf{R},t)\, \Psi(\mathbf{R},0) \end{aligned}$$

where $G(\mathbf{R}',\mathbf{R},t) = \langle \mathbf{R}'|e^{-\hat{H}t}|\mathbf{R}\rangle$ is called the Green's function or the real-space representation of the imaginary-time-evolution operator.
Introducing complete sets of $\hat{H}$ eigenstates,

$$\begin{aligned} G(\mathbf{R}',\mathbf{R},t) &\equiv \langle \mathbf{R}'|e^{-\hat{H}t}|\mathbf{R}\rangle = \sum_i \langle \mathbf{R}'|\Psi_i\rangle\, \langle \Psi_i|e^{-\hat{H}t}|\Psi_i\rangle\, \langle \Psi_i|\mathbf{R}\rangle \\ &= \sum_i \Psi_i(\mathbf{R}')\, e^{-E_i t}\, \Psi_i^*(\mathbf{R}) \quad \text{spectral representation} \end{aligned}$$

but this is not of practical utility since $\Psi_i(\mathbf{R})$ are unknown.
Note that $G(\mathbf{R}',\mathbf{R},0) = \sum_i\, \Psi_i(\mathbf{R}')\, \Psi_i^*(\mathbf{R}) = \delta(\mathbf{R}'-\mathbf{R})$.

Diffusion Monte Carlo

Do not know $\langle \mathbf{R}'|e^{-\hat{H}t}|\mathbf{R}\rangle$ exactly, but can evaluate $G_{\rm rewt}(\mathbf{R}',\mathbf{R},t) \equiv \langle \mathbf{R}'|e^{-\hat{V}t}|\mathbf{R}\rangle$ exactly, if $\hat{V}$ is local in real space, $G_{\rm diffusion}(\mathbf{R}',\mathbf{R},t) \equiv \langle \mathbf{R}'|e^{-\hat{T}t}|\mathbf{R}\rangle$ exactly by introducing momentum representation

But, $e^{-t\hat{H}} \neq e^{-t\hat{T}}e^{-t\hat{V}}$, so repeatedly use short-time approx., $t = M\tau$.

$$G(\mathbf{R}',\mathbf{R},t) = \langle \mathbf{R}'| \left(e^{-t\hat{H}/M}\right)^M |\mathbf{R}\rangle$$

$$= \int d\mathbf{R}_1 \cdots d\mathbf{R}_{M-1} \, \langle \mathbf{R}'|e^{-\tau\hat{H}}|\mathbf{R}_{M-1}\rangle \, \langle \mathbf{R}_{M-1}|e^{-\tau\hat{H}}|\mathbf{R}_{M-2}\rangle \, \cdots \, \langle \mathbf{R}_1|e^{-\tau\hat{H}}|\mathbf{R}\rangle$$

$$e^{-\tau\hat{H}} \approx e^{-\tau\hat{T}}e^{-\tau\hat{V}} + \mathcal{O}(\tau^2) \qquad \text{Trotter breakup}$$

$$e^{-\tau\hat{H}} \approx e^{-\tau\hat{V}/2}e^{-\tau\hat{T}}e^{-\tau\hat{V}/2} + \mathcal{O}(\tau^3)$$

Since the potential energy is diagonal in position space, introducing 2 resolutions of the identity,

$$\langle \mathbf{R}'|e^{-\tau\hat{H}}|\mathbf{R}\rangle = e^{-\tau(V(\mathbf{R}')+V(\mathbf{R}))/2} \underbrace{\langle \mathbf{R}'|e^{-\tau\hat{T}}|\mathbf{R}\rangle}_{\text{next viewgr}} + \mathcal{O}(\tau^3)$$

Short-time Green's function – Kinetic term

Since the kinetic energy is diagonal in momentum space, we can evaluate $\langle \mathbf{R}' | e^{-\tau \hat{T}} | \mathbf{R} \rangle$ by introducing complete sets of momentum eigenstates

$$
\begin{aligned}
\langle \mathbf{R}' | e^{-\tau \hat{T}} | \mathbf{R} \rangle &= \int d\mathbf{P}'\, d\mathbf{P}\, \langle \mathbf{R}' | \mathbf{P}' \rangle \, \langle \mathbf{P}' | e^{-\tau \hat{T}} | \mathbf{P} \rangle \, \langle \mathbf{P} | \mathbf{R} \rangle \\
&= \frac{1}{(2\pi\hbar)^{3N}} \int d\mathbf{P}'\, d\mathbf{P}\, e^{\frac{-i\mathbf{P}'\cdot\mathbf{R}'}{\hbar}}\, \delta(\mathbf{P}' - \mathbf{P})\, e^{\frac{-\tau P^2}{2m}}\, e^{\frac{i\mathbf{P}\cdot\mathbf{R}}{\hbar}} \\
&= \frac{1}{(2\pi\hbar)^{3N}} \int d\mathbf{P}\, e^{\frac{i\mathbf{P}\cdot(\mathbf{R}-\mathbf{R}')}{\hbar}}\, e^{\frac{-\tau P^2}{2m}} \\
&= \frac{1}{(2\pi\hbar)^{3N}} e^{\frac{-m(\mathbf{R}-\mathbf{R}')^2}{2\hbar^2\tau}} \int d\mathbf{P}\, e^{\frac{\tau}{2m}(i\mathbf{P} + \frac{m}{\hbar\tau}(\mathbf{R}-\mathbf{R}'))^2} \\
&= \left(\frac{m}{2\pi\hbar^2\tau}\right)^{\frac{3N}{2}} e^{-\frac{m}{2\hbar^2\tau}(\mathbf{R}-\mathbf{R}')^2} \\
&= \frac{e^{-\frac{1}{2\tau}(\mathbf{R}-\mathbf{R}')^2}}{(2\pi\tau)^{\frac{3N}{2}}} \qquad 3N\text{-dim gaussian of width } \sqrt{\tau} \text{ in a.u.}
\end{aligned}
$$

Diffusion Monte Carlo – Short-time Green's function

Putting the two pieces together

$$G(\mathbf{R}',\mathbf{R},\tau) \;=\; \langle \mathbf{R}'|e^{\tau(E_{\rm T}-\hat{H})}|\mathbf{R}\rangle \approx \frac{1}{(2\pi\tau)^{3N/2}} e^{\left[-\frac{(\mathbf{R}'-\mathbf{R})^2}{2\tau}+\left\{E_{\rm T}-\frac{(v(\mathbf{R}')+v(\mathbf{R}))}{2}\right\}\tau\right]}$$

Diffusion Monte Carlo – Short-time Green's function

Can get the same result directly from the imaginary time Schrödinger Eq:

$$-\frac{1}{2}\nabla^2\psi(\mathbf{R},t) \;+\; (\mathcal{V}(\mathbf{R})-E_{\rm T})\psi(\mathbf{R},t) \;=\; -\frac{\partial\psi(\mathbf{R},t)}{\partial t}$$

Combining the diffusion Eq. and the rate Eq. Green's functions:

$$G(\mathbf{R}',\mathbf{R},\tau) \;\approx\; \frac{1}{(2\pi\tau)^{3N/2}}\, e^{\left[-\frac{(\mathbf{R}'-\mathbf{R})^2}{2\tau}+\left\{E_{\rm T}-\frac{(\mathcal{V}(\mathbf{R}')+\mathcal{V}(\mathbf{R}))}{2}\right\}\tau\right]}$$

The wavefunction,$\psi(\mathbf{R}',t+\tau)$, evolves according to the integral eqquation,

$$\psi(\mathbf{R}',t+\tau) \;=\; \int d\mathbf{R}\, G(\mathbf{R}',\mathbf{R},\tau)\psi(\mathbf{R},t).$$

Columns of $G(\mathbf{R}',\mathbf{R},\tau)$ are not normalized to 1, so weights and/or branching are needed.
The potential energy $\mathcal{V}$ can diverge to $\pm\infty$, so the fluctuations in the weights and/or population are huge!

Expectation values

There is an additional problem that the contribution that various MC points make to expectation values is proportional to $\Psi_{\rm T}(\mathbf{R})$:

$$
\begin{aligned}
E &= \frac{\int d\mathbf{R}\Psi_0(\mathbf{R})H(\mathbf{R})\Psi_{\rm T}(\mathbf{R})}{\int d\mathbf{R}\Psi_0(\mathbf{R})\Psi_{\rm T}(\mathbf{R})} \\
&\approx \frac{\sum_i^{N_{\rm MC}} H(\mathbf{R})\Psi_{\rm T}(\mathbf{R})}{\sum_i^{N_{\rm MC}} \Psi_{\rm T}(\mathbf{R})}
\end{aligned}
$$

This is inefficient for Bosonic systems, and is impossible for Fermionic systems since one gets 0/0.

The problems on previous viewgraph and this one are solved (at the price of biased expectation values) by using importance sampling and fixed-node boundary conditions with the approximate wavefunctions $\Psi_{\rm T}(\mathbf{R})$. In the limit that $\Psi_{\rm T} \to \Psi_0$ the weights of the walkers do not fluctuate at all and every MC point contributes equally to the expectation values.

In order to have finite variance, it is necessary that $\Psi_{\rm T}$ never be nonzero where $\Psi_{\rm G}$ is zero. In fact the usual practice in DMC is $\Psi_{\rm G} = \Psi_{\rm T}$ and so in this section we will not distinguish between them.

Diffusion Monte Carlo – Importance Sampled Green's Function

Importance sampling: Multiply imaginary-time the Schrödinger equation

$$-\frac{1}{2}\nabla^2\Psi(\mathbf{R},t) \;+\; (\mathcal{V}(\mathbf{R})-E_{\rm T})\Psi(\mathbf{R},t) \;=\; -\frac{\partial\Psi(\mathbf{R},t)}{\partial t}$$

by $\Psi_{\rm T}(\mathbf{R})$ and rearranging terms we obtain

$$-\frac{\nabla^2}{2}(\Psi\Psi_{\rm T}) \;+\; \nabla\cdot\left(\frac{\nabla\Psi_{\rm T}}{\Psi_{\rm T}}\Psi\Psi_{\rm T}\right) \;+\; \left(\underbrace{\frac{-\nabla^2\Psi_{\rm T}}{2\Psi_{\rm T}}+\mathcal{V}}_{E_{\rm L}(\mathbf{R})}-E_{\rm T}\right)(\Psi\Psi_{\rm T}) \;=\; -\frac{\partial(\Psi\Psi_{\rm T})}{\partial t}$$

defining $f(\mathbf{R},t)=\Psi(\mathbf{R},t)\Psi_{\rm T}(\mathbf{R})$, this is

$$\underbrace{-\frac{1}{2}\nabla^2 f}_{\text{diffusion}} \;+\; \underbrace{\nabla\cdot\left(\frac{\nabla\Psi_{\rm T}}{\Psi_{\rm T}}f\right)}_{\text{drift}} \;+\; \underbrace{(E_{\rm L}(\mathbf{R})-E_{\rm T})\,f}_{\text{growth/decay}} \;=\; -\frac{\partial f}{\partial t}$$

Since we know the exact Green function for any one term on LHS, an approximation is:

$$\tilde{G}(\mathbf{R}',\mathbf{R},\tau)\approx\frac{1}{(2\pi\tau)^{3N/2}}\,e^{\left[-\frac{(\mathbf{R}'-\mathbf{R}-\mathbf{V}\tau)^2}{2\tau}+\left\{E_{\rm T}-\frac{(E_{\rm L}(\mathbf{R}')+E_{\rm L}(\mathbf{R}))}{2}\right\}\tau\right]}$$

Diffusion Monte Carlo with Importance Sampling

Putting the drift, diffusion and reweighting Green's functions together,

$$\tilde{G}(\mathbf{R}',\mathbf{R},\tau) \approx \int d\mathbf{R}''\; G_{\rm rew}(\mathbf{R}',\frac{\tau}{2})\; G_{\rm dif}(\mathbf{R}',\mathbf{R}'',\tau)\; G_{\rm dri}(\mathbf{R}'',\mathbf{R},\tau)\; G_{\rm rew}(\mathbf{R}',\frac{\tau}{2})$$

$$\boxed{\tilde{G}(\mathbf{R}',\mathbf{R},\tau) \approx \frac{1}{(2\pi\tau)^{3N/2}} e^{\left[-\frac{(\mathbf{R}'-\mathbf{R}-\mathbf{V}\tau)^2}{2\tau}+\left\{E_{\rm T}-\frac{(E_{\rm L}(\mathbf{R}')+E_{\rm L}(\mathbf{R}))}{2}\right\}\tau\right]}}$$

The importance-sampled Green function has $E_{\rm L}(\mathbf{R})$ in the reweighting factor, which behaves MUCH better than the potential, $V(\mathbf{R})$. $V(\mathbf{R})$ diverges to $\pm\infty$ at particle coincidences whereas $E_{\rm L}(\mathbf{R})$ goes to a constant, E_0, as $\Psi_{\rm T} \to \Psi_0$. In addition it has a drift term that keeps the particles in the important regions, rather than relying on the reweighting to achieve that.

Even this does not always work. Why?

Diffusion Monte Carlo with Importance Sampling

Putting the drift, diffusion and reweighting Green's functions together,

$$\tilde{G}(\mathbf{R}',\mathbf{R},\tau) \approx \int d\mathbf{R}''\; G_{\rm rew}(\mathbf{R}',\frac{\tau}{2})\; G_{\rm dif}(\mathbf{R}',\mathbf{R}'',\tau)\; G_{\rm dri}(\mathbf{R}'',\mathbf{R},\tau)\; G_{\rm rew}(\mathbf{R}',\frac{\tau}{2})$$

$$\boxed{\tilde{G}(\mathbf{R}',\mathbf{R},\tau) \approx \frac{1}{(2\pi\tau)^{3N/2}} e^{\left[-\frac{(\mathbf{R}'-\mathbf{R}-\mathbf{V}\tau)^2}{2\tau}+\left\{E_{\rm T}-\frac{(E_{\rm L}(\mathbf{R}')+E_{\rm L}(\mathbf{R}))}{2}\right\}\tau\right]}}$$

The importance-sampled Green function has $E_{\rm L}(\mathbf{R})$ in the reweighting factor, which behaves MUCH better than the potential, $V(\mathbf{R})$. $V(\mathbf{R})$ diverges to $\pm\infty$ at particle coincidences whereas $E_{\rm L}(\mathbf{R})$ goes to a constant, E_0, as $\Psi_{\rm T} \to \Psi_0$. In addition it has a drift term that keeps the particles in the important regions, rather than relying on the reweighting to achieve that.

Even this does not always work. Why?
The above importance sampled Green function leads to an "infinite variance" estimate for systems other than Bosonic ground states!!

Singularities of Green's function

CJU, Nightingale, Runge, JCP 1993

Region	Local energy $E_{\rm L}$	Velocity $\mathbf{V}$
Nodes	$E_{\rm L} \sim \pm\frac{1}{R_\perp}$ for $\Psi_{\rm T}$ $E_{\rm L} = E_0$ for Ψ_0	$V \sim \frac{1}{R_\perp}$ for both $\Psi_{\rm T}$ and Ψ_0
e-n and e-e concidences	$E_{\rm L} \sim \frac{1}{x}$ if cusps not imposed $E_{\rm L}$ finite if cusps are imposed $E_{\rm L} = E_0$ for Ψ_0	V has a discontinuity for both $\Psi_{\rm T}$ and Ψ_0

All the above infinities and discontinuities cause problems, e.g.,

$$\int_0^a dx\, E_{\rm L} = \int_0^a dx \left(\frac{1}{x}\right) = \pm\infty$$

$$\int_0^a dx\, E_{\rm L}^2 = \int_0^a dx \left(\frac{1}{x}\right)^2 = \infty$$

Modify Green's function, by approximately integrating $E_{\rm L}$ and $\mathbf{V}$ over path, taking account of the singularities, at no additional computational cost.

Nonanalyticity of velocity near a node

CJU, Nightingale, Runge, JCP 1993

Linear approximation to $\Psi_{\rm T}$ (knowing $\mathbf{V} = \nabla\Psi_{\rm T}/\Psi_{\rm T}$):

$$\begin{aligned} \Psi_{\rm T}(\mathbf{R}') &= \Psi_{\rm T}(\mathbf{R}) + \nabla\Psi_{\rm T}(\mathbf{R}) \cdot (\mathbf{R}' - \mathbf{R}) \\ &\propto 1 + \mathbf{V} \cdot (\mathbf{R}' - \mathbf{R}) \end{aligned}$$

The average velocity over the time-step τ is:

$$\bar{\mathbf{V}} = \frac{-1 + \sqrt{1 + 2V^2\tau}}{V^2\tau}\mathbf{V} \rightarrow \begin{cases} \mathbf{V} & \text{if } V^2\tau \ll 1 \\ \sqrt{\frac{2}{\tau}}\hat{\mathbf{V}} & \text{if } V^2\tau \gg 1 \end{cases}$$

Infinite local energy near node

Make similar improvement to the growth/decay term of the Green's function by averaging of the local energy over time-step τ.

Discontinuity of velocity at particle coincidences

The e-N coincidence is more important than e-e coincidences because the wavefunction is larger in magnitude there.
Sample from linear combination of drifted Gaussians and exponential centered on nearest nucleus.

Infinite local energy near particle coincidences

Kato, Pure Appl. Math (1957), Pack and Byers-Brown, JCP, (1966), 2^{nd} order, Tew, JCP (2008)

Impose e-N and e-e cusp conditions on the wavefunction, so that divergence in potential energy is exactly canceled by divergence in kinetic energy.

$$\Psi = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} r^l \, f_{lm}(r) \, Y_l^m(\theta, \phi)$$

$$f_{lm}(r) \approx f_{lm}^0 \left[1 + \frac{q_i q_j \mu_{ij} r}{l+1} + O(r^2) \right]$$

with f_{lm}^0 being the first term in the expansion of $f_{lm}(r)$.

Familiar example: e-N cusp for s-state of Hydrogenic atom is $-Z$. e-e cusps are $1/2$ and $1/4$ for $\uparrow\downarrow$ and $\uparrow\uparrow$ respectively. (This is why we chose two of the parameters in the wavefunction in the lab to be -2 and $1/2$.)

Combining with Metropolis to reduce time-step error

Reynolds, Ceperley, Alder, Lester, JCP 1982

$$\underbrace{-\frac{1}{2}\nabla^2 f}_{\text{diffusion}} + \underbrace{\nabla\cdot\left(\frac{\nabla\psi_{\rm T}}{\psi_{\rm T}}f\right)}_{\text{drift}} + \underbrace{\left(E_{\rm L}(\mathbf{R})-E_{\rm T}\right)f}_{\text{growth/decay}} = -\frac{\partial f}{\partial t}$$

If we omit the growth/decay term then $|\Psi_{\rm T}|^2$ is the solution.
But we can sample $|\Psi_{\rm T}|^2$ exactly using Metropolis-Hastings! So, view $G(\mathbf{R}', \mathbf{R}, t)$ as being the proposal matrix $T(\mathbf{R}', \mathbf{R})$ and introduce accept-reject step after drift and diffusion steps.

Relationship of $\tilde{G}(\mathbf{R}', \mathbf{R}, \tau)$ to $G(\mathbf{R}', \mathbf{R}, \tau)$

$\tilde{G}(\mathbf{R}', \mathbf{R}, \tau)$ is a similarity transform of $G(\mathbf{R}', \mathbf{R}, \tau)$ by a diagonal matrix, as discussed when introducing importance sampling at beginning of this lecture in discrete notation.

$$\begin{aligned}
\Psi(\mathbf{R}', \tau) &= \int d\mathbf{R}\; G(\mathbf{R}', \mathbf{R}, \tau)\, \Psi(\mathbf{R}, 0) \\
\Psi_{\rm T}(\mathbf{R}')\Psi(\mathbf{R}', \tau) &= \Psi_{\rm T}(\mathbf{R}') \int d\mathbf{R}\; G(\mathbf{R}', \mathbf{R}, \tau)\, \Psi(\mathbf{R}, 0) \\
&= \int d\mathbf{R}\; \underbrace{\Psi_{\rm T}(\mathbf{R}') G(\mathbf{R}', \mathbf{R}, \tau) \frac{1}{\Psi_{\rm T}(\mathbf{R})}}_{\tilde{G}(\mathbf{R}', \mathbf{R}, \tau)}\; \Psi_{\rm T}(\mathbf{R})\Psi(\mathbf{R}, 0) \\
f(\mathbf{R}', \tau) &= \int d\mathbf{R}\; \tilde{G}(\mathbf{R}', \mathbf{R}, \tau)\, f(\mathbf{R}, 0) \\
\text{where} \quad \tilde{G}(\mathbf{R}', \mathbf{R}, \tau) &= \Psi_{\rm T}(\mathbf{R}') G(\mathbf{R}', \mathbf{R}, \tau) \frac{1}{\Psi_{\rm T}(\mathbf{R})}
\end{aligned}$$

Note: Both $G(\mathbf{R}', \mathbf{R}, \tau)$ and $\tilde{G}(\mathbf{R}', \mathbf{R}, \tau)$ can be for either free or fixed-node boundary conditions. The approximate $\tilde{G}(\mathbf{R}', \mathbf{R}, \tau)$ derived a few viewgraphs back is for fixed-node boundary conditions.

Question to think about before next lecture

In variational (Metropolis) Monte Carlo we have a single unweighted walker. (We could use many walkers but there is no reason to, and in fact it would be slightly less efficient because of the time spent equilibrating multiple walkers.)

In projector Monte Carlo methods, we use multiple weighted walkers.

Why?

Expectation values of operators

We wish to compute the pure (as opposed to mixed) expectation value

$$\langle A \rangle_{\rm pure} = \frac{\langle \Psi_0 | \hat{A} | \Psi_0 \rangle}{\langle \Psi_0 | \Psi_0 \rangle}$$

We consider various cases in order of increasing difficulty:

M.P. Nightingale, in Quantum Monte Carlo Methods in Physics and Chemistry, edited by M.P. Nightingale and CJU

1. $\hat{A}$ commutes with with $\hat{G}$ or equivalently $\hat{H}$ and is near-diagonal in chosen basis. (mixed expectation value)
2. $\hat{A}$ is diagonal in chosen basis. (forward/future walking) Liu, Kalos, and Chester, PRA (1974)
3. $\hat{A}$ is not diagonal in chosen basis, but, $A_{ij} \neq 0$ only when $G_{ij} \neq 0$. (forward/future walking)
4. $\hat{A}$ is not diagonal in chosen basis. (side walking) Barnett, Reynolds, Lester, JCP (1992)

Expectation values of operators

Factor the elements of the importance-sampled projector, $\tilde{G}(\mathbf{R}', \mathbf{R})$, as products of elements of a stochastic matrix/kernel (elements are nonnegative and elements of column sum to 1), $\tilde{T}(\mathbf{R}', \mathbf{R})$, and a reweight factor, $w(\mathbf{R}', \mathbf{R})$.

$$\tilde{G}(\mathbf{R}', \mathbf{R}) \;=\; \tilde{T}(\mathbf{R}', \mathbf{R}) w(\mathbf{R}', \mathbf{R})$$

In the case of DMC

$$\tilde{T}(\mathbf{R}', \mathbf{R}) \;=\; G_{\rm dif}(\mathbf{R}', \mathbf{R}'')\, G_{\rm drift}(\mathbf{R}'', \mathbf{R}) \;=\; \frac{1}{(2\pi\tau)^{3N/2}} e^{-\frac{(\mathbf{R}'-\mathbf{R}-\mathbf{V}\tau)^2}{2\tau}}$$

$$w(\mathbf{R}', \mathbf{R}) \;=\; e^{\left\{ E_{\rm T} - \frac{(E_{\rm L}(\mathbf{R}')+E_{\rm L}(\mathbf{R}))}{2} \right\}\tau}$$

For discrete state space and sparse H, define

$$\tilde{T}(\mathbf{R}', \mathbf{R}) \;=\; \frac{\tilde{G}(\mathbf{R}', \mathbf{R})}{\sum_{\mathbf{R}''} \tilde{G}(\mathbf{R}'', \mathbf{R})}$$

$$w(\mathbf{R}', \mathbf{R}) \;=\; w(\mathbf{R}) \;=\; \sum_{\mathbf{R}''} \tilde{G}(\mathbf{R}'', \mathbf{R})$$

1) $\hat{A}$ commutes with with $\hat{H}$ and is near-diagonal in chosen basis

By *near diagonal* we mean that either:

1. In discrete space $\hat{A}$ is sufficiently sparse that when walker is at state i, $A_{\mathrm{L},i} = \sum_j g_j A j i / g_i$ can be computed sufficiently quickly, or
2. In continuous space $\hat{A}$ has only local and local-derivative terms, e.g., $\frac{-1}{2}\sum_i \nabla_i^2 + V(\mathbf{R})$.

Since $\hat{A}$ commutes with with $\hat{H}$ the mixed estimator equals the pure estimator

$$\langle A \rangle_{\mathrm{mix}} = \frac{\langle \Psi_0 | \hat{A} | \Psi_{\mathrm{T}} \rangle}{\langle \Psi_0 | \Psi_{\mathrm{T}} \rangle} = \frac{\langle \Psi_0 | \hat{A} | \Psi_0 \rangle}{\langle \Psi_0 | \Psi_0 \rangle} = \langle A \rangle_{\mathrm{pure}}$$

1) $\hat{A}$ commutes with with $\hat{H}$ and is near-diagonal in chosen basis

$$
\begin{aligned}
\langle A \rangle &= \frac{\langle \Psi_{\rm T}|\hat{A}|\Psi_0\rangle}{\langle \Psi_{\rm T}|\Psi_0\rangle} = \frac{\langle \Psi_{\rm T}|\hat{A}G^p(\tau)|\Psi_{\rm T}\rangle}{\langle \Psi_{\rm T}|G^p(\tau)|\Psi_{\rm T}\rangle} \\
&= \frac{\sum_{\mathbf{R}_p\cdots\mathbf{R}_0} A\Psi_{\rm T}(\mathbf{R}_p)\left(\prod_{i=0}^{p-1} G(\mathbf{R}_{i+1},\mathbf{R}_i)\right)\Psi_{\rm T}(\mathbf{R}_0)}{\sum_{\mathbf{R}_p\cdots\mathbf{R}_0} \Psi_{\rm T}(\mathbf{R}_p)\left(\prod_{i=0}^{p-1} G(\mathbf{R}_{i+1},\mathbf{R}_i)\right)\Psi_{\rm T}(\mathbf{R}_0)} \\
&= \frac{\sum_{\mathbf{R}_p\cdots\mathbf{R}_0} \frac{A\Psi_{\rm T}(\mathbf{R}_p)}{\Psi_{\rm T}(\mathbf{R}_p)}\left(\prod_{i=0}^{p-1} \tilde{G}(\mathbf{R}_{i+1},\mathbf{R}_i)\right)(\Psi_{\rm T}(\mathbf{R}_0))^2}{\sum_{\mathbf{R}_p\cdots\mathbf{R}_0} \left(\prod_{i=0}^{p-1} \tilde{G}(\mathbf{R}_{i+1},\mathbf{R}_i)\right)(\Psi_{\rm T}(\mathbf{R}_0))^2} \\
&= \frac{\sum_{t=T_{\rm eq}+1}^{T_{\rm eq}+T} A_L(\mathbf{R}_t)W_t}{\sum_{t=T_{\rm eq}+1}^{T_{\rm eq}+T} W_t} \quad \text{since MC pts. from } \left(\prod_{i=0}^{p-1} \tilde{T}(\mathbf{R}_{i+1},\mathbf{R}_i)\right)(\Psi_{\rm T}(\mathbf{R}_0))^2
\end{aligned}
$$

$W_t = \prod_{i=0}^{p-1} w(\mathbf{R}_{t-i},\mathbf{R}_{t-i-1})$ or better $W_t = \prod_{i=0}^{T_{\rm eq}+t-1} w(\mathbf{R}_{T_{\rm eq}+t-i},\mathbf{R}_{T_{\rm eq}+t-i-1})$.

Branching (described later) is used to prevent inefficiency due wide disparity in weight products.

2) Expectation values of diagonal operators that do not commute with $\hat{H}$

DMC straightforwardly gives us

$$\langle A\rangle_{\rm mix} = \frac{\langle \Psi_0|\hat{A}|\Psi_{\rm T}\rangle}{\langle \Psi_0|\Psi_{\rm T}\rangle} = \frac{\int d\mathbf{R}\ \langle \Psi_0|\mathbf{R}\rangle\ \langle \mathbf{R}|\hat{A}|\mathbf{R}\rangle\ \langle \mathbf{R}|\Psi_{\rm T}\rangle}{\int d\mathbf{R}\ \langle \Psi_0|\mathbf{R}\rangle\ \langle \mathbf{R}|\Psi_{\rm T}\rangle} = \frac{\int d\mathbf{R}\ \Psi_0(\mathbf{R})\ A(R)\ \Psi_{\rm T}(\mathbf{R})}{\int d\mathbf{R}\ \Psi_0(\mathbf{R})\ \Psi_{\rm T}(\mathbf{R})}$$

but we want

$$\langle A\rangle_{\rm pure} = \frac{\langle \Psi_0|\hat{A}|\Psi_0\rangle}{\langle \Psi_0|\Psi_0\rangle} = \frac{\int d\mathbf{R}\ \langle \Psi_0|\mathbf{R}\rangle\ \langle \mathbf{R}|\hat{A}|\mathbf{R}\rangle\ \langle \mathbf{R}|\Psi_0\rangle}{\int d\mathbf{R}\ \langle \Psi_0|\mathbf{R}\rangle\ \langle \mathbf{R}|\Psi_0\rangle} = \frac{\int d\mathbf{R}\ \Psi_0(\mathbf{R})\ A(R)\ \Psi_0(\mathbf{R})}{\int d\mathbf{R}\ \Psi_0(\mathbf{R})\ \Psi_0(\mathbf{R})}$$

Two possibilities: Extrapolated estimator and forward walking

1) Extrapolated estimator

$$\begin{aligned}
\langle A\rangle_{\rm DMC} &= \langle A\rangle_{\rm pure} + \mathcal{O}(||\Psi_{\rm T}-\Psi_0||) \\
\langle A\rangle_{\rm VMC} &= \langle A\rangle_{\rm pure} + \mathcal{O}(||\Psi_{\rm T}-\Psi_0||) \\
2\langle A\rangle_{\rm DMC} - \langle A\rangle_{\rm VMC} &= \langle A\rangle_{\rm pure} + \mathcal{O}(||\Psi_{\rm T}-\Psi_0||)^2
\end{aligned}$$

2) Expectation values of diagonal operators that do not commute with $\hat{H}$

Forward or Future Walking

$$\begin{aligned}
\langle A \rangle &= \frac{\langle \Psi_{\rm T} | G^p(\tau) \hat{A} G^{p'}(\tau) | \Psi_{\rm T} \rangle}{\langle \Psi_{\rm T} | G^{p+p'}(\tau) | \Psi_{\rm T} \rangle} \\
&= \frac{\sum_{\mathbf{R}_{p+p'} \cdots \mathbf{R}_0} A(\mathbf{R}_{p'}) \left(\prod_{i=0}^{p+p'-1} \tilde{G}(\mathbf{R}_{i+1}, \mathbf{R}_i) \right) (\Psi_{\rm T}(\mathbf{R}_0))^2}{\sum_{\mathbf{R}_{p+p'} \cdots \mathbf{R}_0} \left(\prod_{i=0}^{p+p'-1} \tilde{G}(\mathbf{R}_{i+1}, \mathbf{R}_i) \right) (\Psi_{\rm T}(\mathbf{R}_0))^2} \\
&= \frac{\sum_{t=T_{\rm eq}+1}^{T_{\rm eq}+T} A(\mathbf{R}_t) W_{t+p}}{\sum_{t=T_{\rm eq}+1}^{T_{\rm eq}+T} W_{t+p}}
\end{aligned}$$

$W_{t+p} = \prod_{i=0}^{p+p'-1} w(\mathbf{R}_{t+p-i}, \mathbf{R}_{t+p-i-1})$ (product over p' past and p future) or better $W_{t+p} = \prod_{i=0}^{T_{\rm eq}+t+p-1} w(\mathbf{R}_{T_{\rm eq}+t+p-i}, \mathbf{R}_{T_{\rm eq}+t+p-i-1})$, (product over entire past and p future generations).

The contribution to the expectation value is: the local operator at time t, multiplied by the weight at a future time $t+p$. Need to store $A(\mathbf{R}_t)$ for p generations.

Usual tradeoff: If p is small, there is some residual bias since $\Psi_{\rm T}$ has not been fully projected onto Ψ_0, whereas, if p is large the fluctuations of the descendent weights increases the statistical noise. (Since we use branching, weight factors from past are not a problem.) For very large p all walkers will be descended from the same ancestor. (Mitochondrial Eve!)

$\hat{A}$ is not diagonal in chosen basis, but, $A_{ij} \neq 0$ only when $G_{ij} \neq 0$

Forward or Future Walking

$$
\begin{aligned}
\langle A \rangle &= \frac{\langle \Psi_T | G^{p-1}(\tau) \hat{A} G^{p'}(\tau) | \Psi_T \rangle}{\langle \Psi_T | G^{p+p'}(\tau) | \Psi_T \rangle} \\
&= \frac{\sum_{\mathbf{R}_{p+p'} \cdots \mathbf{R}_0} \left(\prod_{i=p'+1}^{p+p'-1} \tilde{G}(\mathbf{R}_{i+1}, \mathbf{R}_i) \right) \tilde{A}(\mathbf{R}_{p'+1}, \mathbf{R}_{p'}) \left(\prod_{i=0}^{p'-1} \tilde{G}(\mathbf{R}_{i+1}, \mathbf{R}_i) \right) (\Psi_T(\mathbf{R}_0))^2}{\sum_{\mathbf{R}_{p+p'} \cdots \mathbf{R}_0} \left(\prod_{i=0}^{p+p'-1} \tilde{G}(\mathbf{R}_{i+1}, \mathbf{R}_i) \right) (\Psi_T(\mathbf{R}_0))^2} \\
&= \frac{\sum_{t=T_{\rm eq}+1}^{T_{\rm eq}+T} W_{t+p-1,t+1} \; a(\mathbf{R}_{t+1}, \mathbf{R}_t) \; W_{t,t-p'}}{\sum_{t=T_{\rm eq}+1}^{T_{\rm eq}+T} W_{t+p}}
\end{aligned}
$$

$$
a(\mathbf{R}_{t+1}, \mathbf{R}_t) = \frac{\tilde{A}(\mathbf{R}_{t+1}, \mathbf{R}_t)}{\tilde{T}(\mathbf{R}_{t+1}, \mathbf{R}_t)} = \frac{A(\mathbf{R}_{t+1}, \mathbf{R}_t)}{T(\mathbf{R}_{t+1}, \mathbf{R}_t)}
$$

$$
W_{t_2,t_1} = \prod_{i=t_1}^{t_2-1} w(\mathbf{R}_{i+1}, \mathbf{R}_i)
$$

Again, the product of p' past weights can be replaced by products of weights over entire past.

$\hat{A}$ is not diagonal in chosen basis, and, $\exists$ some $A_{ij} \neq 0$ where $G_{ij} = 0$

Side Walking

Now it becomes necessary to have side walks that start from the backbone walk.

Just as we did for the importance-sampled projector, we factor $\tilde{A}$ into a Markov matrix and a reweighting factor.

The first transition of the side walk is made using this Markov matrix and and the rest of the side-walk using the usual Markov matrix.

The ends of the side-walks contribute to the expectation values.

This method is even more computationally expensive than forward walking, because one has to do an entire side walk long enough to project onto the ground state to get a single contribution to the expectation value.

Mixed and Growth Energy Estimators of the Energy

The *mixed estimator*, $E_{\rm mix}$, is obtained by averaging the configuration eigenvalue of the Hamiltonian, $\hat{H}$ (local energy).

The *growth estimator*, $E_{\rm gr}$, is obtained by averaging the configuration eigenvalue of the projector, $\hat{P}$, and then using the fact that $\hat{P}$ is a function of $\hat{H}$.

If the guiding function $\Psi_{\rm G}$ (used to importance sample $\hat{P}$) is the same as the trial wavefunction $\Psi_{\rm T}$ then the fluctuations of $E_{\rm mix}$ and $E_{\rm gr}$ will be perfectly correlated with each other and there is nothing to be gained by calculating both. (This statement assumes that one does not introduce unnecessary fluctuations in $E_{\rm gr}$ by, for example, integerizing the weights or killing walkers that cross nodes because of time-step error.) In that case the population control bias (to be discussed next) must be positive.

We note in passing that the FCIQMC method (Alavi group) uses the Hartree-Fock wavefunction for $\Psi_{\rm T}$ and the uniform wavefunction for $\Psi_{\rm G}$ and so there is not a strong correlation in the fluctuations of $E_{\rm mix}$ and $E_{\rm gr}$.

Reweighting, Branching, Population Control

1. Since columns of $G(\mathbf{R}', \mathbf{R}, t)$ are not normalized to unity, the "walkers" must carry weights or there must be population of walkers with branching (birth/death events), or, both.
2. If we have a single weighted walker, then a few generations of the walk will dominate and the computational effort expended on the rest of the walk would be largely wasted.
3. It is possible to have a fluctuating population size, with each walker having unit weight, but this leads to unnecessary birth/death events.
4. So, it is best to have a weighted population with approximately, but not exactly equal weights and birth/death events.
5. If left uncontrolled, the size of this population will fluctuate and may get larger than the computer memory or go to zero.
6. Hence, control the population, but this leads to a "population control error" that is inversely proportional to the target population size.
7. There exists a method for removing most of the population control error.

Population control error

The log of the weights of the generations will undergo a random walk and so some generations will have very small or very large weights. So, we have to exercise *population control* by adjusting the weights by a generation-dependent fluctuating factor f. If we are using the exponential projector, $\exp((E_{\rm T} - \hat{H})\tau)$, this is equivalent to adjusting $E_{\rm T}$, but for the purpose of removing the population control error it is better to think in terms of a fixed $E_{\rm T}$ and a fluctuating factor f that needs to be corrected for.
The naive biased estimator for the energy is

$$E_{\rm mix} \quad = \quad \frac{\sum_t^{N_{\rm gen}} \sum_i^{N_w(t)} w_i(t) E_{\rm L}(\mathbf{R}_i(t))}{\sum_t^{N_{\rm gen}} \sum_i^{N_w(t)} w_i(t)}$$

The population control error is proportional to the inverse of the target population size $N_{\rm walk}$. The error arises because of a negative correlation between the energy averaged over the generation and the weight of the generation. When the energy is low, the weight tends to be large and population control artificially reduces the weight and thereby creates a positive bias in the energy. Similarly, when the energy is high, the weight tends to be small and population control artificially increases the weight and this too creates a positive bias in the energy. Since the relative fluctuations in the energy and in the weight go as $1/\sqrt{N_{\rm walk}}$, the relative fluctuations in their covariance goes as $1/N_{\rm walk}$.

So, one way to reduce the population control error is to simply use a large population, and this is what most people do. If one wishes to be sure that the error is sufficiently small, plot the energy versus $1/N_{\rm walk}$ and take the limit $1/N_{\rm walk} \to 0$. But there exists a better way that allows us to estimate and remove most of the population control error within a single run, as described next.

Removing the population control error

Nightingale and Bloete, PRB 1986; CJU, Nightingale, Runge, JCP 1993

The basic idea for correcting the population control error is the following. When we do population control we have a generation-dependent factor by which we change the weights of all the walkers in that generation relative to what the mathematics tells us is correct. So, we keep track of these factors and when computing expectation values we undo these factors for the last several generations. If we undid them for the entire run then we would be back to our original problem, i.e., very large fluctuations in the weights. However, we only need to undo them for a number of generations corresponding to a few times the autocorrelation time to get rid of almost all of the population control bias. In the next viewgraph we explain how to do this and then in the following one we explain a continuous version of the algorithm that is even simpler to implement, though a bit harder to explain (which is why we do them in this order).

Removing the population control error

Nightingale and Bloete, PRB 1986; CJU, Nightingale, Runge, JCP 1993

The expectation value of $w_i(t)$ ought to be

$$E\left[w_i(t)\right] \;=\; \sum_j P_{ij} w_j(t-1) \quad \text{(stochastic equivalent of} \quad w_i(t) = \sum_j P_{ij} w_j(t-1) \;\;)$$

Introduce popul. control factor $f(t) = e^{\tau(E_T - E_{\rm est})} \left(\frac{W(t-1)}{W_{\rm target}}\right)^{1/g}$ & $P_i(t, T_p) = \prod_{p=0}^{T_p} f(t-p)$, where $E_{\rm est}$ is best current estim. of the energy, and, $W(t-1) = \sum_j w_j(t-1)$.

($e^{\tau(E_T - E_{\rm est})}$ compensates for an inaccurate E_T and $\left(\frac{W(t-1)}{W_{\rm target}}\right)^{1/g}$ tries to restore the population weight g generations later.)

Then,

$$E\left[f(t)\; w_i(t)\right] \;=\; \sum_j P_{ij}\; w_j(t-1) \qquad \text{or} \qquad E\left[P_i(t, T_p)\; w_i(t)\right] = \sum_j (P^{T_p})_{ij}\; w_j(t-T_p)$$

So the modified expressions for $E_{\rm mix}$ is:

$$E_{\rm mix} \;=\; \frac{\sum_t^{N_{\rm gen}} P_i(t, T_p) \sum_i^{N_w(t)} w_i(t) E_{\rm L}(\mathbf{R}_i(t))}{\sum_t^{N_{\rm gen}} P_i(t, T_p) \sum_i^{N_w(t)} w_i(t)}$$

This requires us to store a circular buffer of $T_{\rm p}$ population control factors $f(t)$ and iteratively compute two products (with $T_{\rm p}$ and $T_{\rm p+1}$ factors). ($E_{\rm mix}$ requires the first, $E_{\rm gr}$ requires both.)

Removing the population control error

Continuous method

A simpler procedure that does not require a circular buffer is to replace $P_i(t, T_{\rm p})$ by

$$P_i(t,p) = \cdots f(t-2)^{p^2}\ f(t-1)^{p^1}\ f(t)^{p^0} = \prod_{n=0}^{t-1} f(t-n)^{p^n}$$

with p a bit less than 1, calculated recursively at each generation using

$$P_i(t,p) = P_i(t-1,p)^p\ f(t),$$

$$E_{\rm mix} \quad = \quad \frac{\sum_t^{N_{\rm gen}} P_i(t,p)\ \sum_i^{N_w(t)} w_i(t) E_{\rm L}(\mathbf{R}_i(t))}{\sum_t^{N_{\rm gen}} P_i(t,p)\ \sum_i^{N_w(t)} w_i(t)}$$

A rough correspondence between $T_{\rm p}$ in the discrete method and p in the continuous method is established by requiring that $p^{T_{\rm p}} = 1/e$, i.e., $p = e^{-1/T_{\rm p}}$.

Removing the population control error

Continuous method A simpler procedure that does not require a circular buffer is to replace $P_i(t,\, T_{\rm p})$ by

$$P_i(t,p) \;=\; \cdots f(t-2)^{p^2}\; f(t-1)^{p^1}\; f(t)^{p^0} \;=\; \prod_{n=0}^{t-1} f(t-n)^{p^n}$$

$$\text{and} \quad \tilde{P}_i(t,p) \;=\; \cdots f(t-2)^{p^3}\; f(t-1)^{p^2}\; f(t)^{p^1} \;=\; \prod_{n=0}^{t-1} f(t-n)^{p^{n+1}}$$

with p a bit less than 1, calculated recursively at each generation using

$$P_i(t,p) = P_i(t-1,p)^p\; f(t), \qquad \tilde{P}_i(t,p) = \Big(\tilde{P}_i(t-1,p)\; f(t)\Big)^p$$

$$E_{\rm mix} \;=\; \frac{\sum_t^{N_{\rm gen}} P_i(t,p)\; \sum_i^{N_w(t)} w_i(t) E_{\rm L}(\mathbf{R}_i(t))}{\sum_t^{N_{\rm gen}} P_i(t,p)\; \sum_i^{N_w(t)} w_i(t)}$$

$$E_{\rm gr} \;=\; E_{\rm T} + \frac{1}{\tau}\left(1 - \frac{\sum_{t=1}^{N_{\rm gen}-1} P_i(t+1,p)\; \sum_i^{N_w(t+1)} w_i(t+1)}{\sum_{t=1}^{N_{\rm gen}-1} \tilde{P}_i(t,p)\; \sum_i^{N_w(t)} w_i(t)}\right), \qquad \text{for the linear projector}$$

$$E_{\rm gr} \;=\; E_{\rm T} - \frac{1}{\tau}\log\left(\frac{\sum_{t=1}^{N_{\rm gen}-1} P_i(t+1,p)\; \sum_i^{N_w(t+1)} w_i(t+1)}{\sum_{t=1}^{N_{\rm gen}-1} \tilde{P}_i(t,p)\; \sum_i^{N_w(t)} w_i(t)}\right), \qquad \text{for the exponential projector}$$

A rough correspondence between $T_{\rm p}$ in the discrete method and p in the continuous method is established by requiring that $p^{T_{\rm p}} = 1/e$, i.e., $p = e^{-1/T_{\rm p}}$.

Pure State versus Finite Temperature MC methods

So far we have talked about pure-state MC methods. We now give a very brief introduction to a finite-temperature MC method, the path-integral Monte Carlo (PIMC) method. Sometimes PIMC is used to approximate ground states, but that gets expensive since the length of the polymer grows as the inverse temperature. We also very briefly discuss the essence of reptation MC, which is in some sense a hybrid between PMC and PIMC. First we give schematics of the walks in VMC, PMC, PIMC and reptation MC.

Schematic of VMC and PMC

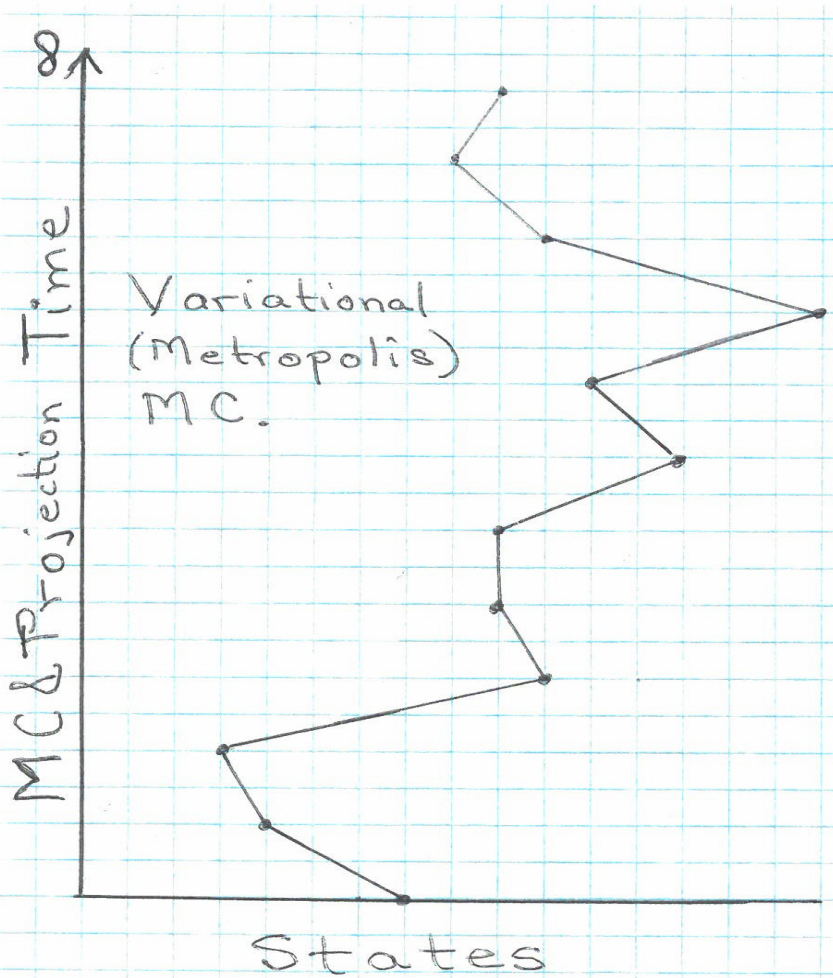

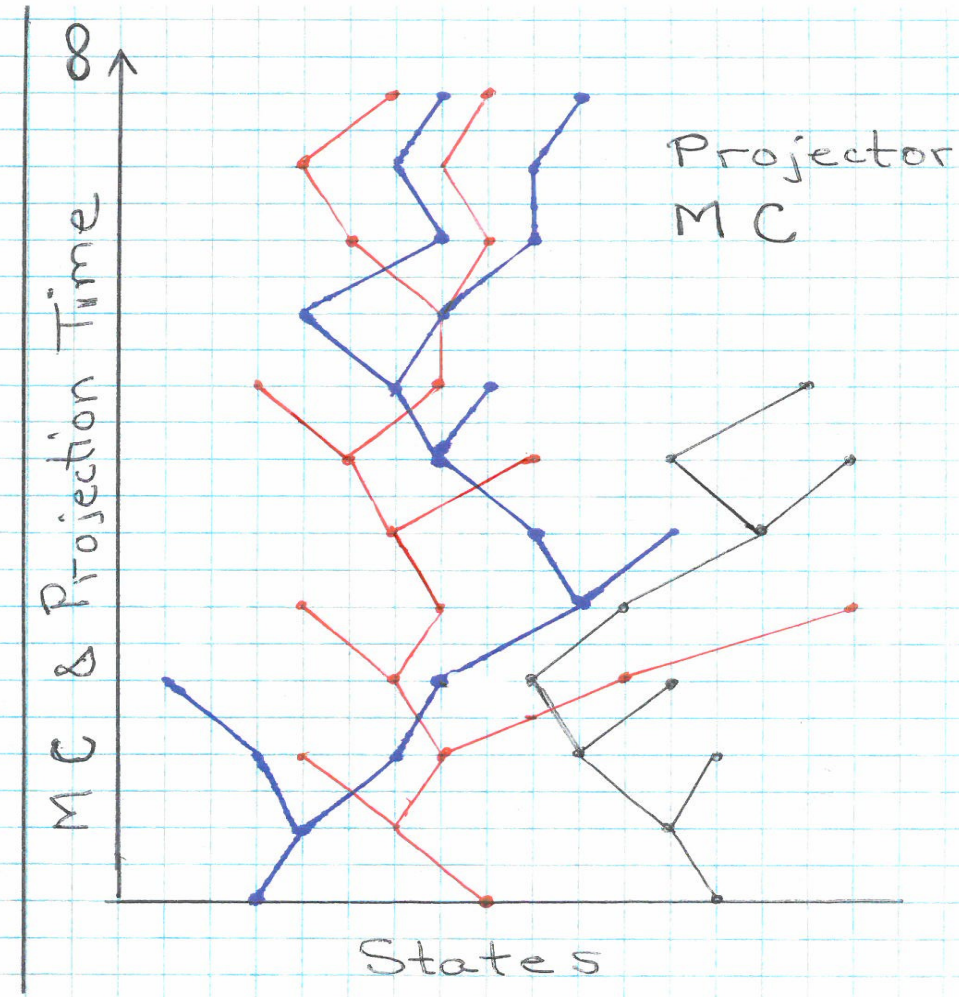

Schematic of PIMC and Reptation MC

Path-Integral M.C.

Average path length $\propto$ projection time/inverse T

Projection time $\neq$ M.C. time

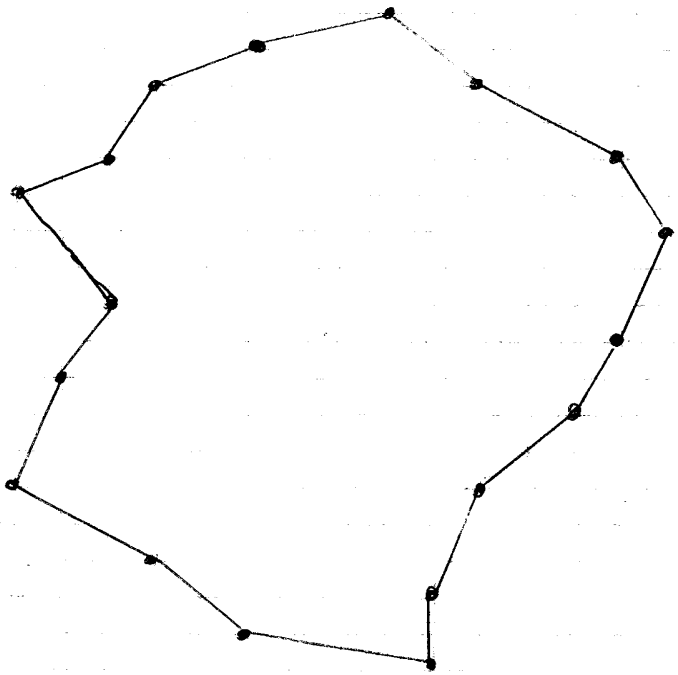

Reptation M.C.

Average path length $\propto$ projection time

Projection time $\neq$ MC time.

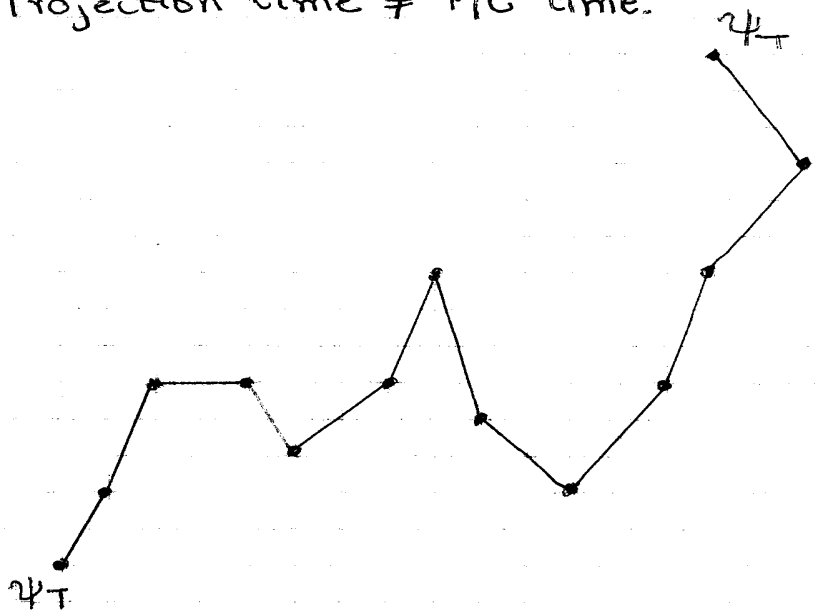

Schematic of VMC, PMC, PIMC and Reptation MC

1. Metropolis MC requires just a single unweighted walker, since the sum of the elements of a column of the Markov matrix add up to 1. PMC requires (for reasons of efficiency) a population of weighted walkers.
   (One could view the Markov matrix in Metropolis MC as a projector. Instead of projecting onto the ground state as in PMC, it instead projects onto the known guiding/trial state.)
2. In VMC and PMC the projection time and the MC time are the same. In PIMC and reptation MC, they are different – the projection time or inverse temperature is finite, whereas the MC time can of course be as large as patience permits.
3. The object being evolved has an extra time dimension in PIMC and reptation MC, compared to VMC and PMC, and so this adds to the computational cost. On the other hand one can compensate for this because the probability density is known, and so one can use Metropolis-Hastings and one has great freedom in the choice of the moves. (Algorithm is correct so long as detailed balance is satisfied.)
4. VMC, PMC and reptation MC are ground state methods, PIMC a finite-temperature method. People do use PIMC to approach the ground state but it becomes very inefficient since the length of the paths become very long.
5. VMC, PMC and reptation MC take advantage of accurate guiding/trial wavefunctions to greatly enhance their efficiency. PIMC does not.

Path-integral Monte Carlo, Finite Temperature

Detailed and excellent review: Ceperley, Rev. Mod. Phys. (1995)

In quantum statistical mechanics the expectation value of an operator A is given by

$$\langle A \rangle = \frac{\mathrm{tr}(\rho \mathrm{A})}{\mathrm{tr}(\rho)}$$

where $\hat{\rho} = e^{-\beta \hat{H}}$ is the quantum density matrix operator, $Z = \mathrm{tr}(\hat{\rho})$ is the partition function, $\hat{H}$ is the Hamiltonian and $\beta = 1/k_B T$.
The density operator in quantum statistical mechanics is identical to the quantum mechanical time evolution operator if we make the identification $t = -i\hbar\beta$.
If trace is evaluated in energy representation then energy is

$$E = \frac{\sum_{i=0}^{\infty} E_i e^{-\beta E_i}}{\sum_{i=0}^{\infty} e^{-\beta E_i}}$$

where E_i are the eigenvalues of $\hat{H}$. For many-body systems we do not know the energy eigenvalues or eigenvectors. PIMC provides a means of evaluating the trace in the coordinate representation.

Path-integral Monte Carlo, Finite Temperature

The expectation value of A in coordinate representation is

$$\langle A \rangle = \frac{\int d\mathbf{R}\, \langle \mathbf{R}|Ae^{-\beta H}|\mathbf{R}\rangle}{\int d\mathbf{R}\, \langle \mathbf{R}|e^{-\beta H}|\mathbf{R}\rangle}$$

Introducing a complete set of position eigenstates, we obtain,

$$\langle A \rangle = \frac{\int d\mathbf{R}\, d\mathbf{R}'\, \langle \mathbf{R}|A|\mathbf{R}'\rangle\langle \mathbf{R}'|e^{-\beta H}|\mathbf{R}\rangle}{\int d\mathbf{R}\, \langle \mathbf{R}|e^{-\beta H}|\mathbf{R}\rangle}$$

If A is diagonal in the coordinate representation then this reduces to

$$\langle A \rangle = \frac{\int d\mathbf{R}\, A(\mathbf{R})\, \rho(R,R,\beta)}{\int d\mathbf{R}\, \rho(R,R,\beta)} \tag{7}$$

where $\rho(R',R,\beta) \equiv \langle \mathbf{R}'|e^{-\beta H}|\mathbf{R}\rangle$ is the density matrix in coordinate representation. This is the same as in diffusion MC, but now imaginary time is interpreted as inverse temperature.

Eq. 7 can be evaluated by the Metropolis method since we derived an explicit, though approximate, expression for $\rho(R',R,\beta)$ when discussing diffusion MC.

$$\rho(R',R,\beta)$$

$$\propto \int d\mathbf{R}_1 \ldots d\mathbf{R}_{M-1}\; e^{-\tau\left[\left\{\left(\frac{\mathbf{R}'-\mathbf{R}_{M-1}}{2\tau}\right)^2+\left(\frac{\mathbf{R}_{M-1}-\mathbf{R}_{M-2}}{2\tau}\right)^2+\ldots+\left(\frac{\mathbf{R}_1-\mathbf{R}}{2\tau}\right)^2\right\}\;+\;\left(\frac{V(\mathbf{R}')}{2}+V(\mathbf{R}_{M-1})+\ldots+V(\mathbf{R}_1)+\frac{V(\mathbf{R})}{2}\right)\right]}$$

Path-integral Ground State (PIGS) / Reptation Monte Carlo

A hybrid between DMC and PIMC

Baroni and Moroni, in NATO book, ed. by M.P, Nightingale and CJU, (1999)

$$\langle A \rangle = \frac{\int d\mathbf{R}\, d\mathbf{R}'\, d\mathbf{R}''\, \langle \Psi_{\rm T}|\mathbf{R}\rangle\, \langle \mathbf{R}|e^{-t\hat{H}}|\mathbf{R}'\rangle\, \langle \mathbf{R}'|A|\mathbf{R}'\rangle\, \langle \mathbf{R}'|e^{-t\hat{H}}|\mathbf{R}''\rangle\, \langle \mathbf{R}''|\Psi_{\rm T}\rangle}{\int d\mathbf{R}\, d\mathbf{R}'\, d\mathbf{R}''\, \langle \Psi_{\rm T}|\mathbf{R}\rangle\, \langle \mathbf{R}|e^{-t\hat{H}}|\mathbf{R}'\rangle\, \langle \mathbf{R}'|e^{-t\hat{H}}|\mathbf{R}''\rangle\, \langle \mathbf{R}''|\Psi_{\rm T}\rangle}$$

Compared to PIMC – Instead of having a closed polymer, which needs to be very long in order to get the ground state have an open polymer, terminated by $\Psi_{\rm T}$. As $\Psi_{\rm T}$ gets better, the length of open polymer can get shorter.

Compared to DMC – Tradeoff between having the complication of moving an entire polymer versus the freedom of using clever Metropolis moves. Less efficient for the energy, but, more efficient for operators that do not commute with $\hat{H}$ if extrapolated estimators are not accurate enough.

It is possible to importance-sample the above, in the same way as in DMC.

Some topics in classical MC and QMC we did not discuss

1. Cluster algorithms – Swendsen-Wang and Wolff
2. Multiple-try Metropolis
3. Multilevel sampling
4. Correlated sampling, umbrella sampling, Wang-Landau
5. Lattice-regularized DMC (Sorella, Casula)
6. Nonlocal pseudopotentials in QMC (Fahy; Mitas, Shirley, Ceperley; Casula)
7. Extended systems (periodic BC, wavefns, finite-size errors, ...) (Foulkes, Needs, ...)
8. Penalty method and coupled electron-ion MC (Ceperley, Dewing, Pierleoni)
9. Zero-variance zero bias (ZVZB) method (Assaraf and Caffarel)
10. Extension of fixed-node to fixed-phase method (Bolton; Ortiz, Jones, Ceperley)
11. Domain Green's function MC. (Kalos)
12. Stochastic series expansion (Sandvik)
13. Loop and worm algorithms
14. Diagrammatic Monte Carlo
15. Impurity solvers for DMRG