COMPUTATIONAL QUANTUM MAGNETISM

Boulder 2003 Summer School

Lecture II- World Line Quantum Monte Carlo

A. Introduction

In this lecture, I will describe the "World–Line" Quantum Monte Carlo algorithm for the Heisenberg and Hubbard models. This method employs path integrals to map the partition function of a quantum mechanical Hamiltonian onto a classical problem in one higher dimension. As we shall see, the "length" of this new "imaginary time" dimension is given by the inverse temperature $\beta = 1/T$. The classical degrees of freedom are the eigenvalues of the original quantum operators as a function of space and imaginary time. One of the attractive features of the approach is that the world–lines trace variables which are associated with the operators in the original quantum Hamiltonian, and therefore offer an intuitive real space picture of the correlations in the system. In contrast, the determinant Monte Carlo algorithm, which is the focus of lecture III, also begins with a mapping of a quantum problem onto a classical problem in one higher dimension, but the resulting classical degrees of freedom are less directly related to the original operators in the Hamiltonian.

The background reading includes a chapter which has a description of the world line algorithm for a single quantum oscillator and for the one dimensional Ising model in a transverse field. These are among the most simple applications of the world-line method. Here, after introducing the general idea of the technique, we will jump directly to the more difficult problem of the Heisenberg Hamiltonian, and its generalizations. Students desiring additional simple examples should examine the background reading.

Before beginning, I should acknowledge that world-line QMC has seen enormous developments over the last decade. 'Loop' algorithms have helped solve the problem of long autocorrelation times by using cleverly constructed non-local monte carlo moves. 'Continuous time' algorithms remove errors associated with the finite inverse temperature discretization mesh. 'Worm' algorithms have been proposed to measure the off-diagonal Greens function. Finally, the 'Stochastic Series Expansion' appears to be especially effective (partly because it incorporates ideas from some of these other approaches). It is beyond the scope of this lecture to describe these developments. However, the student intending serious work in world-line QMC would be well advised to learn them.

<u>B. The General Idea of World-Line QMC</u>

The starting point of QMC is the partition function

$$Z = \operatorname{Tr} e^{-\beta H}.$$
 (1)

Because H contains terms which do not commute, this exponential is difficult to work with.

The world-line approach begins by dividing H into pieces $H = H_1 + H_2 + H_3 + \ldots$ such that, in an appropriate basis, the matrix elements of the exponential of each individual H_{α} can be evaluated. To isolate the H_{α} from each other, one discretizes β into small increments $\beta = L\Delta\tau$ and approximates,

$$Z = \operatorname{Tr} e^{-\beta H} = \operatorname{Tr} [e^{-\Delta \tau H}]^{L} \approx \operatorname{Tr} [e^{-\Delta \tau H_{1}} e^{-\Delta \tau H_{2}} e^{-\Delta \tau H_{3}} \dots]^{L}.$$
 (2)

The errors associated with this 'Suzuki-Trotter' approximation can be systematically eliminated by increasing L (reducing $\Delta \tau$). Complete sets of states in which the exponentials of H_{α} are calculable are introduced in between all the exponentials, and the matrix elements evaluated. The partition function Z is thereby reduced to the calculation of the sum over all these complete sets of states of the product of the *numbers* which result from the computation of the matrix elements. Finally, standard classical monte carlo methods are employed to evaluate the sum. A lightning review of classical monte carlo is contained in the appendix.

C. World Line QMC for the Spin–1/2 XXZ Hamiltonian

Because it makes the discussion no more complex, we consider a generalization of the Heisenberg model, the quantum XXZ Hamiltonian, which has richer physics. We will restrict ourselves to one dimension, because that simplifies our discussion considerably.

$$H = J \sum_{i} [(S_{x,i}S_{x,i+1} + S_{y,i}S_{y,i+1}) + \lambda S_{z,i}S_{z,i+1}] - B_z \sum_{i} S_{z,i}$$

= $J \sum_{i} [\frac{1}{2}(S_{+,i}S_{-,i+1} + S_{-,i}S_{+,i+1}) + \lambda S_{z,i}S_{z,i+1}] - B_z \sum_{i} S_{z,i}$ (3)

As before, $S_{x,i}, S_{y,i}$ and $S_{z,i}$ are quantum spin-1/2 operators at each site *i*, and $S_{+,i}, S_{-,i}$ are the associated raising and lowering operators. B_z is an external field in the *z* direction. When $|\lambda| > 1$ this model is in the Ising universality class and exhibits a finite temperature phase transition in zero field to a state with long range order, in dimensions greater than d = 1. For $|\lambda| < 1$, the model is in the XY universality class, with a finite temperature phase transition in zero field to a state with long range order in dimensions greater than d = 2. A Kosterlitz-Thouless phase transition to a state with spin correlations which decay with a power-law in d = 2 occurs at finite T. The isotropic Heisenberg point $\lambda = 1$ has long range order in d = 2 only at T = 0.

QMC methods played a crucial role in proving that long range order exists in the ground state of the 2-d spin-1/2 Heisenberg model. The spin-1/2 XXZ Hamiltonian is isomorphic to a lattice system of hard-core bosons with a near neighbor interaction. Recent experiments on "Mott" transitions in optically trapped atoms has led to quite a bit of recent simulation work on interacting Bose systems.

A convenient choice of H_{α} in the case of the XXZ Hamiltonian is the "checkerboard" decomposition, illustrated in Fig. 1,

$$H_1 = J \sum_{i \text{ odd}} \left[\frac{1}{2} (S_{+,i} S_{-,i+1} + S_{-,i} S_{+,i+1}) + \lambda S_{z,i} S_{z,i+1} \right] - \frac{1}{2} B_z \sum_{i} S_{z,i}$$

FIG. 1: The division of the one-dimensional XXZ Hamiltonian into H_1 and H_2 is illustrated for an N = 8 site, 1-d lattice. H_1 consists of independent two-site pieces linking pairs of sites $(1,2)(3,4)(5,6)\ldots$ H_2 consists of independent two-site pieces linking pairs of sites $(2,3)(4,5)(6,7)\ldots$ Periodic boundary conditions in the spatial direction connect the last site, i = 8to the first, i = 1.

To see why this is a useful division, introduce complete sets of states which are eigenstates of the z component of spin at each imaginary time slice. The terms in H_{α} involving the z components of spin act on these states and are immediately converted to numbers. The partition function becomes,

$$Z = \sum_{\{S_{z,i,l}\}} \exp[-\lambda J \Delta \tau \sum_{i,l} S_{z,i,l} S_{z,i+1,l}]$$

$$\langle S_{z,1,1} S_{z,2,1} \dots S_{z,N,1} | e^{-\Delta \tau H_{ex,1}} | S_{z,1,2} S_{z,2,2} \dots S_{z,N,2} \rangle$$

$$\langle S_{z,1,2} S_{z,2,2} \dots S_{z,N,2} | e^{-\Delta \tau H_{ex,2}} | S_{z,1,3} S_{z,2,3} \dots S_{z,N,3} \rangle$$

$$\langle S_{z,1,2} S_{z,2,2} \dots S_{z,N,2} | e^{-\Delta \tau H_{ex,1}} | S_{z,1,3} S_{z,2,3} \dots S_{z,N,3} \rangle$$

$$\cdots$$

$$\langle S_{z,1,2L} S_{z,2,2L} \dots S_{z,N,2L} | e^{-\Delta \tau H_{ex,2}} | S_{z,1,1} S_{z,2,1} \dots S_{z,N,1} \rangle, \qquad (5)$$

with the abbreviation $H_{ex,\alpha}$ for the exchange terms in the Hamiltonian H_{α} . How come the z components of spin $S_{z,i,l}$ have two indices now? The first index is the usual spatial site i = 1, 2, ..., N. The second index is the 'imaginary time' l = 1, 2, ..., 2L which labels the complete state inserted at the particular point l in the string of exponentials. The appearance of this extra index l is the technical realization of the statement that a d dimensional quantum problem is mapped into a d + 1 dimensional classical one. In other words, the partition function is a sum over a set of variables $\{S_{z,i,l}\}$ which lives in 2-d (if our original quantum problem was in 1-d). Note that we have to introduce 2L time slices in this checkerboard decomposition in one dimension. The factor L comes from the discretization of β , while the factor of 2 is a consequence of the number of pieces into which H was divided.

To see that the matrix elements are calculable, the crucial observation is that H_1 and H_2 consist of independent two-site pieces. Thus the matrix elements factorize,

$$\langle S_{z,1,l} S_{z,2,l} \dots S_{z,N,l} | e^{-\Delta \tau H_{\text{ex},\alpha}} | S_{z,1,l+1} S_{z,2,l+1} \dots S_{z,N,l+1} \rangle$$

= $\prod_{i \text{ odd}} \langle S_{z,i,l} S_{z,i+1,l} | e^{-\Delta \tau \frac{1}{2} J(S_{+,i} S_{-,i+1} + S_{-,i} S_{+,i+1})} | S_{z,i,l+1} S_{z,i+1,l+1} \rangle,$ (6)

and one must only evaluate the corresponding two site matrix elements. Using the same notation as in lecture one, the results are,

$$\langle + + | e^{-\Delta \tau \frac{1}{2}(S_{+,1}S_{-,2}+S_{-,1}S_{+,2})} | + + \rangle = 1 \langle - - | e^{-\Delta \tau \frac{1}{2}(S_{+,1}S_{-,2}+S_{-,1}S_{+,2})} | - - \rangle = 1, \langle + - | e^{-\Delta \tau \frac{1}{2}(S_{+,1}S_{-,2}+S_{-,1}S_{+,2})} | + - \rangle = \cosh(\frac{J}{2}\Delta\tau), \langle - + | e^{-\Delta \tau \frac{1}{2}(S_{+,1}S_{-,2}+S_{-,1}S_{+,2})} | - + \rangle = \cosh(\frac{J}{2}\Delta\tau), \langle + - | e^{-\Delta \tau \frac{1}{2}(S_{+,1}S_{-,2}+S_{-,1}S_{+,2})} | - + \rangle = -\sinh(\frac{J}{2}\Delta\tau), \langle - + | e^{-\Delta \tau \frac{1}{2}(S_{+,1}S_{-,2}+S_{-,1}S_{+,2})} | + - \rangle = -\sinh(\frac{J}{2}\Delta\tau),$$

$$\langle - + | e^{-\Delta \tau \frac{1}{2}(S_{+,1}S_{-,2}+S_{-,1}S_{+,2})} | + - \rangle = -\sinh(\frac{J}{2}\Delta\tau),$$

$$\langle - + | e^{-\Delta \tau \frac{1}{2}(S_{+,1}S_{-,2}+S_{-,1}S_{+,2})} | + - \rangle = -\sinh(\frac{J}{2}\Delta\tau),$$

$$\langle - + | e^{-\Delta \tau \frac{1}{2}(S_{+,1}S_{-,2}+S_{-,1}S_{+,2})} | + - \rangle = -\sinh(\frac{J}{2}\Delta\tau),$$

$$\langle - + | e^{-\Delta \tau \frac{1}{2}(S_{+,1}S_{-,2}+S_{-,1}S_{+,2})} | + - \rangle = -\sinh(\frac{J}{2}\Delta\tau),$$

$$\langle - + | e^{-\Delta \tau \frac{1}{2}(S_{+,1}S_{-,2}+S_{-,1}S_{+,2})} | + - \rangle = -\sinh(\frac{J}{2}\Delta\tau),$$

$$\langle - + | e^{-\Delta \tau \frac{1}{2}(S_{+,1}S_{-,2}+S_{-,1}S_{+,2})} | + - \rangle = -\sinh(\frac{J}{2}\Delta\tau),$$

$$\langle - + | e^{-\Delta \tau \frac{1}{2}(S_{+,1}S_{-,2}+S_{-,1}S_{+,2})} | + - \rangle = -\sinh(\frac{J}{2}\Delta\tau),$$

$$\langle - + | e^{-\Delta \tau \frac{1}{2}(S_{+,1}S_{-,2}+S_{-,1}S_{+,2})} | + - \rangle = -\sinh(\frac{J}{2}\Delta\tau),$$

$$\langle - + | e^{-\Delta \tau \frac{1}{2}(S_{+,1}S_{-,2}+S_{-,1}S_{+,2})} | + - \rangle = -\sinh(\frac{J}{2}\Delta\tau),$$

$$\langle - + | e^{-\Delta \tau \frac{1}{2}(S_{+,1}S_{-,2}+S_{-,1}S_{+,2})} | + - \rangle = -\sinh(\frac{J}{2}\Delta\tau),$$

with all other matrix elements zero. The only non-zero matrix elements are those between states which have the same $S_{z,tot} = S_{z,1} + S_{z,2}$ in the two time slices. This is a consequence of the fact that $[S_{+,1}S_{-,2} + S_{-1}, S_{+,2}, S_{z,1} + S_{z,2}] = 0$.

Let's summarize. The partition function of the quantum XXZ model in 1-d has been expressed as a sum over classical degrees of freedom, $\{S_{z,i,l}\}$ on a 2-d space time lattice,

$$Z = \sum_{\{S_{z,i,l}\}} W[\{S_{z,i,l}\}]$$
(8)

The summand W consists of two pieces: an exponential of products of pairs of $S_{z,i,l}$ on adjacent spatial sites *i* for each of the imaginary time slices *l*, and a product of matrix elements.

Since the variables $\{S_{z,i,l}\}$ are just real numbers, as are the weights $W[\{S_{z,i,l}\}]$, we may follow the usual (classical) monte carlo procedure for generating configurations of $\{S_{z,i,l}\}$. This consists of suggesting changes to $\{S'_{z,i,l}\}$ which are then accepted with probability, $p = \min[1, W(\{S'_{z,i,l}\}/W(\{S_{z,i,l}\}))$. This 'Metropolis algorithm' yields configurations $\{S_{z,i,l}\}$ with probability proportional to $W(\{S_{z,i,l}\})$. (See appendix.)

The off-diagonal matrix elements are negative if the XY coupling is antiferromagnetic, J > 0. This is an instance of the "sign problem" in quantum monte carlo and reflects the conceptual difficulty encountered if the weights $W(\{S_{z,i,l}\})$ can be negative. If the XY coupling is ferromagnetic, J < 0, then there is no sign problem. In the case of a near-neighbor antiferromagnetic coupling on a bipartite lattice, we can eliminate the sign problem by rotating the spin operators on one sublattice. This changes the sign of J (but not the sign of the λJ term), yet does not alter the physics since it simply corresponds to a different choice of direction of the local axes of spin quantization.

If the signs cannot be eliminated by such a rotation or some other means, then the world– line method will not work for the Hamiltonian in question. Examples are antiferromagnetic models on non–bipartite lattices like the triangular lattice, or models with longer range antiferromagnetic coupling, such as the two–dimensional, square lattice, " J_1-J_2 " Heisenberg model which has a next near–neighbor antiferromagnetic coupling J_2 across the diagonal of a square. This sign problem is the most fundamental limitation to quantum simulation techniques. It arises not only in the world–line algorithm, but also in the determinant approach, in ground state projection methods, etc.

It is useful to picture the structure of the checkerboard break-up by drawing the (1+1)dimensional array of spins and shading the squares corresponding to bonds across which a piece of the Hamiltonian acts. Thus H_1 connects spatial sites $(1, 2), (3, 4), \ldots (N-1, N)$ from time slice l to l+1 where l is odd, and likewise H_2 connects spatial sites $(2, 3), (4, 5), \ldots (N, 1)$ from time slice l to l+1 where l is even. (The last pair (N, 1) is present if we have periodic boundary conditions which connect the first and last sites in our Hamiltonian.) The resulting "checkerboard lattice" is illustrated in Fig. 2.



FIG. 2: The "checkerboard" which arises by shading those plaquettes of the space-imaginary time lattice across which time evolution operators act. In the figure, the number of spatial sites N = 8and number of imaginary time slices 2L = 6. Periodic boundary conditions connect sites and slices at the edges of the lattice, illustrated here by the identification of the boundaries x = 0 and $\tau = 0$ with x = 8 and $\tau = 6$ respectively. World lines can traverse only diagonals of shaded squares. A typical monte carlo move which pulls a world-line across an unshaded square is shown.

That the matrix elements vanish unless the sum of the z components of spin at the top and bottom of a shaded square are equal puts constraints on the terms in the sum over $S_{z,i,l}$ which contribute to the partition function. These conservation laws can best be visualized as follows: Draw lines connecting the sites on which the z component of spin is positive. Since the number of up spin sites is conserved from time slice to time slice, the result is a set of continuous "world–lines". The world–lines can cross only the shaded squares of the checkerboard lattice, since it is on these that the Hamiltonian acts. The periodic boundary conditions inherent in the trace require that these world–lines also connect continuously from the last time slice, l = 2L, to the first, l = 1. This same construction occurs in the monte carlo path integral formulation for problems of interacting bosons and fermions where the conservation law is particle number, and the name "world–line" is especially appropriate. We need to formulate monte carlo moves which respect the restricted class of spin configurations which is allowed. A move which flips a single spin $S_{z,i,l} \rightarrow -S_{z,i,l}$ would always give rise to zero matrix elements if the original configuration had non-zero Boltzmann weight. Put differently, it would "break" a world-line. It is easy to see (Fig. 2) that moves which "pull" a world-line across an *unshaded* square of the checkerboard lattice preserve all the local conservation laws and result in configurations of nonzero weight, assuming the original configuration was allowed. Four spin variables are changed in such an update, and the values of the matrix elements on four shaded plaquettes are modified. This means that the decision making process is local, and updating all the degrees of freedom on the lattice scales linearly with the lattice size. An acceptance-rejection step using the Metropolis algorithm in which the move is accepted with probability max(1,R), where R is the ratio of the product of new to old values of the matrix elements on the four modified plaquettes, along with the exponentials from the S_z operators.

It is useful explicitly to write down $W(\{S_{z,i,l}\})$, $W(\{S'_{z,i,l}\})$, and R for the two configurations of Fig. 2. There are 24 matrix elements, one for each of the shaded plaquettes, and a the term from the S_z operators obtained by counting up the pairs of parallel spins on each time slice. Before pulling the line,

$$W(\{S_{z,i,l}\}) = [\sinh(\frac{J}{2}\Delta\tau)]^6 \left[\cosh(\frac{J}{2}\Delta\tau)\right]^{18} e^{-8\lambda J\Delta\tau} e^{-12B\Delta\tau}$$
(9)

Meanwhile,

$$W(\{S'_{z,i,l}\}) = [\sinh(\frac{J}{2}\Delta\tau)]^4 \left[\cosh(\frac{J}{2}\Delta\tau)\right]^{20} e^{-12B\Delta\tau}$$
(10)

The ratio of these weights is

$$R = \left[\coth(\frac{J}{2}\Delta\tau)\right]^2 e^{+8\lambda J\Delta\tau} \tag{11}$$

It is favorable to perform the move as far as the matrix elements of H_{ex} are concerned, since their product is increased. The move also increases the term arising from the coupling in the z direction, as long as it is antiferromagnetic ($\lambda J > 0$). The magnetic field B_z does not affect R because the two configurations have the same total magnetization.

This world-line algorithm for one-dimensional quantum spin-1/2 systems can easily be generalized to higher dimension. There are different possibilities for dividing up the Hamiltonian. For d = 2, one might break H into four pieces corresponding to odd and even links in each of the x and y directions of the spatial lattice. This requires 4L intermediate states, where $L = \beta/\Delta\tau$. The allowed spin configurations are even more restricted, and not completely trivial to visualize. Pulling world-lines across an unshaded plaquette in either the x or y directions are still allowed moves. These will now change the values of eight spins and eight of the plaquettes of the lattice. However, these moves do not exhaust the full phase space. Moves which introduce a local twist of the lines are also permitted, and should be included to ensure ergodicity. Indeed, the density of these local twists has been used to characterize the various phases of the Hamiltonian. It is also possible to break H into only two pieces, each consisting of independent four site terms. The resulting matrix elements are a bit more complicated to evaluate, but the variance in the QMC is reduced since more of the sum is done analytically. Also, this approach involves introducing only half as many time slices (2L), and provides an easier means to keep track of local twists. This second method is also useful in the study of models with "ring exchange" terms, which have recently been proposed.

It is possible to introduce variance reduction techniques into world-line monte carlo. Rather than suggesting a move which would pull a line across an unshaded plaquette each time that plaquette is encountered in sweeping through the lattice, one can modify the suggestion probability based on the configuration of spins on neighboring sites. In general this is done according to some insight into the expected correlations, reducing the rate of suggestion of moves whose acceptance would violate expected structures. Of course, any such modification of the suggestion probability must also be appropriately accounted for in the Metropolis acceptance/rejection step so as to preserve detailed balance. Thus the algorithm always remains exact, to within Trotter errors, and only the variance and equilibration are affected.

We can now summarize the procedure for generating configurations in a world–line simulation of the quantum spin-1/2 XXZ Hamiltonian: Initialize the lattice to a configuration which respects local conservation laws. Then suggest moves, like pulling world lines across unshaded plaquettes, which result in allowed configurations, and accept or reject them using the Metropolis algorithm.

Snapshots of the spin configurations in the course of the simulation of the XXZ Hamiltonian are given in Fig. 3 for J = 1 and $\lambda = 0, 1, 2$. These provide intuitive pictures of the underlying correlations. For example, we can see the antiferromagnetic spin order in the z direction build up as λ increases.

While, these snapshots provide qualitative pictures of the physics, we need to describe how to measure operator expectation values quantitatively. It turns out there are constraints on what one can calculate. This is a serious drawback of the world–line algorithm, especially in comparison with the determinant approach. Consider evaluating,

$$\langle \hat{A} \rangle = Z^{-1} \operatorname{Tr}[\hat{A}e^{-\beta H}].$$
(12)

If the operator \hat{A} is diagonal in the basis of complete sets of intermediate states, the procedure is simple. \hat{A} acts on the state $\langle S_{z,1,1} S_{z,2,1} \dots S_{z,N,1} |$ yielding a number $a[\{S_{z,i,1},\}]$ without altering the state. The operator $e^{-\beta H}$ then generates in the numerator the same sequence of matrix elements as in the partition function in the denominator. In other words,

$$\langle A \rangle = \frac{\sum_{\{S_{z,i,l}\}} W[\{S_{z,i,l}\}]a[\{S_{z,i,l}\}]}{\sum_{\{S_{z,i,l}\}} W[\{S_{z,i,l}\}]}.$$
(13)

Thus if configurations are generated with weight proportional to W, one simply accumulates the numbers $a[\{S_{z,i,1}\}]$ and averages them over the course of the measurement sweeps in the simulation. The periodic nature of the trace allows the operator A to be inserted at



FIG. 3: "Typical" snapshots of world–line configurations for the XXZ Hamiltonian on an N = 8 site lattice at J = 1 with $\lambda = 0, 1, 2$, inverse temperature $\beta = 2$ (with L = 8), and magnetization per site $m_z = 0$. The world–lines follow the trajectories of the up spins, as described in the text. As λ becomes large the world–lines increasingly tend to occupy every other spatial site, reflecting a growth of antiferromagnetic order.

any point in the string of matrix elements. To improve statistics, one can measure values $a[\{S_{z,i,l}\}]$ on all time slices.

Expectation values of diagonal operators $\hat{A}(0)\hat{B}(\tau)$ which are offset in imaginary time (here $\hat{B}(\tau) = e^{\tau \hat{H}}\hat{B}e^{-\tau \hat{H}}$) can also be measured for $\tau = n\Delta\tau$ by inserting \hat{A} and \hat{B} into positions in the product of matrix elements separated by imaginary time τ and accumulating the numbers $a[S_{z,1,l}, S_{z,2,l}, \dots S_{z,N,l}] b[S_{z,1,l+n}, S_{z,2,l+n}, \dots S_{z,N,l+n}]$ which arise. Again we can improve statistics by inserting the pair A, B anywhere in the string of incremental time evolution operators as long as they are separated by imaginary time $n\Delta\tau$.

Measuring matrix elements of operators which are not diagonal in the basis of intermediate states is harder, and, indeed, often not possible. Consider $\langle S_{+,i}S_{-,i+1} + S_{-,i}S_{+,i+1} \rangle$, an expectation value which is needed in determining the energy. This operator acts on an intermediate state vector and modifies it, so that the matrix elements in the numerator and denominator are no longer identical. This means that our ratio of integrals is no longer of the form of an integral of some "weight function" in the denominator and the product of the same weight function and a "measurement" in the numerator. We can fix this by multiplying and dividing by the matrix element in the denominator which is missing in the numerator. Then the same product appears in both places, and the measurement is made by accumulating the ratio of the new matrix element to the old one. That is,

$$\langle S_{+,i}S_{-,i+1} + S_{-,i}S_{+,i+1} \rangle =$$

$$\left\langle \frac{\langle S_{z,i,l} S_{z,i+1,l} | (S_{+,i}S_{-,i+1} + S_{-,i}S_{+,i+1}) e^{-\Delta \tau \hat{H}_j} | S_{z,i,l+1} S_{z,i+1,l+1} \rangle}{\langle S_{z,i,l} S_{z,i+1,l} | e^{-\Delta \tau \hat{H}_\alpha} | S_{z,i,l+1} S_{z,i+1,l+1} \rangle} \right\rangle_{\rm MC}.$$
 (14)

where $\langle \rangle_{\text{MC}}$ denotes a monte carlo average of the indicated ratio of matrix elements. We have again exploited the fact that the operator can be inserted at any point l in the string of intermediate states. H_{α} is whatever piece of the Hamiltonian happens to be acting at that point. In fact, to get a nonzero result, H_{α} must be that piece of the Hamiltonian which connects the spatial sites (i, i + 1) in propagating from l to l + 1.

We can make this a little more concrete by noting the matrix element,

$$\langle S_{z,i,l} S_{z,i+1,l} | (S_{+,i} S_{-,i+1} + S_{-,i} S_{+,i+1}) e^{-\tau H_{\alpha}} | S_{z,i,l+1} S_{z,i+1,l+1} \rangle,$$
(15)

vanishes if $S_{z,i,l} + S_{z,i+1,l} = 0, 2$. (Recall local conservation laws require that $S_{z,i,l+1} + S_{z,i+1,l+1} = S_{z,i,l} + S_{z,i+1,l}$. In this case we have either no world-lines propagating up the shaded square, or else two world-lines propagating up the square. No spin-flips are possible, and it is natural that the expectation value vanish. If, on the other hand, the plaquette has $S_{z,i,l} + S_{z,i+1,l} = 1$, then there is a non-zero contribution. The ratio of expectation values is $\tanh(J\Delta\tau/2)$ if the plaquette's single world-line is propagating vertically upward, and $\coth(J\Delta\tau/2)$ if the single world-line traverses the plaquette diagonal. That the kinetic energy is large in the latter case is intuitively appealing when the model being simulates is a fermion or boson Hamiltonian: a world-line hopping diagonally across the plaquette makes a large contribution to the expectation value of the kinetic energy, while a world-line moving upwards (without hopping) makes a small contribution to the kinetic energy.

Measuring an operator which does not conserve particle number locally, is even more complicated, since the resulting ratio is ill-defined. More precisely, an operator like $S_{+,i}S_{-,i+j} + S_{-,i}S_{+,i+j}$ where j > 1 will result in a sequence of states in the numerator which is completely different from that in the denominator. It is possible to solve this problem by conducting two parallel simulations, one with "broken" world-lines, but in practice the technique is complicated, and also results in measurements which are noisy. For this reason, most world-line simulations of the simple sort described here, measure only operators which do not break world-lines. These issues are closely related to the discussion of "forward-" and "side-" walking in diffusion QMC, and are addressed by the development of 'worm' algorithms.

We comment that it is sometimes possible to measure correlation functions which at first appear to break world lines by inserting different types of intermediate states in the simulation. For example, to measure an operator like $S_{+,i}S_{-,i+j} + S_{-,i}S_{+,i+j} = 2(S_{x,i}S_{x,i+j} + S_{y,i}S_{y,i+j})$, which is not diagonal in the basis of $S^{z,i}$, one could insert as intermediate states complete sets of eigenstates of $S_{x,i}$ in the path-integral. Whether or not this trick works depends on the sign problem. Because the XXZ Hamiltonian does not commute with $S_{x,tot}$, the eigenstates $|S_{x,1} = +, S_{x,2} = +\rangle$ and $|S_{x,1} = -, S_{x,2} = -\rangle$ are mixed, in addition to a mixing of $|S_{x,1} = +, S_{x,2} = -\rangle$ and $|S_{x,1} = -, S_{x,2} = +\rangle$. A straightforward calculation shows that there is no sign problem if one is in the Heisenberg or Ising limits $\lambda \geq 1$. In spite of the restrictions on the operators we can measure, we are still able to extract a lot of interesting physics from world–line simulations, as shown briefly below.

We have one final caveat to make about the world-line algorithm. The moves described above, in which world-lines are locally deformed, are not ergodic. For example, they never change the total magnetization. This is so because whole world lines are never created or destroyed, only changed in shape. There is also another, more subtle, quantity that these moves conserve: the winding number. Periodic boundary conditions require that the world lines connect across the last and first time slices in a continuous fashion. They can do this is a trivial way, with each world-line moving generally upward through the lattice (see Fig. 4a), but they can also satisfy the periodic boundary conditions in imaginary time with configurations in which there is a net flow of world-lines across the right or left spatial sides of the system (see Fig. 4b). The latter configurations are said to have non-zero winding number. Moves which are local deformations of the world lines cannot change the winding number, that is, they can never evolve from Fig. 4a to Fig. 4b.



FIG. 4: World-line configurations with winding number W = 0 [a] and W = 1 [b]. Both satisfy all the local conservation laws, as well as periodicity in imaginary time. No sequence of local distortions of the world-lines can evolve the system from [a] to [b].

As with the equivalence of the fixed and fluctuating particle number ensembles, in the thermodynamic limit it can be argued that the fixed winding number sector yields the same physics as allowing the winding number to fluctuate. Perhaps the simplest way to see this is that if there were no spatial periodic boundary conditions in the kinetic energy term, then winding around the lattice would not be possible. But we certainly expect that periodic and open boundary cinditions to give the same result in the thermodynamic limit. Therefore we expect the inclusion of non-zero winding moves to be unnecessary in the thermodynamic limit. It has also been been verified by explicit calculation with and without moves which

change the winding number that expectation values of observables are the same in the thermodynamic limit. An ancillary benefit of loop algorithms, whose primary virtue is the rapid exploration of phase space, is that they also allow for moves connecting the different magnetization and winding number sectors.

D. Some Results

One of the most fundamental issues in the study of the spin-1/2 antiferromagnetic Heisenberg model was whether the ground state in two dimensions exhibits long range 'Neel' order. This was of great interest subsequent to the discovery of high temperature superconductivity because of the antiferromagnetism present in the CuO_2 layers of that material. QMC played an important role in demonstrating that the ground state did have long range order, as opposed to alternatives like 'resonating valence bond' states.

A variant of this problem, which we have studied extensively, is what happens to this Neel state when an external field is applied in the z direction which changes the magnetization away from $m_z = \frac{1}{N} \sum_i S_{z,i} = 0$. Such a field obviously breaks the full rotational symmetry of the Heisenberg model. Somewhat counterintuitively, it favors magnetic order in the xy plane! The reason is that a configurations of spins which are lined up antiferromagnetically in the xy plane can tilt slightly up in the z direction and pick up field energy proportional to the tilt angle. If the spins are lined up antiferromagnetically in the z direction, the energy picked up by rotation of the spins which point in the opposite to the field is quadratic in the rotation angle.

As a fascinating consequence, the application of a field can induce a finite temperature Kosterlitz-Thouless phase transition, even though the system orders only at T = 0 in the absence of the field.

We wanted to understand the answer to the following question: Suppose we took the XXZ model with $\lambda > 1$ so that order in the z direction is preferred. What happens in that case when a field is applied in the z direction? Which effect wins, $\lambda > 1$ or the field B_z ? Our interest in this problem actually came from the fact that the XXZ model is identical to hard-core bosons on a lattice, interacting with a near-neighbor repulsion. The question enunciated above turns out to be completely equivalent to asking which of superfluid or charge density wave order wins in the boson model as the density is varied.

Fig. 5 shows some results for the antiferromagnetic structure factor in the z direction $S_{zz}(\pi,\pi) = \frac{1}{N} \sum_{ij} \langle S_{z,i} S_{z,j} \rangle$, and the spin stiffness ρ_s^{xy} in the xy direction, as a function of magnetization per site m_z . Both are probes of long range magnetic order. The use of the spin stiffness to determine the xy order is a consequence of the limitation on the world-line algorithm's ability to measure the spin correlations $\langle S_{x,i} S_{x,j} \rangle$ directly. Instead the change in the energy in response to a twist of the boundary conditions, ρ_s^{xy} , is measured by examining the average winding number in the simulation. In the absence of a field, when $m_z = 0$, there is only order in the z direction, a consequence of the fact that $\lambda = \frac{3}{2} > 1$. However, when B_z is turned on and m_z moves away from 0, nonzero stiffness develops in the xy plane, and the z order decreases. Finite size scaling of the z structure factor reveals that long range

order remains for m_z nonzero, out to about the point where the linear shape of $S_{zz}(\pi,\pi)$ vs m_z converts to quadratic.

The conclusion appears to be that simultaneous xy and z order occurs when $\lambda > 1$ and $B_z \neq 0$. But more careful analysis of the dependence of the energy on the magnetization reveals that the system is thermodynamically unstable, and rather than exhibiting simultaneous order of the two types, instead phase separates into regions of long range z and long range xy order. This instability appears in Fig. 6 in the form of a negative magnetic susceptibility.



FIG. 5: The magnetic structure factor in the z direction, $S_{zz}(\pi,\pi)$ and spin stiffness in the xy direction ρ_s^{xy} as a function of magnetization per site m_z , in the Ising limit of the XXZ model $(\lambda = \frac{3}{2})$. The lattice is 8x8 and the inverse temperature $\beta = 6$. The system develops long range xy order $(\rho_s^{xy} \neq 0)$ when $m_z \neq 0$. It also appears to retain long range z order.

Interestingly, we found that if one considers the Heisenberg model with a term $J_2 \sum_{\langle \langle ik \rangle \rangle} S_{z,i} S_{z,k}$ which couples the z component of spin on next-near neighbor sites (i.e. sites which are diagonally across from each other on a plaquette of the lattice) then the ground state can exhibit simultaneous order in the xy plane and in the z direction, in which components of spin form alternating up and down S_z lines. In boson language, superfluid order coexists with a pattern of boson occupation in which full and empty rows alternate.

Such issues of phase separation are of course intensely studied for the Hubbard Hamiltonian, and variants like the t - J model as well, where one is interested in the clustering of holes doped into an antiferromagnet.

A final comment concerns the dramatic jump in the field B_z required to change the magnetization around $m_z = 0$ in Fig. 6. If you imagine starting from a state of all spins down, $m_z = -\frac{1}{2}$, then flipping spins up is energetically favorable with an antiferromagnetic interaction. So, in fact, a negative field is required to prevent spins from flipping up. It



FIG. 6: z component of magnetization per site m_z as a function of magnetic field B_z for the XXZ model in the Ising limit $(\lambda = \frac{3}{2})$. The slope of m_z vs B_z is the susceptibility, and is negative near $m_z = 0$. This indicates a thermodynamic instability and phase separation.

becomes slightly less energetically favorable to flip spins up as m_z increases, and so B_z increases gradually with m_z . However, at $m_z = 0$ one has finally a situation where, for the first time, further up flips break antiferromagnetic bonds, an energetically costly thing to do. To force the system into such a state, B_z needs to take a big jump to positive values.

This jump is the analog of the jump in the chemical potential μ needed to change the particle density in the Hubbard model that occurs at the Mott transition, when doubly occupied sites first appear on the lattice.

E. World–Lines for Interacting Fermions

Our second illustration of the world-line algorithm is for lattice fermions. We present only a short discussion because, as we shall see, the sign problem makes fermion world line QMC impractical except in one dimension. Consider first a model of one-dimensional spinless fermions with a repulsion V between near-neighbor sites.

$$H = -t \sum_{i} (c_{i}^{\dagger} c_{i+1} + c_{i+1}^{\dagger} c_{i}) - \mu \sum_{i} n_{i} + V \sum_{i} n_{i} n_{i+1}.$$
 (16)

 c_i^{\dagger} and c_i are fermion creation and destruction operators at site *i* and n_i is the number operator. The formulation of the world-line algorithm resembles that of the quantum spin-1/2 Hamiltonian just described very closely. We break up the exponential of $-\beta H$ and insert complete sets of occupation number states. The only non-zero matrix elements are,

$$\langle 00 | e^{t\Delta\tau(c_i^{\dagger}c_{i+1}+c_{i+1}^{\dagger}c_i),} | 00 \rangle = 1,$$

$$\langle 10 | e^{t\Delta\tau(c_i^{\dagger}c_{i+1}+c_{i+1}^{\dagger}c_i),} | 10 \rangle = \cosh(t\Delta\tau), \qquad \langle 10 | e^{t\Delta\tau(c_i^{\dagger}c_{i+1}+c_{i+1}^{\dagger}c_i),} | 01 \rangle = \sinh(t\Delta\tau),$$

$$\langle 01 | e^{t\Delta\tau(c_i^{\dagger}c_{i+1}+c_{i+1}^{\dagger}c_i),} | 10 \rangle = \sinh(t\Delta\tau), \qquad \langle 01 | e^{t\Delta\tau(c_i^{\dagger}c_{i+1}+c_{i+1}^{\dagger}c_i),} | 01 \rangle = \cosh(t\Delta\tau),$$

$$\langle 11 | e^{t\Delta\tau(c_i^{\dagger}c_{i+1}+c_{i+1}^{\dagger}c_i),} | 11 \rangle = 1.$$

$$(17)$$

These matrix elements are identical to those arising in the spin-1/2 XXZ model. This similarity in the world–line algorithm for spin, boson, and fermion models reflects the exact mappings which exist between these problems in one dimension.

For more general fermion Hamiltonians, world–line algorithms differ from those for quantum spins (and bosons) in two ways. First, in the fermion case, one is typically interested in models in which the operators also carry a spin index. (Of course there are instances of quantum spin and boson models where there are several spin or boson species, but such cases are the exception, whereas for fermions they are the rule.) When there are only density– density interactions of the form $n_{i\uparrow}n_{j\downarrow}$ coupling the two spin species, that is, no spin flip hopping terms like $c_{i\uparrow}^{\dagger}c_{j\downarrow}c_{j\downarrow}$ or more complicated ('Hund's Rules') interactions like $c_{i\uparrow}^{\dagger}c_{i\downarrow}c_{j\downarrow}c_{j\uparrow}$, then the path–integral which arises involves two *separate* checkerboard lattices. That is, the off–diagonal matrix elements involve only one spin species at a time, and the coupling between the spin species is through the diagonal terms which factor out of the matrix elements as numbers. Thus one suggests moves which distort fermion world–lines of one spin species leaving the other spin species' world–lines unchanged, though this static configuration of the other spin species does enter the acceptance–rejection decision through the diagonal interaction terms.

The second, and much more profound, difference is in the sign problem. Boson and quantum spin operators have non-zero commutation relations on the same site, but *commute* on different sites. This means that the sign of the matrix elements entering the simulation is determined completely locally, for example by the explicit solution of the two site problem which arises after the checkerboard decomposition. Fermion operators, on the other hand, *anticommute* on different sites. Additional minus signs can arise in getting the fermion creation operators in their canonical order after a hopping process.

To be more specific, recall that in expressing a fermion occupation number state like $|10110...\rangle$ a convention for the order in which the creation operators act on the vacuum must be chosen. For example we might define $|11010...\rangle = c_1^{\dagger}c_2^{\dagger}c_4^{\dagger}...|vac\rangle$. Consider the action of $c_3^{\dagger}c_4$ on this state. We anticommute the *pair* of operators through c_1^{\dagger} and c_2^{\dagger} , but since we are moving two operators together there are no sign changes. At this point, the destruction operator for site 4 meets the creation operator on site 4 and they cancel. The crucial point to observe is that this whole process leaves the creation operators in their canonical order $c_1^{\dagger}c_2^{\dagger}c_3^{\dagger}...|vac\rangle$ so we get the occupation number state $|11100...\rangle$ with no minus sign.

What would happen if we instead acted with a longer range hopping $c_5^{\dagger}c_2$ on this state? This pair of operators would anticommute through c_1^{\dagger} without introducing a minus sign. c_2^{\dagger} would then meet and destroy c_2 , but c_5^{\dagger} would be left out of its canonical order, that is, to the left of c_4^{\dagger} . Anticommuting it through yields $-|10011...\rangle$. This minus sign makes the probability with which one would like to generate the associated world-line configurations *negative*, which prevents useful simulations. The general rule is that one gets a minus sign if there is an odd number of occupied sites "in between" the two states that are connected by the hopping, where "in between" means the order of creation operators in the convention chosen for acting on the vacuum state. One does not need to have long-range hopping to encounter this problem. If one has a two-dimensional lattice and has a convention for occupation number states in which the operators for rows in the x direction are adjacent, the hopping in the y direction will be between sites which are not "neighbors" in the string of creation operators. Such hopping processes will yield minus signs.

Fermion world-line simulations are not possible except strictly in one dimension. Even a one-dimensional lattice with a single "impurity" orbital onto which the fermions can hop will have a sign problem. As we have discussed, hoppings which are longer range than near-neighbor cannot be handled. An exception is the case when only interactions, and not kinetic energy terms couple one-dimensional chains along which the fermions hop. Thus the world-line algorithm for fermions is almost exclusively applied to models in one-dimension.

F. Conclusions

The world–line QMC algorithm is a powerful approach to the simulation of lattice quantum spin, boson, and fermion models. It is considerably more pictorial than other QMC methods like the determinant algorithm. It also has a local action which results in a nominally linear scaling with system size for a sweep through the lattice updating all the degrees of freedom.

However, the technique also has a number of drawbacks: The difficulty in measuring certain observables, the sign problem which prevents the study of frustated spin models or fermions in greater than one dimension, and long equilibration and autocorrelation times. A number of these difficulties have been solved by the construction of loop, continuous time, worm, and stochastic series expansion, algorithms. But the sign problem has not been solved.

G. Acknowledgements

Preparation of these lectures was supported by the National Science Foundation under grants NSF–DMR–9985978 and NSF–ITR–0313390.

H. Appendix: A Lightning Review of Monte Carlo

Nature generates sequences of states for the experimentalist according to the Boltzmann distribution. The experimentalist then looks at those states and makes measurements, averaging over as long a time (as many states) as needed to get results to a desired level of

accuracy. The goal of Monte Carlo is to devise a method where the computer plays a similar role to that of nature, generating configurations according to some desired probability. In the following, we will consider a general probability for the state *i* given by $p_i = w_i / \sum_i w_i$. In most classical statistical mechanics problems $w_i = e^{-\beta E_i}$, the Boltzmann distribution, but we have seen in QMC that w_i could have some more complex structure.

The Monte Carlo procedure is very simple.

(i.) Take an initial configuration i and propose a change to a new configuration j.

(ii) Accept the new configuration with probability $\min(1, w_j/w_i)$. Notice that you do not need to know $\sum_i w_i$ to compute this probability.

(iii) Repeat the process many times.

There are of course many subtleties: What sort of procedure is followed to "suggest changes"? How many configurations do you need to generate and do you, as in some experiments, need to allow the system time to equilibrate? The reader is referred to textbooks for the answers to these detailed questions.

The rough justification of the monte carlo method is the following. Denote by T_{ji} the probability of making a transition from state *i* to state *j*.

(i) The matrix T has the properties $T_{ji} \ge 0$ and $\sum_j T_{ji} = 1$. It is not difficult to show that such a *stochastic* matrix has maximum eigenvalue $\lambda = 1$.

(ii) The eigenvector with this maximal eigenvalue is w_i . This can be seen by noting that the choice of acceptance probability p guarantees detailed balance: $T_{ji}w_i = T_{ij}w_j$ for all pairs of states ij. As a consequence $\sum_i T_{ji}w_i = \sum_i T_{ij}w_j = w_j \sum_i T_{ij} = 1 \cdot w_j$.

(iii) The repeated use of T to generate states can be thought of as the repeated application of the matrix T, a process which, for any matrix, projects out the eigenvector of largest eigenvalue. Since w is that eigenvector, repeated use of T will generate states with probability proportional to w_i .