THE APPLICATION OF QUANTUM MONTE CARLO TO PROBLEMS IN ELECTRONIC STRUCTURE

D. M. CEPERLEY
National Center for Supercomputing Applications,
Department of Physics and the Beckman Institute,
University of Illinois at Urbana-Champaign
1110 W. Green St., Urbana, Il, 61801, USA

ABSTRACT. A overall view of the quantum Monte Carlo methods which have been applied to the calculation of electronic properties of condensed matter systems is given. The dependance of statistical and systematic errors on trial functions and computer time is discussed. Generic problems and opportunities are identified.

1. Introduction

The long range goal of research into quantum Monte Carlo (QMC) methods is to be able to accurately predict physical properties of condensed matter from the fundamental equations for the electrons and nuclei. These methods can, in principle, calculate exact properties of many-body systems, but almost all applications have been thus far limited to the ground states of small molecules and solids, such as hydrogen, which do not have core electrons. The methodology has been actively discussed for 25 years; the first large-scale application was a variational calculation of McMillan[1] in 1965, followed by an exact GFMC calculation of Kalos[2] both of the microscopic properties of superfluid ^{4}He. However, except for a few exceptions, the QMC methods have not yet made a major impact in material science. A job that QMC that already does well is to provide very accurate benchmark calculations for more approximate methods such as those of local density functional theory(LDA). A much more ambitious task for QMC, by analogy with classical simulations, is to act as a black box, where the Hamiltonian can be inserted into the top of the machine and the exact quantum many-body properties come out the bottom, without significant human intervention. The QMC methods are not sufficiently developed for this approach to work. This article will review the methods, the progress, the outstanding opportunities and attempt to make some prognostications on the future of this method.

2. Quantum Monte Carlo Methods

In this section I will briefly list the stochastic methods which have been used to solve electronic structure problems. I will only discuss methods to directly solve the Schroedinger equation and not ones for the simplified Heisenberg or Hubbard models. The list will only contain "exact" stochastic methods that have been applied to solving the non-relativistic Schroedinger equation for a finite many-body systems of electrons and nucleons.

2.1. VARIATIONAL MONTE CARLO

Variational Monte Carlo (VMC) is the simplest and most direct method simulation method for quantum systems. Let $\Psi(R)$ be a trial wave function. The form almost universally used is of the Bijl-Jastrow-Slater (BJS)[3] form:

$$\Psi(R)= \exp(- \Sigma \, u(r_{ij})) \, \| \phi_k (r_i) \| \qquad\qquad (1)$$

where $\|\phi_k(r_i)\|$ is a model one-body determinant which can be obtained in Hartree-Fock calculation and u(r) is a pair wave function between electrons i and j. A VMC calculation consists of sampling $|\Psi(R)|^2$ and evaluating expectations such as the energy as an average over the sampled configurations[4]. For a bose system like 4He, with no fermion determinant, the VMC calculation is essentially identical to a classical molecular dynamics or Monte Carlo simulation. For fermion systems, the only new feature is the evaluation of the determinant, which takes order N^3 operations. Some identities from linear algebra can be used to speed up that evaluation for QMC[5].
   The main advantage of the VMC method over traditional HF or LDA methods is that any computable function can used for a trial function and there is still some hope that very accurate, yet simple trial functions can be found. The pair trial function of Eq. (1) can directly account for most electron-electron correlation, something which would need an expansion into an exponentially large number of Slater determinants of one-body functions to achieve. As we shall see in the next section, accuracies of the total energy of 0.3 ev/atom are typical of variational energies for this trial function.
   The disadvantage with the VMC method is that it is hard to optimize the trial function since the energy is a non-linear function of the variational parameters and all quantities are obtained with statistical errors. Correlated sampling methods[4] give smooth functions of the energy with respect to the free parameters in the trial function so that one can take derivatives and thus go in the direction of the optimal trial function. But there is still noise so that the there always a large error in the estimate of the optimal trial function parameters. This is not much of a problem in estimating energies and energy differences but it is much more of a problem in other properties. There is some empirical evidence that it is more robust to minimize the variance of the local energy (i.e. $|H\Psi(R)|^2$) to obtain variational parameters[6]. There have been very few attempts to obtain the optimal trial function within a class of wave functions. Only recently[7] have we learned that the optimal pair product trial function for the Ne atom will give 85% of the correlation energy (it misses by 1.6 ev). It is now clear that in order to achieve accuracies better than 0.3 ev/atom in total energy of valence electrons will require going beyond the two-body level in the trial function. Whether that is computationally feasible remains to be seen.
   One might hope that in the calculation of relative energy differences, high accuracy could still be obtained by these relatively crude pair trial functions by using the same practical philosophy prevailing elsewhere: missing correlation energy at the variational level will tend to be independent of the placement of the ions, ionization state or excitation level state etc., and cancellation of errors will occur. Nonetheless, it will still be hard to achieve reliable predictions at the level of 0.1ev with VMC methods and pair trial functions. Human nature is such that one tends to stop optimizing wave functions, adding basis functions etc., when expected results have been obtained. Also the VMC method systematically favors simpler states. An example is the phase transition of electrons at low

density. Variationally it is much simpler to describe a Wigner crystal than a highly correlated electron fluid. Thus variational results for the Wigner transition[8-9] are a factor of 3 too high in density. For these reasons it is wise to search for more automatic methods that are less subject to bias.

2.2. PROJECTOR MONTE CARLO

A more robust QMC method is based on using the operator exp(-Ht) to project out from a given trial function all high energy components, leaving the ground state at large imaginary 'time', t..This projection can be accomplished with a stochastic algorithm because spatial matrix elements <R|exp(-tH)|R'> are always positive, except those introduced by fermi statistics. It was realized by Kalos[2] that any knowledge of the ground state wave function can be used to make the results statistically much more reliable. This is called importance sampling with a guiding function, $\Psi(R)$. The projection operator with importance sampling is

$$f(R,R';t) = \Psi(R)<R|\exp(-tH)|R'>/\Psi(R').$$

Usually the best available variational trial function is used both to guide the random walk and as the initial condition on the projection. There are several varieties of algorithms that have been used to accomplish this projection, the two most common being Diffusion Monte Carlo[10] (DMC) and Green's Function Monte Carlo[4] (GFMC).

To motivate DMC let us write the evolution equation satisfied by the above operator in the form:

$$df(R;t)/dt = \lambda \ \mathbf{grad} \ [\ \mathbf{grad} \ f - D(R) \ f \ ] - [E_L(R) - E_T \ ] \ f(R;t) \qquad (2)$$

where $E_L(R) = \Psi^{-1}H\Psi$ is the local energy, $D(R) = 2\Psi^{-1}\mathbf{grad}\Psi$ is the "quantum force" of the guiding function, $\lambda = h^2/2m$ and E_T is the "trial energy" adjusted to keep the total normalization of f asymptotically constant. Then interpreting f(R;t) as a probability density, its evolution is the result of three processes corresponding to the terms on the right hand side: diffusion, drift and branching. In DMC one usually deals with an ensemble of random walks, typically 100 to 1000 separate systems. The elementary moves consist of adding small random displacements to the electron positions, of drifting the positions in the direction of the quantum force and of duplicating or deleting members of the ensemble. In practice, the evolution is usually accomplished by assuming the branching and drift are constant within a single move. This leads to a systematic time-step error but there are variants of the method correct to various orders of time step. We use a generalized Metropolis algorithm[10], where trial moves are occasionally rejected. Then time step errors are proportional to the difference between the exact wave function and the trial functions and the algorithm is very robust.

Such time step errors can be completely eliminated by the GFMC algorithm with the addition of an extra process in the random walk, the process of promoting a given random walk from one "generation" to the next. Special algorithms[11] have been developed for Coulomb potential. These are only called for if high accuracy is demanded.

The energy is obtained very accurately in DMC because with a good importance function, the local energy is nearly constant. However other properties, not commuting with the Hamiltonian have a possible systematic bias because averages are obtained from

the mixed distribution $f=\phi\Psi$ where ϕ is the exact wave function instead of from the square of the ground state wave function. In principle the bias can be removed with either the forward walking method or with the simpler, but approximate, linear extrapolation[4] but in practice it is difficult to ensure that convergence is reached.

Fermi statistics pose a very difficult problem to all projection methods[12] for electronic systems in 2 or 3 dimensions. Fermi statistics are not unique; any many-body excited state will have essentially the same problem. To project out an excited state, one begins with the initial state orthogonal to all lower-lying states. But this will be non-positive and cannot be simply interpreted as a probability distribution. Alternatively one can antisymmetrize the projecting function but then the matrix elements needed in the random walk become non-positive. One must assign plus and minus signs to random walks and any expectation value will then be expressed as the difference between its value on the plus and minus walks. This is called the transient estimate method. It is an exact method, however, it is statistically unstable in the sense that for large projection times, the signal-to-noise ratio decreases exponentially fast. The exponent is proportional to the total excitation energy, so that it becomes worse, either for higher excitation energies, or in the case of fermi statistics, for more electrons. Accurate results can be obtained, given good guiding functions, for, on the order of 50 fermions or less, but larger fermion systems are prohibitively slow.

The stable, yet approximate, fixed-node method[9,13] is more commonly used. In this method, one places the extra boundary condition on the unknown 'exact' solution that it vanish whenever the trial function does. Then the Schroedinger equation can be solved in the nodal region delimited by the nodes of the trial function and antisymmetry no longer has to be directly taken into account. The boundary condition raises the energy, one obtains the lowest energy consistent with the fixed-nodes. Because the energy is second order in the placement of the nodes, and because all the 'boson-like' correlation is exactly put in, this method turns out to be about one order of magnitude more accurate than VMC; if HF-like nodes are used. This boundary condition is implemented in very simple fashion: moves which cross the nodes of the trial function are not allowed. The release-node method[11] is a combination of the fixed-node and transient estimate method. It is a simple way of directly calculating the difference between the fixed-node energy and the exact energy.

There are other fermion methods used with lattice models. The so-called Stratonvitch-Hubbard transformation[14] is used to replace the electron-electron interaction with a auxiliary (random) field which must be sampled over. There have been some attempts[15] to apply this to continuum problems, but so far with little success. The difficulty is in finding a method which does not have minus sign problems, and for which a realistic electron-electron correlation does not cause wild fluctuations in the auxiliary fields. Also one does not know how to apply importance sampling beyond the one-body level. On the other hand, such methods lend themselves to problems where one wants to freeze all but certain energy bands. An extreme case is the Hubbard model of a solid

2.3. PATH INTEGRAL MONTE CARLO

Path integral Monte Carlo[16], (aka world-line Monte Carlo) is an exact algorithm is calculate thermal averages, not single state properties and may in the future be applied to electronic structure problems. One expands the trace of the density matrix into a path in imaginary time. A single classical particle turns into a ring "polymer." The number of polymer beads is proportional to the inverse temperature. Quantum statistics comes into the path integral as an exchange of particles as they go around the path. The problem of

superfluidity of bose liquids becomes mapped into a classical system of ring polymer which can hook up according to any possible permutation of particle labels. A generalized Metropolis algorithm is used to sample the paths and permutations. Superfluity and bose condensation are reflected by paths which stretch across a macroscopic distance in the fluid.

The chief advantage of the path integral formulation is that it allows calculations of thermal properties. One can simulate finite temperature, quantum, phase transitions. It is not as efficient in calculating the ground state energy as the projector method. However, it is more robust in calculating other properties than the energy, since the bias of the guiding function does not have to be removed. Fermi statistics poses the same problem for the path integral method as it does for other exact quantum methods; antisymmetry is taken into account by assigning a negative sign to walks which involve an odd permutation of particle labels. For large systems *or* as the temperature goes to zero, the parity of the permutation becomes random. This implies the signal-to-noise ratio will go to zero exponentially fast in both the particle number and in the inverse temperature. Currently we are in the process of developing a fixed-node path integral Monte Carlo algorithm for fermi systems[17] which will be applicable to electronic structure calculations at high temperature (say 1000 K). The temperature is high because one has to work down from temperatures on the order of a Rydberg.

3. Pseudopotentials in QMC

If all electrons are included in a QMC simulation, the time step will be controlled by that of the 1s electrons and the statistical error will be dominated by the fluctuations of the core electrons[18]. Of course, the interesting physics is usually in the valence region and on the energy scale of the valence electrons. It is wise to eliminate as many electrons as is practical. Pseudo-potentials(pp) are widely used in most *ab initio* methods. However, there is good reason to reexamine their form for use in QMC. The QMC method can compute exact properties of a Hamiltonian and, if we are to spend large amounts of computer time doing benchmark calculations, it would be highly desirable to have a good approximation to the correct Hamiltonian for the purpose of comparison with experimental data and ultimately to make reliable predictions.

Here I will discuss only two of the possible ways of eliminating core electrons. Let me first recall the usual form of the pp used in LDA calculations. The electron-core interaction is typically assumed to be the non-local operator:

$$h(r) = \sum_{l,m} v_l(r) \, Y_{lm}(\Omega) \int d\Omega' Y_{lm}^*(\Omega') \qquad (3)$$

where $v_l(r)$ is the core-electron interaction in the lth angular momentum channel. There is no difficulty in applying this potential in a VMC method[19]. One simply has an extra integral to perform periodically during the random walk, the integral of an electron over a sphere around the nearest ion. But non-locality is a major inconvenience in PMC because if the non-local part of the potential becomes positive, as it often does, signs, which cause the "fermion disease", are introduced even for a single electron. Two different methods have been made around devised to allow DMC simulations with realistic pps.

Bachelet et. al.[20] have constructed a local pseudo-hamiltonian where the effective mass tensor of the electrons varies and is non-isotropic in the core:

$$h(r) = -\mathbf{grad}\, A(r)\, \mathbf{grad} + v(r).$$

Here $A(r)$ is an effective mass tensor and $v(r)$ is a local potential. Then standard fixed-node and release-node methods can be used except the diffusion process inside an ion core differs in the radial and tangential directions. Within the constraint that the effective mass must always remain positive, one can fix the interaction of an ion with an electron in three separate angular momentum channels. The QMC results are in reasonable agreement with experiment for the third row atoms (Na through Cl) dimers and even solids[24], giving much better binding energies that those from LDA or HF calculations. To date, a satisfactory derivation of the ph from first principles has not yet been accomplished so one does not know how well such a Hamiltonian will be able to describe these atoms. But there is a very serious problem in extending the pseudo-hamiltonian transition metal elements because by construction the first electron of a ph must lead to an s-like ground state. Thus the pseudoization for an atom like Cu must extend to a Ne like core which causes the Monte Carlo calculations to be too costly of computer time.

One can directly use DMC methods with a non-local pp by evaluating the non-locality on a good trial function[21]. This introduces a new approximation but the effect of the approximation will vanish as the trial function approaches the exact wave function. The simplest variant of this approach is to evaluate the non-locality on a pure determinantal trial function. We[22] have used an optimized Bijl-Jastrow-Slater trial function to minimize the "locality approximation". The integral of the trial function over the sphere is done using standard quadrature formulas. For example, with a Si atom, one can use a 6-point formula, with the points along the coordinate axes, and for the Cu atom an 18-point formula. The use of the locality approximation is intimately tied to the use of the fixed-node approximation since otherwise fluctuations in the energy around the nodes of the trial function will be severe and unphysical. The upper bound property of the fixed-node approximation no longer holds when the original potential is non-local. But this seems to be the simplest approach for DMC calculations for something like solid Cu if one only wants to include the 3d and 4s electrons in the simulation.

Currently we are using QMC in two different ways in the theory of pseudopotentials. First, to directly calculate an effective many-body pseudopotential. To date we have used pp derived from LDA or HF theory, and these are not adequate because the precision of QMC calculations is much higher. The error caused by the use of an approximate pp is the largest systematic effect. One can formulate a many-body pp in terms of "natural orbitals", single body Green's functions and single body density matrices. Such pseudopotentials are optimal in the linear response regime. We are using QMC on full core atoms to calculate the natural orbitals and hope thereby to derive pseudoatoms with identical natural orbitals outside the atomic cores.

Second, we[23] are using QMC to test the pp. We have determined that LDA derived pp overbind molecules and HF derived pp underbind molecules. The optimal PP should be somewhere in between. Major problems have been discovered in using LDA derived pp for the Cu atom[22]. Of course it is not yet clear whether any possible pp of the form of Eq. (3) will be able to calculate relative energies accurate to 0.01 ev. Core polarization terms are necessary to obtain a good description in some cases and they are straightforward to put into the DMC calculation.

4. Errors in Quantum Monte Carlo

In this section I will discuss what has been learned in recent years about the magnitude of errors that are present in Quantum Monte Carlo calculations and how much computer time

is needed to achieve an interesting level of precision. There are many kinds of errors that can be controlled through careful programing. Some examples are the inadequacies of the random number generator and the floating point precision of the computer. The large errors due to a pseudopotential has been discussed in the previous section.

Finite size effects are often the largest errors when one wants to compare simulations done with typically several hundred fermions to a physical system in the thermodynamic limit. Some of the finite scaling methods developed for classical statistical mechanics are applicable to quantum system, but fermion system have often a more complex scaling behavior. As an illustration, consider how a fermi surface with necks and gaps will appear in a sequence of finite systems in periodic boundary conditions. Also, because a simulation of a fermion system is slower than that of a classical system the range of system sizes to test the scaling behavior is much reduced. The largest simulation that we have attempted has used on the order of 500 fermions. To perform simulations on larger systems of delocalized electrons will require a major jump in algorithms and computers. Thus there is little possibility at the present of studying subtle fermion phase transitions and careful control of finite size effects is in the future. Happily the finite size effects in many systems can be studied at the mean field level and used to correct the simulation results.

Let us now consider the effects of the two other uncontrollable errors, the statistical error and the systematic error by which I mean the limited function space in VMC or the fixed-node error in DMC.

It is shown in ref. [18] and confirmed by a variety of empirical studies, that the statistical error is controlled by the "missing" energy in the trial function:

$$\sigma \geq \sqrt{\frac{2(E_V - E_0)}{\tau P}} \qquad (4)$$

where E_V is the energy of the trial function, E_0 the ground state energy, τ the effective time step of the VMC or FN-DMC run and P is the total number of steps. A trial function is optimal for use in QMC if the missing energy times the computer time per step is minimized.

What is plotted as squares in Figure (1) is the amount of computer time needed for the statistical error to equal this "missing" energy for a variety of systems. The diagonal line shows on this log-log plot the usual $T^{-1/2}$ decrease with computer time. If the trial function is made more accurate the point will move both down and to the right according to Eq. (4); i. e. the statistical error will converge faster but it is usually worth running a more accurate trial function longer since the systematic error is reduced also. One sees that the variational results are typically accurate to 0.1 to 0.5 ev/atom. The trial function for Si_2 is particularly bad. The increase in computer time as the system gets larger is mainly due to the increase in computational effort which roughly scales as the cube of the number of atoms. Generally one will want to get statistical errors an order of magnitude smaller than the missing energy because in calculating differences in energy much of the missing correlation energy will cancel. As an example, in the calculation of the binding energy of solid Si if we subtract the solid variational energy per atom from the atomic variational energy one makes an error of 0.13 ev/atom while the solid energy is in error by 0.33 ev/atom. In the case of Si it is worth doing variational calculations to a statistical accuracy of at least 0.1 ev/atom.

The pluses in Figure (1) show the computer time needed to have the statistical error equal the error of the fixed-node approximation. It is difficult to accurately calculate this fixed-node error for many-body systems but estimates range between 0.05 and 01 ev/atom if Hartree-Fock or LDA trial functions are used. Again one might expect that differences in
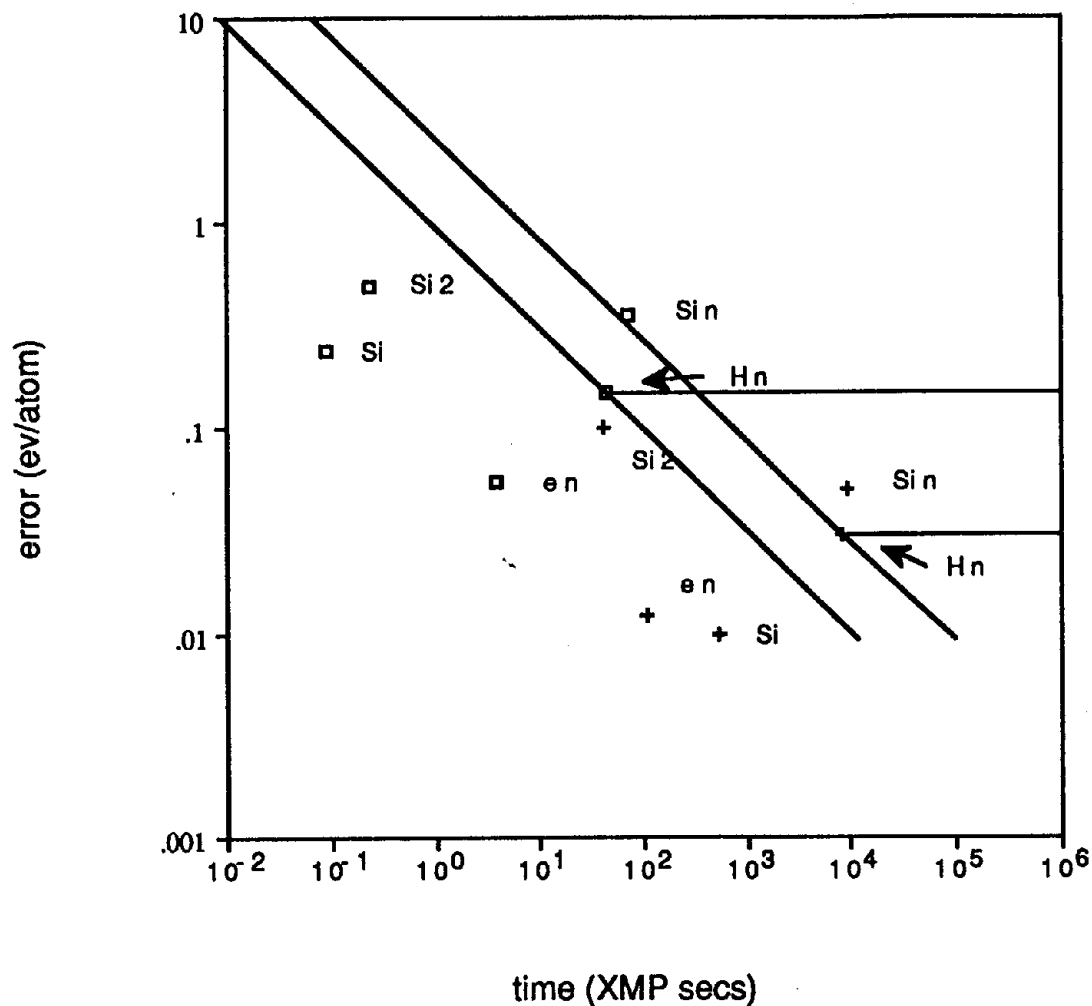
Figure 1: The systematic error in ev/atom for 5 different systems as a function of the CRAY-XMP time in seconds needed so that the statistical error equals the systematic error. The squares are variational Monte Carlo results with a two-body trial function and the pluses are fixed-node diffusion Monte Carlo results. The systems plotted are a Si atom with the pseudo-hamiltonian[20], (4 electron system) the Si_2 dimer, and solid silicon, Si_n[24] where n=64, solid hydrogen, H_n[25] where n=125 at rs=1.3 and the electron gas[9] at rs=2. The fixed-node error for bulk Si and H is roughly estimated, for the other systems it is computed with the release-node procedure. The solid lines show the evolution of the error for the VMC and FN-DMC solid H calculations. The error is dominated by the statistical error for times less than the statistical error, for larger times it is dominated by the systematic error.

energy will have a 3 times smaller systematic error than the total energy so that one can expect to calculate binding energies accurate to order 0.01 ev. Thus QMC is already capable of performing many interesting condensed matter calculations with current algorithms and computers. Such a calculation for solid H will take 27 XMP hours. Surprisingly, the computer time needed to get the errors down to 0.01 ev/atom for 64 atoms of Si is only 30 times needed needed for the Si_2 molecule. The FN-DMC result converges at a slower rate than the VMC result simply because the FN-DMC algorithm must use a smaller time-step, τ, to accurately compute the branching factor and drift terms.

5. Conclusions and Prospects

The only QMC work to have large impact on the electronic structure community was the calculation of the ground state energy of the electron gas[9]. In this section I would like to examine some of the reasons why these methods have not made a broader impact. First of all, electronic systems are relatively well described by mean field theories (MF) (by this I mean both LDA and HF methods) while such is not the case for helium physics; there the MF level approximation is not even useful. Applications of QMC have in fact, usually been in areas where mean field theory does relatively well. An example is in the ground state energy of small molecules. Variational calculations are essentially trying to get rid of excited state components in a MF trial function and this is possible for a molecule since there is a reasonably large energy gap to the next relevant excited state. It should also be mentioned there has been many-orders of magnitude more effort devoted to the MF and perturbation theory methods, software and training so that there is effectively an enormous entry price for a new methods to become competitive.

The biggest successes should be for systems where mean field theory does not work so well, because of cooperative phase transitions, degenerate ground states or other many-body effects or when one cannot make the reduction to the Born-Oppenheimer potential surface . The electron gas provides one example. The QMC results are widely taken as the standard for the LDA exchange and correlation energy. The QMC calculation are so accurate because the trial function for a homogeneous systems is much more accurate than it would be an an inhomogeneous case. With importance sampling QMC is more-or-less a direct method for calculating differences with respect to HF. At low densities, the predictions of the Wigner crystallization density[8-9] stand as one of the first, and, as yet untested, predictions of QMC. There the failure of MFT is plain since it is difficult to study on the same level of approximation a strongly correlated electronic liquid and a localized quantum crystal. But QMC also has serious difficulties, because of the bias of the importance sampling and the FN approximation. It is very likely that the true ground state of the electronic liquid has considerably richer structure than the simple polarized and unpolarized liquid studied in those earlier QMC calculations, as is suggested by the rich phase diagram of liquid and solid ^3He.

Many-body hydrogen illustrates another successful QMC study[25], but one that has as yet not realized its full potential. In this first QMC calculation of a real solid, very simple atomic and molecular trial functions were used, thus treating, on the same footing, the molecular and atomic phases. While MFT calculations are expected to be accurate in the atomic phase, they do a very poor job of describing the metal-insulator transition. An advantage of QMC is that the large zero point motion of the protons can be directly put in, without any approximation or substantial change in the method. Thus we were able to study with the same method, the atomic-molecular transition, the metallization, and the orientational ordering of the molecules. There were some aspects which were clearly less

than adequate in that first calculation. First, the use of very simple rather than the full MFT trial functions resulted in a considerable loss of accuracy at both the VMC and FN-DMC level. (The missing energy was on the order of 0.6ev/atom rather than 0.1ev/atom with the best MFT trial functions.) Second, computer resources in 1980 were really not adequate to handle the 2-time scale problem introduced by allowing the protons to move. We are now continuing these calculations with optimized trial functions.

What are some of the methodological problems which must be solved before QMC has a larger impact in condensed matter physics? The only really fundamental problem is, I believe, the fermion problem. But as the figure shows, the FN calculations could easily be adequate to obtain accuracies of the order of 0.01 ev/atom, even without solving this problem. But because of its variational character, this could still be not adequate in describing phase transitions, etc. In the past there has not been enough effort to optimize trial functions so that the real potential of VMC and DMC is not yet evident. It is just not that difficult to find the optimal BJS trial function, but what one needs is a standard code that can do this reliably and routinely. An extremely interesting problem is to use QMC techniques to learn the mathematical structure of the many-fermion wave function. For example, it is known that incorporating some special backflow and three-body terms can drastically improve the description of the ground state of liquid ^{3}He but nothing similar has been tried with electronic systems. In a similar vein, no one has yet devised an algorithm to determine that determinant which minimizes the fixed-node energy. One expects that the stronger implementation of the variational principle and resulting self-consistency will have the same beneficial effect that it does in all other electronic structure methods. One should not give up on more rigorous fermion algorithms. The release-node algorithm is after all practically convergent for some systems of 50 fermions and we are testing out some new ideas to extend the convergence to larger systems.

One important problem in QMC-theory is to directly calculate ionic forces with sufficient accuracy that QMC could be used for a classical MD or MC simulation of atoms and molecules. An important goal is an exact version of the Car-Parinello algorithm to benchmark the LDA approximation in a disordered environment. An example of a problem that can probably only be solved with QMC is to calculate a rate reaction that is strongly modified by a disordered environment like a liquid. In general, one wants to learn how to do various calculations of energy differences, for example, the energy to create a particle-hole excitation in a solid. The naive method would be to perform one simulation in the ground state and another in the excited state and subtract to find the excitation energy. For such a scheme to work the statistical error of the energy of the whole system must be reduced to the level of 0.1 ev. This then is N^2 times as costly, where N is the number of atoms needed to give a reasonable description of the bulk. Correlated sampling methods reduce the time needed by a large amount but a full calculation at the DMC level of energy differences in a many-body system has not yet been attempted.

Also, we must remind ourselves that experimental results are usually obtained by measuring dynamical properties and thus calculations of the ground state energy are not all that interesting. It is clearly important to focus on reliable calculation of other properties than the energy and to really understand how the choice of the trial function affects the calculated long-range order. True quantum dynamics is a very difficult problem. Progress has been made in the more limited goal of equilibrium dynamics but I would like point out these approaches will not work if the fixed-node approximation is made.

Some related problems where QMC could have major impact is in the description of the metal-insulator transition and on systems where electronic correlation is very large such as in hydrogen or on surfaces and defects. Putting muons, excitons or positrons into a solid is clearly an area where QMC will have unique capabilities since with those heavier

particles one is in an intermediate regime where neither MFT or the adiabatic approach is satisfactory.

QMC will continue to be important for benchmark calculations, particularly in low symmetry situations where explicit methods may have basis set problems. The possibility of finite temperature fermion QMC opens up new applications in the area of astrophysics and plasma physics and possibly for fermion superfluids. Finally we would like to point out that new massively parallel computers could deliver 3 to 4 orders-of-magnitude more performance than present computers and that QMC calculations can easily be adapted to these machines regardless of whether the architecture it is a loosely coupled network of RISC workstations or a SIMD machine.

6. Acknowledgments

7. References

1. Kalos. M. H., Levesque, D. and Verlet, L., *Phys. Rev.* A9,2178 (1974).
2. McMillan, W. L., *Phys. Rev.* A **138**, 442 (1965).
3. Bijl, A. *Physcia* 7,869 (1940); R. Jastrow, *Phys. Rev.* **98**, 1479 (1955).
4. Ceperley, D.M., and Kalos, M.H., "Quantum Many-Body Problems," in <u>Monte Carlo Methods in Statistical Physics</u>, ed. K. Binder, Springer-Verlag, 1979.
5. Ceperley, D., Chester, G.V., and Kalos, M.H., *Phys. Rev. B*, **16**, 3081 (1977).
6. Umrigar, C. J., Wilson, K. G., and Wilkins, J. W., *Phys. Rev. Letts.* 60,1719 (1988).
7. Schmidt, K. E. and Moskowitz, J. W., *J. Chem. Phys.* **93**, 4172 (1990).
8. Ceperley, D., *Phys. Rev. B*, **18**, 3126 (1978).
9. Ceperley, D.M., and Alder, B.J., *Phys. Rev. Letts.*, **45**, 566 (1980).
10. Reynolds, P.J., Ceperley, D.M., Alder, B.J., and Lester Jr., W.A., *J. Chem. Phys.*, **77**, 5593 (1982)
11. Ceperley, D., *J. Comp. Phys.*, **51**, 404 (1983). Ceperley, D.M., and Alder, B.J., *J. Chem. Phys.*, **81**, 5833 (1984).
12. Schmidt, K. E. and Kalos, M. H. in <u>Applications of the Monte Carlo Method in Statistical Physics</u>, Ed. K. Binder (Springer-Verlag, 1984).
13. Anderson, J. B. , *J. Chem. Phys.* 65, 4121 (1976).
14. Stratonovitch, R. L., *Dokl. Akad. Nauk. SSSR* 115,1097(1957)., Hubbard, J., *Phys. Rev. Letts.* 3, 77 (1959).
15. White, S. R., Wilkins, J. W., and Wilson, K.G., *Phys. Rev. Letts.* **56**, 5024 (1988).
16. Ceperley, D.M., and Pollock, E.L., *Phys. Rev. Letts.*, **56**, 351 (1986).Ceperley, D. M., and Pollock, E. L., "Path-Integral Computation Techniques for Superfluid ^{4}He", <u>Proceedings of Elba Conference on Monte Carlo Methods for Physics,</u>1990. Schmidt, K. E., and Ceperley, D. M., "Monte Carlo Techniques for Quantum Fluids, Solids and Droplets", in <u>Monte Carlo Methods</u> Vol 3, ed. K. Binder, Springer-Verlag, 1991.
17. Ceperley, D. M., "Fermion Nodes", in press *J. Stat. Phys.* (1991).

18. Ceperley, D., *J. of Stat. Phys.*, **43**, 815 (1986).

19. Fahy, S., Wang, X. W., and Louie, S. G., *Phys. Rev. Letts.* **61**,1631 (1988); Phys. Rev. **B42**, 3503 (1990).

20. Bachelet, G.B., Ceperley, D.M., and Chiochetti, M., *Phys. Rev. Letts.*, **62**, 2088 (1989).

21. Hurley, M. M. and Christiansen, P. A., *J. Chem. Phys.* **86**, 1069 (1987), Hammond, B. L., Reynolds, P. J. and Lester, W. A. *J. Chem. Phys.* **87**,1130 (1987).

22. Mitas, L., Shirley E. L., and Ceperley, D. M., "Nonlocal Pseudopotentials and Diffusion Monte Carlo" submitted to *J. Chem. Phys.* (1991)

23. Shirley, E., Martin, R. M., Bachelet G. B., and Ceperley, D. M., *Phys. Rev.* **B 42**, 5057(1990).

24. Li, X.-P., Ceperley, D. M., and Martin,R. M., "Calculation of the Equation of State of Solid Si by Green's Function Monte Carlo", submitted *Phys. Rev. Letts.* 1990.

25. Ceperley, D.M., and Alder, B.J., *Phys. Rev. B.*, **36**, 2092 (1987).