Transition metal oxides using quantum Monte Carlo

2007 J. Phys.: Condens. Matter 19 343201

(http://iopscience.iop.org/0953-8984/19/34/343201)

IOP PUBLISHING JOURNAL OF PHYSICS: CONDENSED MATTER

J. Phys.: Condens. Matter **19** (2007) 343201 (13pp) doi:10.1088/0953-8984/19/34/343201

TOPICAL REVIEW

Transition metal oxides using quantum Monte Carlo

Lucas K Wagner[1]

Center for High Performance Simulation and Department of Physics, North Carolina State University, Raleigh, NC 27695, USA

E-mail: lkwagner@berkeley.edu

Received 9 May 2007, in final form 22 June 2007
Published 20 July 2007
Online at stacks.iop.org/JPhysCM/19/343201

Abstract

The transition metal–oxygen (TM–O) bond appears prominently throughout chemistry and solid-state physics. Many materials, from biomolecules to ferroelectrics to the components of supernova remnants contain this bond in some form. Many of these materials' properties strongly depend on fine details of the TM–O bond and intricate correlation effects, which make accurate calculations of their properties very challenging. We present quantum Monte Carlo, an explicitly correlated class of methods, to improve the accuracy of electronic structure calculations over that of more traditional methods like density functional theory. We find that unlike s–p type bonding, the amount of hybridization of the d–p bond in TM–O materials is strongly dependent on electronic correlation.

(Some figures in this article are in colour only in the electronic version)

Contents

[1] Present address: 366 Le Conte Hall, #3700, Berkeley, CA 94720, USA.

1. Introduction

Transition metal chemistry is a particularly exciting area of research, with applications from astrophysics to biology to potential inexpensive high-efficiency solar cells and high-temperature superconductivity. Because of the partially filled d shell, transition metals can form many types of bonds and can also exhibit ferroelectric and ferromagnetic ordering. Transition metal oxides (TMOs) are particularly interesting because they are one of the most common transition metal complexes, and they exhibit most of the above effects. This rich physics is quite difficult to describe theoretically, however, since electronic correlation is very strong in these materials. Current approximate density functional theories tend to perform quite poorly on transition metals, particularly in comparison to its quite good accuracy on elements with s- and p-type bonding. Problematic quantities are not hard to find; they include the dipole moment in molecules, binding (or cohesive energies), the lattice constants of perovskites, high-pressure behavior, and band gaps/excitation energies.

Rather than attempting to improve the approximate density functional, quantum Monte Carlo (QMC) approaches take a different direction—explicitly treating the electronic correlation in a wavefunction-based approach, while maintaining reasonable scaling with system size. It can be made to scale from O(1) to O(N^3) in the number of electrons [1], depending on the quantity of interest. QMC attains very low upper-bound energies on medium-sized electronic problems (up to thousands of electrons at the time of writing), and has been used as a benchmark method on s–p systems [2]. Since it treats the electronic correlation explicitly in the many-body wavefunction, it is a promising method for strongly correlated TMO systems.

The goal of this review is to summarize the current state of the art of QMC as applied to TMOs. This is a fairly new field, with few calculations. Most of these calculations have benchmarked the method to determine the accuracy that one should expect. This accuracy has generally been quite high on most of the quantities studied, particularly for energetics. In the course of this work, it has also been determined what trial function (starting guess, as explained in the methods section) is necessary to obtain this accuracy. The upper-bound property of diffusion Monte Carlo has been critical in this success. By this, we have also learned that the electronic correlation in transition metal oxides is entangled with the d–p orbital hybridization in these materials.

2. Quantum Monte Carlo

The most common flavours of quantum Monte Carlo that have been used on TMOs are variational, diffusion, and reptation Monte Carlo (VMC, DMC, and RMC, respectively). We will summarize them here; one can find a more complete review in [3]. Another flavour, auxiliary field Monte Carlo (AFQMC) [4], has been used in a few calculations, but will not be discussed here.

VMC is a direct application of the variational theorem. We write the many-body wavefunction as a function of many-body coordinates $\mathbf{R} = [\mathbf{r}_1, \mathbf{r}_2, \ldots, \mathbf{r}_{N_e}]$ and a set of variational parameters $\mathbf{P}$. One then approximates the ground-state wavefunction by minimizing

the energy expectation value

$$E(\mathbf{P}) = \int \Psi^*(\mathbf{R}, \mathbf{P}) H \Psi(\mathbf{R}, \mathbf{P})\, \mathrm{d}\mathbf{R}, \tag{1}$$

assuming that the wavefunction is normalized. For a complicated variational ansatz such as we will introduce later, this integral cannot be evaluated analytically. One can, however, evaluate it using Monte Carlo by rearranging the integral to read

$$E(\mathbf{P}) = \int |\Psi(\mathbf{R}, \mathbf{P})|^2 \frac{H\Psi(\mathbf{R}, \mathbf{P})}{\Psi(\mathbf{R}, \mathbf{P})}\, \mathrm{d}\mathbf{R}. \tag{2}$$

Since $|\Psi(\mathbf{R}, \mathbf{P})|^2$ is a probability distribution function, one can sample it using Markov chain Monte Carlo and evaluate the energy expectation value as an average over the local energy $E_L(\mathbf{R}) = \frac{H\Psi(\mathbf{R})}{\Psi(\mathbf{R})}$. The lowest-energy approximate wavefunction is then found by minimizing the energy. In practice, a combination of energy and the variance of the local energy [5] or variance only [6] is optimized.

Many wavefunctions can be used with VMC, since the only requirement is that one can evaluate the wavefunction and its derivatives quickly. For the work covered in this paper, we start with a Slater determinant of one-particle orbitals, D, or a linear combination of Slater determinants. We then multiply D by the explicitly correlated inhomogeneous Jastrow correlation factor e^U to obtain the Slater–Jastrow variational wavefunction De^U. We write

$$U = \sum_{ijI} u(r_{iI}, r_{jI}, r_{ij}) \tag{3}$$

where the lower-case indices stand for electronic coordinates, and the upper-case indices are ionic coordinates. There is considerable choice on how to expand u; for concreteness, we show one expansion that performs well enough and has been applied to TMOs. The correlation factor is expanded in the Schmidt–Moskowitz form [7]:

$$\begin{aligned} u(r_{iI}, r_{jI}, r_{ij}) = &\sum_k c_k^{ei} a_k(r_{iI}) + \sum_m c_m^{ee} b_k(r_{ij}) \\ &+ \sum_{klm} c_{klm}^{eei} (a_k(r_{iI}) a_l(r_{jI}) + a_k(r_{jI}) a_l(r_{iI})) b_k(r_{ij}), \end{aligned}$$

where the a_k and b_k functions are written as

$$\frac{1 - z(r/r_{\mathrm{cut}})}{1 + \beta_k z(r/r_{\mathrm{cut}})}, \tag{4}$$

with different β_k for the different types of functions. The polynomial $z(x) = x^2(6 - 8x + 3x^2)$ is chosen so the functions go smoothly to zero at $r_{\mathrm{cut}} = 7.5$ bohr. The β_k and all the expansion coefficients c^{ei}, c^{ee}, and c^{eei} are optimized. If there are multiple determinants, their coefficients can also be optimized. We then use the VMC wavefunction as a trial function for RMC or DMC.

DMC and RMC are based on the so-called imaginary time Schrödinger equation

$$-\frac{\mathrm{d}\Psi(\mathbf{R}, \tau)}{\mathrm{d}\tau} = (H - E_0)\Psi(\mathbf{R}, \tau), \tag{5}$$

which has a steady-state solution Φ_0, the lowest energy eigenfunction with eigenvalue E_0 as long as $\Psi(\mathbf{R}, 0)$ has a non-zero overlap with Φ_0. All non-steady-state solutions converge exponentially to the eigenstate Φ_0 as τ goes to infinity. Transforming to an integral equation, we have

$$\Phi_0(\mathbf{R}_1) = \lim_{\tau \to \infty} \int G(\mathbf{R}_1, \mathbf{R}_0, \tau) \Psi_{\mathrm{T}}(\mathbf{R}_0)\, \mathrm{d}\mathbf{R}_0, \tag{6}$$

where G is the Green's function of the imaginary-time Schrödinger equation and $\Psi_{\mathrm{T}}(\mathbf{R}_0)$ is the trial wavefunction that we obtain from VMC. Solving for the exact G for large τ is as difficult as

solving for Φ_0, so we choose some constant small value of τ for which we know G accurately (for example, see [3, 8]), and compound the operations (suppressing the τ dependence of G):

$$\Phi_0(\mathbf{R}) = \lim_{n\to\infty} \int G(\mathbf{R}, \mathbf{R}_n) \dots G(\mathbf{R}_1, \mathbf{R}_0)\Psi_T(\mathbf{R}_0)\, d\mathbf{R}_0\, d\mathbf{R}_1 \dots d\mathbf{R}_n. \tag{7}$$

Each application of G is interpreted as a stochastic process, in the same way that the diffusion equation can be mapped onto Brownian particles and vice versa (in fact, for a free particle, the Hamiltonian is $-\frac{1}{2}\nabla^2$ and the simulation is a diffusion process).

DMC performs a simulation of these random particles for large n. All implementations of DMC use a particularly clever importance sampling transformation by multiplying the imaginary-time Schrödinger equation (equation (5)) by the trial function $\Psi_T(\mathbf{R})$ and working with the time-dependent function $\Psi_T(\mathbf{R})\Psi(\mathbf{R}, \tau)$. Since the time dependence is the same, it eventually obtains samples distributed according to the probability distribution function $P_{R_\infty}(\mathbf{R}) = \Phi_0(\mathbf{R})\Psi_T(\mathbf{R})$. This transformation improves the efficiency of the calculation by several orders of magnitude [3] by using information that we already have about the ground state in the form of a trial function. The final probability distribution function can be used to evaluate the ground-state energy as follows:

$$\langle E_0 \rangle = \int d\mathbf{R}\Psi_T(\mathbf{R})\Phi_0(\mathbf{R})\frac{H\Psi_T(\mathbf{R})}{\Psi_T(\mathbf{R})}, \tag{8}$$

since Φ_0 is an eigenstate of H and H can operate forwards or backwards. Any operators that do not commute with the Hamiltonian will have expectation values that are biased, only becoming unbiased in the limit of $\Psi_T = \Phi_0$.

We can remove the error in these operators by using reptation Monte Carlo [9, 10], where the random walk is performed in the space of paths: $s = [\mathbf{R}_0, \mathbf{R}_1, \dots, \mathbf{R}_{n-1}, \mathbf{R}_n]$. We sample the path probability distribution

$$\Pi(s) = \Psi_T(\mathbf{R}_0)G(\mathbf{R}_0, \mathbf{R}_1) \dots G(\mathbf{R}_{n-1}, \mathbf{R}_n)\Psi_T(\mathbf{R}_n). \tag{9}$$

This can be interpreted in several different ways. If we examine the distribution at $\mathbf{R}_0$, we can view the samples of Green's functions as acting on $\Psi_T(\mathbf{R}_n)$, and therefore $P_{R_0}(\mathbf{R}_0) = \Psi_T(\mathbf{R}_0)\Phi_0(\mathbf{R}_0)$. This is the same distribution as we obtain in DMC as the path length goes to infinity. Alternatively, since G is symmetric on exchange of the two $\mathbf{R}$ coordinates, the probability distribution of $\mathbf{R}_n$ is the same. Finally, we can split the path in two, one projecting on $\Psi_T(\mathbf{R}_0)$, and the other projecting on $\Psi_T(\mathbf{R}_n)$. We then have

$$\begin{aligned} P_{R_{n/2}}(\mathbf{R}_{n/2}) = {} & (G(\mathbf{R}_{n/2}, \mathbf{R}_{n/2-1}) \dots G(\mathbf{R}_1, \mathbf{R}_0)\Psi_T(\mathbf{R}_0)) \\ & \times (G(\mathbf{R}_{n/2}, \mathbf{R}_{n/2+1}) \dots G(\mathbf{R}_{n-1}, \mathbf{R}_n)\Psi_T(\mathbf{R}_n)) = \Phi_0^2(\mathbf{R}_{n/2}) \end{aligned}$$

for $n \to \infty$, which allows us to obtain correct expectation values of operators that do not commute with the Hamiltonian.

3. Geometry optimization

In TMO materials, it is particularly useful to be able to optimize the geometry of the system within QMC. The usual way of doing this in mean-field calculations is to calculate the forces on the atoms and use one of many minimization routines. Unfortunately, there are not yet any reliable methods to calculate the force within diffusion Monte Carlo, despite much work in that direction [10–13]. These methods all require high-accuracy trial wavefunctions, which we usually do not have for transition metals. Thus, with the current state of the art, we are only able to optimize a few key degrees of freedom using the total energies from DMC calculations and line minimization. Even this must be done carefully because of the statistical uncertainty

in the DMC energy. What follows is the scheme used in the work presented here, which has been found to be quite robust.

According to Bayes' theorem, given a model M and a set of data D, the probability of the model given the set of data is

$$P(M|D) = \frac{P(D|M)P(M)}{P(D)}. \tag{10}$$

$P(D)$ is an unimportant normalization constant and $P(M)$ is called the prior distribution, which we are free to set to reflect the *a priori* probability distribution on the set of models. One usually sets $P(M) = 1$, the unbiased maximum entropy/least knowledge condition. In the case of normally distributed data on a set of points $\{x_1, x_2, \ldots, x_N\}$,

$$P(D|M) \propto \exp\left[-\sum_i (M(x_i) - D(x_i))^2/2\sigma^2(x_i)\right], \tag{11}$$

where $\sigma(x)$ is the statistical uncertainty of $D(x)$.

For example, in the case of bond lengths, we can limit our space of models to $M(x) = c_1 + c_2x + c_3x^2$, for x close to the minimum bond length. This is equivalent to setting the prior distribution equal to one for all quadratic functions and to zero for non-quadratic functions. One then calculates several data points $D(x)$ with statistical uncertainties $\sigma(x)$. The probability distribution function of the bond length b is then obtained by calculating the marginal distribution

$$p(b) = \frac{\int \delta(-c_2/2c_3 - b)P(D|M)P(M)\,\mathrm{d}c_1\,\mathrm{d}c_2\,\mathrm{d}c_3}{\int P(D|M)P(M)\,\mathrm{d}c_1\,\mathrm{d}c_2\,\mathrm{d}c_3}. \tag{12}$$

This integral is only three-dimensional, and as such could be calculated by a grid method, but it is convenient to calculate it by Monte Carlo, by sampling $P(D|M)P(M)$ and binning the bond length. The probability distribution function for the bond length is typically a Gaussian function to high accuracy, so it can be described as a mean value with a statistical uncertainty.

To make this scheme more efficient, we would like to calculate QMC energies as far away from the minimum as possible while still maintaining accuracy. This is because the energy changes much more quickly far from the minimum, which mitigates the stochastic uncertainties. That is, the energy scale is larger far from the minimum, so less precision is necessary. Thus, we should use a fitting function that is valid as far from the minimum as possible, while containing as few parameters as possible. For minimum-energy geometries, it has been found [14, 15] that the Vignet or modified Morse potentials are quite good for this purpose.

4. Approximations

4.1. Pseudopotentials

In QMC, we can increase the efficiency significantly by using pseudopotentials to replace the core electrons with an effective potential. This has the effect of removing the large fluctuations near the core, which do not contribute much to the valence electrons' correlation, which is important for chemical properties. This introduces two approximations in the technique: first, the pseudopotential itself, and second, the small localization error [16] in diffusion Monte Carlo.

It has been found that small-core pseudopotentials are necessary for high accuracy on transition metals [17, 18]. On the 3d metals, which are the primary focus in this paper, this means a Ne-core pseudopotential. The reason for this is that the 3d electrons occupy much the

same space as the semicore 3p and, to a lesser extent, the 3s electrons. Since the 3d electrons are strongly affected by bonding, they in turn interact with the semicore. This interaction will change with correlation and chemical environment, so we must include the semicore electrons in accurate electronic structure calculations. This is not unique to QMC and is generally done in density functional theory where high accuracy is needed [19].

4.2. Finite size errors

When performing calculations for extended systems such as crystals, it is necessary to introduce periodic boundary conditions. This is an approximation on two levels. The first is the standard one-body level that is corrected by using reciprocal-space sampling (i.e., k-points). The second level is inherent in a many-body correlated method, where the periodic boundary conditions force the electron to interact unphysically with its periodic image. This is similar to the finite simulation cell error in classical molecular dynamics simulation. This is typically corrected by either modifying the Coulomb interaction to remove the spurious interaction [20] or by a correction [21, 15]. Neither of these methods has clearly been demonstrated to be superior, and both methods or similar ones have been used successfully. Even with these corrections, a QMC calculation of an extended system usually involves on the order of 40 to 100 atoms, regardless of the size of the primitive cell, followed by extrapolation to infinite size.

4.3. Fixed node

The algorithms described above are exact when the wavefunction can be written as a positive function, since then $\Psi_T\Phi_0$ is a probability distribution function. For fermions, it is not usually the case that Ψ_T has the same zeros as the exact ground state, so we make the fixed-node approximation, where the nodal surface of the exact wavefunction are assumed to be the same as the trial wavefunction. This approximation typically results in recovering 90–95% of the correlation energy, and can be relaxed, but at the cost of exponential scaling of the system size [3].

Given that the pseudopotential localization approximation is usually quite small for energy differences [22], we are mostly concerned with the fixed-node error. The Jastrow factor does not change the nodes of the wavefunction, so in the method outlined above, the nodes (and thus the final accuracy) are fixed to be the nodes of the Slater determinant of orbitals from the mean-field method. It is currently not feasible to vary the orbital expansion directly for a large system, since the number of parameters grows to the thousands for even moderately sized systems. However, partial optimizations can be done, and, as we shall see, they are very effective for TMO systems.

5. TMO molecules

Simple molecular systems are excellent starting points for the study of transition metal oxides, since they are small enough to study carefully in a reasonable amount of time, and are also treatable by accurate but expensive quantum chemistry techniques like the coupled cluster. This provides an additional much-needed data point to compare the accuracy of the various electronic structure methods.

5.1. Near-optimal one-particle orbitals

Wagner and Mitas [23] performed the first calculations using DMC on simple two-atom transition metal oxides (TiO and MnO), and found a strong dependence of the calculated

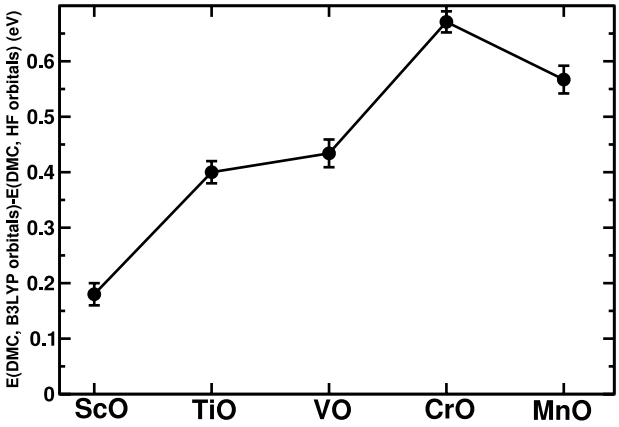

Figure 1. The energy gain in DMC from using B3LYP orbitals as a function of the metal monoxide. The line is a guide to the eye. Taken from [25].

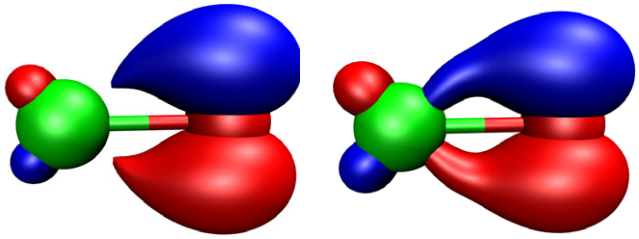

Figure 2. The d–p hybridization orbital (doubly occupied) isosurface for TiO in Hartree–Fock (left) and B3LYP (right). B3LYP enhances the hybridization significantly, which leads to lower energy in QMC. Figure generated using VMD and POV-Ray [26, 27].

binding energy on the orbitals used in the Slater determinant. They used the B3LYP hybrid DFT/Hartree–Fock functional, and varied the percentage of Hartree–Fock mixing. They found the optimal percentage to be very close to the semi-empirical value fitted by Becke for his B3PW potential [24]. We have plotted the energy gain of B3LYP orbitals versus Hartree–Fock for the first five transition metal monoxide molecules in figure 1. Upon examining the orbitals, they found a large difference in the d–p hybridization for both TiO (figure 2) and MnO. This is a direct consequence of the importance of electronic correlation in transition metals.

To understand the importance of the one-particle orbitals, one can conceptually divide the total energy in three parts, each described by a different part of the wavefunction:

- One-body and antisymmetry: the Slater determinant.
- Two-body electron interaction: Jastrow factor.
- Higher orders: implicit diffusion Monte Carlo wavefunction.

The first part, the Slater determinant, determines the nodes of the wavefunction and therefore the ultimate accuracy of the calculation. Empirically, in materials containing only s- and p-type elements, these three parts are almost independent of each other—the Hartree–Fock orbitals are close to optimal for a Slater–Jastrow wavefunction. In transition metal oxides, however, this situation changes, and the two-body and higher interactions are strong enough to change the one-body part significantly. In TMOs, this effect seems to be largely in the d–p hybridization between oxygen and the transition metal.

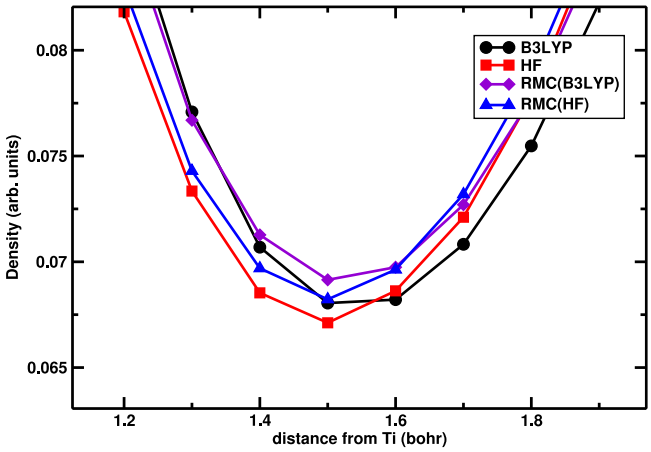

Figure 3. The density of the TiO molecule projected onto the Ti–O axis in the bonding region for various methods.

Table 1. Binding energies of the first five transition metal monoxides by different theoretical methods, along with RMS deviations from the experiment (all in eV). Statistical uncertainties in units of 10^{-2} eV are shown in parentheses for Monte Carlo and experimental results. Zero-point energy corrections are estimated to be much less than the uncertainty in experiment. There are too few AFQMC data to calculate meaningful RMS values.

Method	ScO	TiO	VO	CrO	MnO	RMS
LDA [28]	9.09	9.13	8.48	6.26	6.51	2.19
CCSD(T) [29]	6.71	6.64	6.13	4.20	3.43	0.31
TPSSh [28]	7.11	7.18	6.44	4.45	4.62	0.38
DMC [25]	7.06(3)	6.81(3)	6.54(3)	3.98(2)	3.66(3)	0.21
AFQMC [30]	—	7.02(21)	—	—	3.79(34)	—
Exp [31]	7.01(12)	6.92(10)	6.44(20)	4.41(30)	3.83(8)	0

By using the reptation Monte Carlo algorithm, we can obtain the unbiased one-particle density within the fixed-node approximation (figure 3), which gives further insight into the importance of correlation in the one-particle density. QMC tends to enhance the density in the bonding region (the hybridization) over both Hartree–Fock and B3LYP, but is not able to completely repair the erroneous Hartree–Fock density because of the fixed-node approximation. This is the reason for the large energy gain from using B3LYP orbitals to fix the nodal surface.

5.2. Energetic performance

The total energy of a system is quite important for determination of lowest-energy spin states, competing phases, reactions, etc, and is a place where traditional density functional theory has encountered difficulties on transition metal oxides. In table 1, we compare the binding energy obtained by DMC using B3LYP orbitals to DFT in the local density approximation (LDA), meta-GGA TPSSh and coupled cluster. We find excellent accuracy, with the root-mean-square (RMS) deviations of DMC within the experimental uncertainty for most materials. CrO is the only molecule with a large deviation from experiment; however, it is not very far outside the experimental uncertainty. DMC is also able to consistently obtain a minimum-energy bond length with errors below 0.01 Å (table 2), better than any other published result.

Table 2. Bond lengths in Å for the first five transition metal monoxide molecules.

Method	ScO	TiO	VO	CrO	MnO	RMS
LDA [28]	1.644	1.597	1.564	1.584	1.602	0.033
CCSD(T) [29]	1.680	1.628	1.602	1.634	1.66	0.011
TPSSh [28]	1.659	1.613	1.582	1.612	1.628	0.012
DMC [25]	1.679(2)	1.612(3)	1.587(3)	1.617(4)	1.652(4)	0.008
Exp [31]	1.668	1.623	1.591	1.621	1.648	0

Table 3. Dipole moments in Debye. The fixed-node RMC results have been obtained with a single determinant of B3LYP orbitals. See text for an analysis of the errors involved for the case of TiO.

Method	ScO	TiO	VO	CrO	MnO
LDA [28]	3.57	3.23	3.10	3.41	—
CCSD(T) [29]	3.91	3.52	3.60	3.89	4.99
TPSSh [28]	3.48	3.43	3.58	3.97	—
RMC [25]	4.61(5)	4.11(5)	4.64(5)	4.76(4)	5.3(1)
Exp [37]	4.55	3.34(1) [38]	3.355	3.88	—

5.3. Dipole moments

While energetics are very important for electronic structure calculations, one is also often interested in non-energetic properties, such as dipole moments. There has been little work done on such things within QMC, even in the context of simpler s and p systems. To our knowledge, the only study of dipole moments other than on TMOs is of the CO molecule [32]. A primary reason for this lack of calculations is that, until the development of RMC, there has not been an easy to implement method to obtain expectation values without the mixed-estimator bias. The commonly used methods, pure diffusion Monte Carlo and forward-walking [33–35], do not scale well with the system size [36], since they suffer from increased fluctuations of weights as the number of particles increases. One can also use extrapolated estimation, where the expectation value of an operator is estimated as $\langle \mathcal{O} \rangle = 2\langle \mathcal{O} \rangle_{\mathrm{DMC}} - \langle \mathcal{O} \rangle_{\mathrm{VMC}}$, but that method introduces an additional approximation that one would like to avoid if possible.

RMC, on the other hand, scales quite well, and is easily applicable to medium-sized systems such as TMO molecules. As we have noticed above, the electronic correlation and hybridization are very intertwined, and therefore, the electronic correlation and dipole moment are also closely related. In table 3, we report the dipole moments for the first five transition metal monoxides using RMC with B3LYP orbitals. RMC obtains dipole moments much higher than that found in experiment, which is somewhat surprising given the high accuracy seen in energetic properties. We will explore the fixed-node approximation and its effect on the dipole moment in the next section.

5.4. Beyond the Slater–Jastrow form

In this section, we explore one of the biggest advantages of the QMC method—the ability to go beyond a Slater–Jastrow trial function if needed. As we saw in the previous section, RMC with the Slater–Jastrow trial function does not obtain dipole moments in agreement with experiment. The dipole moment is very sensitive to electronic correlation, and we wish to perform as accurate a calculation as possible to approach the exact value. We can do this in QMC by expanding the wavefunction in determinants. We write the trial wavefunction as

$$\Psi_{\mathrm{T}}(\mathbf{R}) = \left(\sum_i c_i D_i \right) e^U, \tag{13}$$

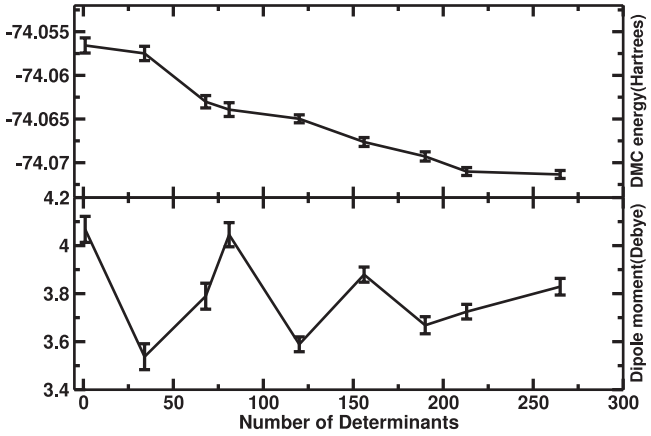

Figure 4. The number of determinants versus the energy and dipole moment for TiO. The dipole moments are shifted downwards by 0.1 Debye to correct for the pseudopotential error.

where the D_i are determinants of one-particle orbitals, e^U is the Jastrow factor, and the c_i are variational parameters. These determinants and the initial coefficients are taken from a configuration interaction (CI) calculation, and the coefficients are reoptimized using variational Monte Carlo in the presence of the Jastrow factor.

This last reoptimization step is crucial, since the DMC energy increases if the CI coefficients are kept constant. This is a result of the strong correlation of these systems—the first-order correlations are taken care of by the Jastrow factor, which the CI calculation tries to describe (inefficiently) with determinants.

In figure 4, we see the convergence of this expansion for TiO. The energy has a smooth convergence in the number of determinants, but the dipole moment oscillates significantly, with smaller oscillations as the number of determinants increases. The final result is approximately 3.8(1) Debye, a significant change from the Slater–Jastrow trial wavefunction, but still quite far from the experimental value of 3.34(1) Debye. While this calculation is probably not at the exact limit, the dipole moment does not appear to change enough to reconcile with experiment. Somewhat reassuringly, though, the coupled-cluster value also predicts a larger value for the dipole moment, so it is possible that the experiment may be in error. More studies of non-energy properties using quantum Monte Carlo are sorely needed, however, to obtain an estimate of the expected accuracy.

6. Solids

Calculations on extended TMO systems using QMC are particularly challenging, since QMC suffers not only from one-body finite-size effects (i.e., that described by k-point sampling), but also from many-body finite-size effects, which require large simulation cells. For this reason, complete studies such as those reported above for molecules are not usually attainable, and most work is still in progress. We will discuss a few preliminary studies and a few private communications of work that remains unpublished at the time of this writing. Clearly, the details of the calculations may change, so this section is meant more as a comment on the current state of the art.

Using QMC, there have been studies of the antiferromagnet NiO [39, 40] and MnO [17]. Except for Tanaka [39], who performed a very rough optimization of the lattice constant within

Table 4. Cohesive energies for several materials using QMC, all calculated per formula cell. Also listed are the optimal mean-field orbitals if reported. LDA is the local density approximation of DFT, and PBE0 is a hybrid functional.

Material	DMC binding energy (eV)	Experimental	Mean-field orbitals
NiO [40]	9.442(2)	9.5	Hartree–Fock
MnO [17]	9.40(5)	9.5	
$BaTiO_3$ [42]	31.2(3)	31.57	LDA
FeO [41]	9.47(4)	9.7	PBE0 [43]

variational Monte Carlo, all the published studies calculated only the cohesive energy, which comes quite close to experiment (table 4) for the materials available. In the very recent work of Kolorenc and Mitas [41], they obtain similar accuracy for the cohesive energy of FeO and also obtain the correct ordering of phases for that material, which DFT mispredicts. In most of these materials, researchers have found a large dependence on the mean-field orbitals used, with the optimal orbitals ranging from Hartree–Fock to LDA. Apparently, there is no universal optimal mean-field method.

Wagner and Mitas [42] have also reported using the Bayesian optimization scheme to find the minimum-energy lattice constant of $BaTiO_3$, which is well known to be underestimated by over 1% in the LDA to density functional theory, and overestimated by a similar margin in the gradient corrections. This 1% error in the lattice constant can affect the calculated spontaneous polarization up to 50%, so even this small error is not acceptable for a truly first-principles description of this material. DMC obtains a cubic lattice constant in error only by 0.015 ± 0.005 Å, which is somewhat less than half a per cent, a significant improvement over the density functional results. Also, in $BaTiO_3$, there is an energy gain in DMC of $\sim$1 eV/formula cell by using LDA orbitals instead of Hartree–Fock orbitals, and they report that it is due to a similar change in d–p hybridization that is seen in the transition metal monoxide molecules.

7. Conclusions

On the systems that have been tested thus far, QMC offers unprecedented accuracy in a completely first-principles and scalable method, particularly in the energetics of the systems. The d–p hybridization of transition metal oxides is strongly affected by electronic correlation. Using QMC methods, we can clearly see this, both by investigating the minimum-energy orbitals and by examining the one-particle density and dipole moment within QMC. The dipole moment in particular is strongly affected by the level of correlation present in the quantum mechanics approximation.

On TMO molecules, we have a significant gain in the total energy on expansion into determinants, of about 0.5 eV. This means that we are relying on cancellation of errors for the high accuracy of QMC, although to a much lesser degree than post-Hartree–Fock approaches and DFT. We see this error in the dipole moment, which does not benefit from cancellation of errors. On the molecules, however, we can use a brute-force approach by expanding in determinants and come quite close to the true ground state. However, this kind of expansion will ultimately fail for large systems, since the number of determinants grows very quickly with system size. In order to reliably check the QMC results, it is vital to develop new reasonably scaling wavefunctions that go beyond the Slater–Jastrow form. Some work has been done in this direction with the RVB [44], Pfaffian [45], and backflow [46, 47] wavefunctions in QMC. These wavefunctions' accuracy should be tested on TMO systems in the future. Equally important are optimization schemes within VMC that can systematically minimize the energy

with respect to the wavefunctions' parameters, despite the stochastic nature of VMC, which is under serious investigation [5, 48]. Finally, we need to be able to calculate forces within QMC accurately and efficiently. The current state of the art is not sufficient to treat transition metal oxides [15], and the Bayesian method of geometry optimization is only efficient for a few dimensions.

The future looks promising for QMC calculations of TMO solids, with the only drawback that the calculations are very expensive on today's computers, since one must use a large supercell. However, the scaling with system size is quite favourable, and QMC is very easy to operate in parallel, so it can take advantage of low-cost processors. It has already been shown for a few important transition metal oxide solids that QMC can obtain binding energies and other energetic properties with excellent accuracy, well worth the additional cost when high accuracy is needed. It remains to be seen how well the method performs for non-energetic properties, and what sort of trial wavefunctions are necessary to obtain even higher accuracy.

Acknowledgments

I would like to acknowledge Lubos Mitas, Jindrich Kolorenc, and Michal Bajdich for their support and discussions in much of the work discussed in this article, as well as E Ertekin and V Srinivasan for their comments on the article itself. I am also grateful for an NSF Graduate Research Fellowship and NSF grant EAR-0530110 for funding during the course of this work.

References

[1] Williamson A J, Hood R Q and Grossman J C 2001 Linear-scaling quantum Monte Carlo calculations *Phys. Rev. Lett.* **87** 246406
[2] Grossman J C 2002 Benchmark quantum Monte Carlo calculations *J. Chem. Phys.* **117** 1434
[3] Foulkes W M C, Mitas L, Needs R J and Rajagopal G 2001 Quantum Monte Carlo simulations of solids *Rev. Mod. Phys.* **73** 33
[4] Zhang S and Krakauer H 2003 Quantum Monte Carlo method using phase-free random walks with Slater determinants *Phys. Rev. Lett.* **90** 136401
[5] Umrigar C J and Filippi C 2005 Energy and variance optimization of many-body wavefunctions *Phys. Rev. Lett.* **94** 150201
[6] Umrigar C J, Wilson K G and Wilkins J W 1988 Optimized trial wavefunctions for quantum Monte Carlo calculations *Phys. Rev. Lett.* **60** 1719
[7] Schmidt K E and Moskowitz J W 1990 Correlated Monte Carlo wavefunctions for the atoms He through Ne *J. Chem. Phys.* **93** 4172–8
[8] Umrigar C J, Nightingale M P and Runge K J 1993 A DMC method with small timestep errors *J. Chem. Phys.* **99** 2865
[9] Baroni S and Moroni S 1999 Reptation quantum Monte Carlo: a method for unbiased ground-state averages and imaginary-time correlations *Phys. Rev. Lett.* **82** 4745
[10] Pierleoni C and Ceperley D M 2005 *Chem. Phys. Chem.* **6** 1872
[11] Filippi C and Umrigar C J 2000 Correlated sampling in quantum Monte Carlo: a route to forces *Phys. Rev.* B **61** R16291
[12] Assaraf R and Caffarel M 2003 Zero-variance zero-bias principle for observables in quantum Monte Carlo: application to forces *J. Chem. Phys.* **119** 10536
[13] Chiesa S, Ceperley D M and Zhang S 2005 Accurate, efficient, and simple forces computed with quantum Monte Carlo methods *Phys. Rev. Lett.* **94** 036404
[14] Maezono R, Ma A, Towler M D and Needs R J 2007 Equation of state and Raman frequency of diamond from quantum Monte Carlo simulations *Phys. Rev. Lett.* **98** 025701
[15] Wagner L K 2006 Quantum Monte Carlo for transition metal systems: method developments and applications *Thesis*
[16] Mitas L, Shirley E L and Ceperley D M 1991 Nonlocal pseudopotentials and diffusion Monte Carlo *J. Chem. Phys.* **95** 3467
[17] Lee J W, Mitas L and Wagner L K 2004 Quantum Monte Carlo study of MnO solid *Preprint* cond-mat/0411247

[18] Flad H-J and Dolg M 1997 Probing the accuracy of pseudopotentials for transition metals in quantum Monte Carlo calculations *J. Chem. Phys.* **107** 7951–9
[19] King-Smith R D and Vanderbilt D 1994 First-principles investigation of ferroelectricity in perovskite compounds *Phys. Rev.* B **49** 5828–44
[20] Kent P R C, Hood R Q, Williamson A J, Needs R J, Foulkes W M C and Rajagopal G 1999 Finite-size errors in quantum many-body simulations of extended systems *Phys. Rev.* B **59** 1917–29
[21] Chiesa S, Ceperley D M, Martin R M and Holzmann M 2006 Finite-size error in many-body simulations with long-range interactions *Phys. Rev. Lett.* **97** 076404
[22] Casula M, Filippi C and Sorella S 2005 Diffusion Monte Carlo method with lattice regularization *Phys. Rev. Lett.* **95** 100201
[23] Wagner L K and Mitas L 2003 A quantum Monte Carlo study of electron correlation in transition metal oxygen molecules *Chem. Phys. Lett.* **370** 412
[24] Becke Axel D 1993 Density-functional thermochemistry. III. The role of exact exchange *J. Chem. Phys.* **98** 5648–52
[25] Wagner L K and Mitas L 2007 Energetics and dipole moment of transition metal monoxides by quantum Monte Carlo *J. Chem. Phys.* **126** 034105
[26] Humphrey W, Dalke A and Schulten K 1996 Vmd-visual molecular dynamics *J. Mol. Graph.* **14.1** 33–8
[27] 2004 Persistence of Vision Pt. Ltd. Persistence of vision (tm) raytracer
[28] Furche F and Perdew J P 2006 The performance of semilocal and hybrid density functionals in 3d transition-metal chemistry *J. Chem. Phys.* **124** 044103
[29] Bauschlicher C W and Maitre P 1995 Theoretical study of the first transition row oxides and sulfides *Theor. Chim. Acta* **90** 189
[30] Al-Saidi W A, Krakauer H and Zhang S 2006 Auxiliary-field quantum Monte Carlo study of TiO and MnO molecules *Phys. Rev.* B **73** 075103
[31] Merer A J 1989 *Annu. Rev. Phys. Chem.* **40** 407
[32] Schautz F and Flad H-J 1999 Quantum Monte Carlo study of the dipole moment of CO *J. Chem. Phys.* **110** 11700–7
[33] Caffarel M and Claverie P 1988 Development of a pure diffusion quantum Monte Carlo method using a full generalized Feynman–Kac formula I. Formalism *J. Chem. Phys.* **88** 1088–99
[34] Caffarel M and Claverie P 1988 Development of a pure diffusion quantum Monte Carlo method using a full generalized Feynman–Kac formula II. Applications to simple systems *J. Chem. Phys.* **88** 1100–9
[35] *Monte Carlo Methods in Ab Initio Quantum Chemistry* 1994 (Singapore: World Scientific)
[36] Assaraf R, Caffarel M and Khelif A 2000 Diffusion Monte Carlo methods with a fixed number of walkers *Phys. Rev.* E **61** 4566–75
[37] Steimle T C 2000 *Int. Rev. Phys. Chem.* **19** 455
[38] Steimle T C and Virgo W 2003 The permanent electric dipole moments of the X3D, E3P, A3P and B3P states of titanium monoxide, TiO *Chem. Phys. Lett.* **381** 30
[39] Tanaka S 1995 Variational quantum Monte-Carlo approach to the electronic structure of NiO *J. Phys. Soc. Japan* **64** 4270
[40] Needs R J and Towler M D 2003 The diffusion quantum Monte Carlo method: designing trial wavefunctions for NiO *Int. J. Mod. Phys.* B **17** 5425 (also available at http://www.tcm.phy.cam.ac.uk/∼mdt26/publications.html)
[41] Kolorenc J and Mitas L 2007 private communication
[42] Wagner L K and Mitas L 2006 unpublished work
[43] Adamo C and Barone V 1999 Toward reliable density functional methods without adjustable parameters: the pbe0 model *J. Chem. Phys.* **110** 6158–70
[44] Casula M and Sorella S 2003 Geminal wavefunctions with Jastrow correlation: a first application to atoms *J. Chem. Phys.* **119** 6500–11
[45] Bajdich M, Mitas L, Drobny G, Wagner L K and Schmidt K E 2006 Pfaffian pairing wavefunctions in electronic-structure quantum Monte Carlo simulations *Phys. Rev. Lett.* **96** 130201
[46] Drummond N D, Rios P L, Ma A, Trail J R, Spink G G, Towler M D and Needs R J 2006 Quantum Monte Carlo study of the Ne atom and the Ne+ ion *J. Chem. Phys.* **124** 224104
[47] Rios P L, Ma A, Drummond N D, Towler M D and Needs R J 2006 Inhomogeneous backflow transformations in quantum Monte Carlo calculations *Phys. Rev.* E **74** 066701
[48] Umrigar C J, Toulouse J, Filippi C, Sorella S and Hennig R G 2007 Alleviation of the fermion-sign problem by optimization of many-body wavefunctions *Phys. Rev. Lett.* **98** 110201