Quantum annealing of an Ising spin-glass by Green's function Monte Carlo

Lorenzo Stella[1,*] and Giuseppe E. Santoro[1,2]
[1]*International School for Advanced Studies (SISSA) and INFM Democritos National Simulation Center, Via Beirut 2-4, I-34014 Trieste, Italy*
[2]*International Centre for Theoretical Physics (ICTP), P.O. Box 586, I-34014 Trieste, Italy*
(Received 5 April 2006; revised manuscript received 8 December 2006; published 13 March 2007)

We present an implementation of quantum annealing (QA) via lattice Green's function Monte Carlo (GFMC), focusing on its application to the Ising spin glass in transverse field. In particular, we study whether or not such a method is more effective than the path-integral Monte Carlo- (PIMC) based QA, as well as classical simulated annealing (CA), previously tested on the same optimization problem. We identify the issue of importance sampling, i.e., the necessity of possessing reasonably good (variational) trial wave functions, as the key point of the algorithm. We performed GFMC-QA runs using such a Boltzmann-type trial wave function, finding results for the residual energies that are qualitatively similar to those of CA (but at a much larger computational cost), and definitely worse than PIMC-QA. We conclude that, at present, without a serious effort in constructing reliable importance sampling variational wave functions for a quantum glass, GFMC-QA is not a true competitor of PIMC-QA.

DOI: 10.1103/PhysRevE.75.036703 PACS number(s): 02.70.Uu, 02.70.Ss, 75.10.Nr, 07.05.Tp

I. INTRODUCTION

Quantum annealing (QA) [1,2] is an optimization method based on the idea of searching for the ground state of some *classical* Hamiltonian by adiabatically switching off an appropriate source of quantum fluctuations, in much the same way as temperature would do in classical thermal annealing (CA) [3]. The same approach is also known as adiabatic quantum computation [4] in the quantum information community.

A very popular QA approach is based on an imaginary-time quantum Monte Carlo (QMC) implementation, i.e., the path-integral Monte Carlo (PIMC) approach. A certain success in the application of PIMC-QA has been obtained in most of the cases studied: the folding of off-lattice polymer models [5,6], the random Ising model [7,8], and the random-field Ising model ground-state search [9], Lennard-Jones clusters optimization [10,11], and the traveling salesman problem [12]. Nevertheless, a counterexample exists [13], where PIMC-QA performs definitely *worse* than simple CA: the 3-Boolean-satisfiability (3-SAT) problem [14], which is a prototype of a large class of hard combinatorial optimization problem [the so-called nondeterministic polynomial (NP) complete class (see Ref. [15])].

In order to understand its features in detail, in a recent paper [16] we have studied the PIMC-QA performance focusing our attention on a simple, but highly instructive, toy problem: the double-well potential. There, we learned a few possible dangers of the PIMC-QA method: (i) The unavoidably finite temperature T, which provides a thermal lower limit to the average residual energies (i.e., the energy minimum estimated by QA minus its exact value). (ii) The severe sampling difficulties (i.e., ergodicity breaking), which possibly occur close to a Landau-Zener crossing [3].

We propose here to investigate a QA algorithm based on a different QMC technique, as an alternative to PIMC-QA. A natural choice is provided by *Green's function Monte Carlo* (GFMC). GFMC is different from PIMC in that it can directly sample the ground state (i.e., $T=0$) of a quantum Hamiltonian, avoiding, in principle, the first of the PIMC drawbacks. However—contrary to PIMC—GFMC is not a completely unbiased scheme: It demands the knowledge of a good approximation to the ground-state wave function to achieve proper efficiency [17]. So, the result of this comparison of GFMC against PIMC is *a priori* not clear.

A sensible test for this new GFMC-QA algorithm is provided by the random Ising model, a real optimization problem already addressed through PIMC-QA, as well as CA, in the recent past [7,8]. The Hamiltonian of the random Ising model in transverse field is

$$H(\Gamma) = -\sum_{\langle i,j\rangle} J_{i,j}\sigma_i^z\sigma_j^z - \Gamma\sum_i \sigma_i^x = H_{cl} + H_{kin}, \qquad (1)$$

where $\Sigma_{\langle i,j\rangle}$ indicates a sum over nearest neighbors, $J_{i,j}$ are random nearest-neighbor Ising coupling constants, and σ_i^z, σ_i^x are Pauli's matrices on lattice site i. If we denote by $\{S_i\}$ a generic spin configuration (where $S_i=\pm 1$ are the eigenvalues of the σ_i^z matrix), the challenging (classical) function we want to minimize is just given by the first term in Eq. (1), $E_{cl}(\{S_i\})=\langle\{S_i\}|H_{cl}|\{S_i\}\rangle$, which here plays the role of a *potential energy*. The second term in Eq. (1), $H_{kin}=-\Gamma\Sigma_i\sigma_i^x$, is the source of quantum fluctuations, which then plays the role of a *kinetic energy*. In particular, we will concentrate our efforts on a representative instance, which has been extensively analyzed in Refs. [7,8]. This is a two-dimensional (2D) instance, on a 80×80 square lattice, being the couplings $J_{i,j}$ drawn from a flat distribution in $[-2,2]$. The choice of a 2D case is motivated by the fact that the Ising glass ground-state search without a longitudinal magnetic field—is actually a

*Present address: Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, United Kingdom.

polynomial problem [18] in 2D, where very efficient *branch and cut* algorithms (see Ref. [19]) are known to find the true optimal-state energy E_{GS}. This provides for a clear benchmark for an annealing study.

The goal is to simulate the imaginary-time evolution defined by Eq. (1) starting from a large value of the transverse field Γ, which is then monotonically ramped down to zero in a certain *finite* time τ. So, Γ is the *annealing parameter* of the simulation. The adiabatic theorem [20] ensures that for "slow enough" evolution, the exact unperturbed (i.e., $\Gamma=0$) ground state should be eventually recovered. We note that the annealing schedule might also have no characteristic time τ—for instance, a power-law decrease of Γ. Remarkable convergence theorems [21] have been recently proven for both PIMC-QA and GFMC-QA, if a power-law annealing schedule is allowed: nevertheless, the bound on the power-law exponent decreases as the size of the systems increases [21], which might make a truly convergent power-law annealing impracticably slow. We will use here a linear annealing $\Gamma(t)=\Gamma_0(1-t/\tau)$ with a finite τ to compare directly our results with previous publications using the same schedule [7,8], originally inspired by its being close in spirit to the actual experimental realization of QA in Ising ferroglasses [22,23].

The rest of the paper is organized as follows: In Sec. II we present the main ideas of a GFMC-based QA approach. In Sec. III we present the results of our variational studies, showing the inherent difficulties associated to the selection of good wave functions for a disordered quantum system. In Sec. IV we discuss the GFMC-QA results and compare them with previous PIMC-QA and CA data on the same problem. Finally, in Sec. V we give some concluding remarks.

II. GREEN'S FUNCTION MONTE CARLO QUANTUM ANNEALING: IDEAS

Every QA algorithm is based on the iterative solution of an appropriate quantum dynamics, which for the problem at hand, is the Schrödinger dynamics associated to the Hamiltonian in Eq. (1). The basic evolution step (from now on $\hbar=1$) is

$$\psi(t+\Delta t)=e^{-iH[\Gamma(t)]\Delta t}\psi(t), \tag{2}$$

which should be iterated during the simulation until the algorithm finally stops when the annealing parameter Γ is reduced to zero. The total number of annealing steps is referred to as the *annealing time* of the simulation.

As argued in Ref. [24], an imaginary-time quantum evolution is (for our optimization purposes) equally good, and most likely even superior, to the standard real-time evolution. In practice, it is better to use the imaginary-time evolution operator, $\exp(-H[\Gamma(t)]\Delta t)$, instead of $\exp(-iH[\Gamma(t)]\Delta t)$ in Eq. (2): the imaginary-time evolution naturally tends to *filter out* the corresponding ground state of $H[\Gamma(t)]$ from the state it is applied to [17,24]. In this sense, the Green's function Monte Carlo (GFMC) is just a stochastic technique, which implements such a form of imaginary-time propagation [17]. However, the process underlying this iterated method is *not* a properly defined Markov chain [25], and so, it cannot be immediately simulated by a standard Monte Carlo approach [26]. In order to solve this problem, the phase space must be extended: a state is then defined by means of its position x—the configuration $x=\{S_i\}$ in the Hilbert space of the system—and its *weight* w—which is essentially related to $\psi(x)$—[the pair (x,w) is often referred to as a *walker*.] In practice, to improve convergence and stability, many walkers are evolved at the same time and then periodically reconfigured according to a well-defined stochastic process called *branching* [17,25]. The analogy of such a many-walker GFMC with a genetic-like algorithm is worth noting [27]. Each walker (x,w) plays the role of an individual that propagates (mutates) increasing or decreasing its *fitness*—represented by its weight w—with branching favoring the survival of those with highest fitness (largest w).

The final, important ingredient that makes the algorithm work is the so-called *importance sampling* [26]. It can also be seen, in the genetic analogy mentioned before, as a way of proposing mutations that—instead of being equally probable—are biased by a function which guides the system towards the most representative configurations. This is obtained by a guess of the exact ground state, the so-called *trial wave function* $\psi_T(x)$. As we see later, a good trial function can drastically improve the quality of a GFMC simulation [28]. In the next section we will describe two choices of $\psi_T(x)$ that we have tested for the Ising case, and the difficulties encountered.

III. VARIATIONAL WAVE FUNCTIONS FOR THE ISING SPIN GLASS

The first idea that comes to mind is a kind of "mean-field" (MF) wave function, made up of a product of single-site factors such as

$$|\psi_T^{(MF)}\rangle=\prod_{i=1}^{N}\left(\frac{e^{+h_i/2}|\uparrow\rangle_i+e^{-h_i/2}|\downarrow\rangle_i}{\sqrt{2\cosh(h_i)}}\right), \tag{3}$$

where $\{h_i\}$, the local fields in the z direction on each site i, are variational parameters to be optimized for each given value of the transverse field Γ. The optimization of the $\{h_i\}$ amounts to finding the minimum of the variational energy, $E_T^{(MF)}=-\Sigma_{\langle i,j\rangle}J_{i,j}m_im_j-\Gamma\Sigma_i\sqrt{1-m_i^2}$, where $m_i=\tanh(h_i)$ are the local magnetizations [27]. As it turns out, this optimization can be easily done only for large enough Γ, where the solution with $m_i=0$, representing all spins aligned along the x direction, is found. This solution survives down to some value Γ_{cr} of the transverse field, below which nontrivial solutions—i.e., with nonvanishing local magnetizations $m_i\neq 0$—start to appear. Because of the many similarities between our model in the low-Γ region and the well-known Weiss mean-field approach to the classical random Ising model [29], we expect the number of local minima to be very large. In a sense, minimizing $E_T^{(MF)}$ in the glassy phase ($\Gamma<\Gamma_{cr}$) is not easier than finding the classical ground state of the target problem at $\Gamma=0$, even if a continuous set of variables, $m_i\in(-1,1)$, is employed instead of the original discrete one, $S_i=\pm 1$. A numerical confirmation of this simple analogy is reported in Ref. [27].

A second, quite natural choice of ψ_T is a Boltzmann-type wave function of the form

$$\psi_T^{(\beta)}(\{S_i\}) \propto e^{-(\beta/2)E_{cl}(\{S_i\})}, \tag{4}$$

where β is the variational parameter to be optimized, and $E_{cl}(\{S_i\})$ (see Sec. I) is the classical energy of a given spin configuration $\{S_i\}$. Once again, for large Γ we expect to find $\beta=0$ (i.e., all configurations equally present in ψ_T), while by decreasing Γ, larger and larger values of β will favor configurations $\{S_i\}$ where the classical potential energy $E_{cl}(\{S_i\})$ has a local minimum until we get, for $\Gamma=0$, to the asymptotic limit $\beta\to\infty$ (ideally), required by a wave function that is perfectly localized in the global minimum (whose energy is indicated as $E_{GS}=N\epsilon_{GS}$, N being the number of sites).

To calculate the expectation value of the energy with the Boltzmann-type choice in Eq. (4), as a function of the single parameter β, we used a standard variational Monte Carlo (VMC) algorithm [17] with single spin-flip moves. Figure 1 shows (top panel) the optimal value β_{opt} of β, which minimizes the variational energy $E_{tot}^{(Boltz)}=\langle\psi_T^{(\beta)}|H|\psi_T^{(\beta)}\rangle$, for several values of the transverse field Γ. Surprisingly, β_{opt} saturates for small Γ to about $\beta_{opt}\approx 2$, instead of showing the expected $\beta_{opt}\to+\infty$ behavior. This is clearly an effect of a severe ergodicity loss of the VMC algorithm, which is not difficult to understand. For a given β the VMC samples the thermal distribution $|\psi_T^{(\beta)}(x)|^2=e^{-\beta E_{cl}(x)}$. We are therefore searching for the *effective temperature* $1/\beta$, which provides the best approximation to the wave function of the ground state of a quantum Ising glass at nonzero Γ. Now, from classical spin-glass physics [30] we know that a threshold energy E_{th} exists below which the system has a finite complexity, i.e., it displays an exponentially ($\sim\exp N$) large number of metastable minima. Close to this threshold energy, the relaxation of any local algorithm towards equilibrium becomes exceedingly slow (the algorithm gets stuck for a long time in every visited minimum) and the measured average quantities are no longer representative of their true thermodynamical values. Evidently, for $\Gamma\to 0$, the variational algorithm is not visiting the regions near the true minimum of the classical energy, but is wandering in a high-energy band of metastable states, separated by moderate energy barriers. In such a case, a small and finite value of β allows one to still overcome such barriers, so as to find slightly more favorable local minima, while perfect localization ($\beta\to+\infty$) in a *wrong* excited state would lead to an average bigger residual energy.

The central and bottom panels in Fig. 1 show the optimal variational energies $\epsilon_{tot}^{(Boltz)}=\langle\psi_T^{(\beta_{opt})}|H|\psi_T^{(\beta_{opt})}\rangle/N$, and the variational residual energy $\epsilon_{res}^{(Boltz)}=\langle\psi_T^{(\beta_{opt})}|H_{cl}|\psi_T^{(\beta_{opt})}\rangle/N-\epsilon_{GS}$ corresponding to the optimal β shown in the top panel, for several values of transverse field Γ. For large Γ values, the variational total energy (center panel) is linear in Γ, as it should be, since the transverse field kinetic term dominates in the quantum paramagnetic phase [see Eq. (1)], while the variational residual energy per site is of order 1. By decreasing Γ, we notice that the variational residual energy saturates, for small Γ, to finite nonzero values of order 0.03, in agreement with the previously noted saturation in the optimal β_{opt}. A closer inspection shows that the variational residual energy is actually *nonmonotonic* for $\Gamma<0.25$, again an artifact of sampling difficulties. Notice, however, that this saturation value is definitely below the best (down to $\Gamma=0.01$) results provided by the mean-field approach introduced above, which is of order 0.035 (shown for comparison by a dashed horizontal line) [27]. Therefore, with all its pitfalls, the Boltzmann-type trial wave function in Eq. (4) provides, at low Γ, an approximation of the true ground state, which is marginally better than that obtained by the mean-field *Ansatz*, Eq. (3). Moreover, $\psi_T^{(\beta)}$ is also much better behaved, and easier to optimize.

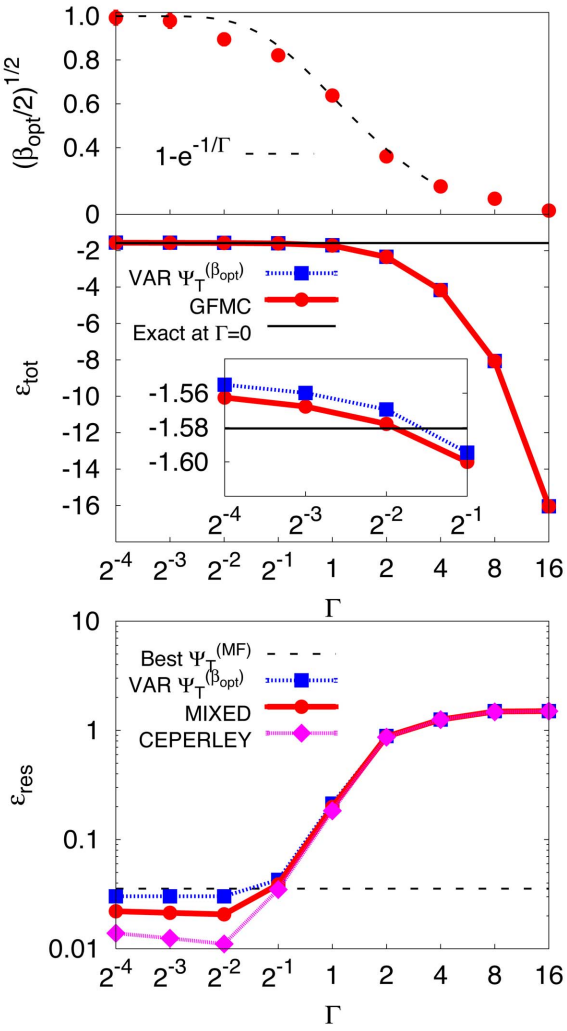

FIG. 1. (Color online) (Top) Plot of the optimal β, β_{opt}, for the "Boltzmann" trial wave function $\psi_T^{(\beta)}$ defined in Eq. (4), for several values of Γ. The dashed line is the fit employed in Sec. IV. (Center) Optimal variational energies $\epsilon_{tot}^{(Boltz)}$ corresponding to the β_{opt} shown in the top panel, and the GFMC estimate of the total energy per spin. The inset magnifies the small-Γ region, where little differences are noticeable. (Bottom) The variation residual diagonal energy $\epsilon_{res}^{(Boltz)}$ corresponding to the β_{opt} shown in the top panel, and two GFMC estimators of the residual diagonal energy: the mixed and the Ceperley one (see the text). The dashed horizontal line represents the best residual energy ever achieved, for $\Gamma>0.01$, by employing the mean-field trial wave function in Eq. (3).

As a last quality test for the Boltzmann-type trial wave function, we performed GFMC simulations *with importance sampling* in order to estimate the ground-state properties of the Hamiltonian in Eq. (1), again for several *fixed* values of Γ. Details can be found in Ref. [27].

The middle panel of Fig. 1 shows the estimate of the total energy (per spin) for some given value of Γ, compared to the correspondent variational results. The inset allows one to appreciate the differences between the two results in the small Γ region. As expected, GFMC allows one to gently improve the variational findings.

In the bottom panel we plot data regarding the residual diagonal energy. We report both the *mixed estimate* [17,31]—labeled as "mixed"—and the *Ceperley estimate* [31]—labeled as "Ceperley"—for this observable as obtained by GFMC simulations. Once again, they are consistently lower than the correspondent variational results, which are also displayed in the same panel. On the other hand, the nonmonotonic behavior of the residual energy data for small $\Gamma < 0.25$, already noted for the pure variational results, should ring a bell about the quality of the Boltzmann-type wave function, and the efficiency of the sampling in that region.

IV. GFMC QUANTUM ANNEALING

We finally present the results of the GFMC-based QA approach, where the transverse field Γ in Eq. (1) is decreased stepwise during the simulation, while at the same time, the importance sampling Boltzmann-type wave function is changed according to the corresponding value of the variational parameter $\beta_{opt}(\Gamma)$. [Practically, we used for $\beta(\Gamma)$ the fitting function shown in Fig. 1 (upper panel.)]

As a benchmark, we will compare GFMC-QA outcomes with the path-integral Monte Carlo quantum annealing (PIMC-QA) and classical simulated annealing (CA) results described in Refs. [7,8]. We reduce the coupling Γ in Eq. (1) at each Monte Carlo step (MCS) in a linear way: we start from an initial large enough value of the transverse field, $\Gamma_0 = 2.5$, and then we set $\Gamma_n = \Gamma_0(1 - n/\tau)$ during the nth MCS $(0 \leqslant n < \tau)$ [32]. τ is the total annealing time measured as the total number of MCS performed by the algorithm. We used 20 walkers, and performed branching at each MCS, because the low-Γ region is affected by severe weight instabilities, which would otherwise make the algorithm completely unstable (for the initial, large-Γ part of the annealing one could consider branching less often, as weights are well under control; this makes a negligible difference). Finally, we made use of a continuous-time approach, sampling directly the probability of generating an off-diagonal move according to a Poisson's process [31].

The fact that importance sampling is indeed a crucial ingredient is demonstrated, for our case, in Fig. 2, where we show the data obtained by GFMC-QA *without importance sampling* (top curve) compared with the ones obtained with importance sampling. Quite evidently, the residual energy obtained without importance sampling is terribly bad. Figure 2 also shows the best residual energy per spin ever reached during the annealing simulation, for several values of τ averaged over ten independent repetitions of the whole annealing process (due to computer-time limitations, only a single run is shown for the largest, $\tau > 10^8$, annealings). For comparison, the CA and PIMC-QA data obtained in Ref. [7] are also shown. Notice first that the τ axes of the three calculations are completely unrelated: the GFMC τ is measured in units in which a MCS is just a *single spin-flip*, while MCS for the CA and PIMC-QA are intended as *sweeps of the entire lattice* of N spins (including all the 20 Trotter slices, for the PIMC case [7]). For this reason, we also present the CA and PIMC-QA data in a shifted time axis where τ is multiplied by the number of sites (here $N = 80^2$, rightmost curves). Although the GFMC-QA data are strictly below both the CA and the PIMC-QA data, on the same per spin time unit (i.e., compared to the shifted CA and PIMC-QA data), it is clear that the GFMC slope is still *worse* than that of PIMC-QA, and indeed surprisingly similar to CA. Moreover, the CPU time needed for a single spin flip in GFMC is much larger than the corresponding single-spin move in CA or PIMC-QA (each GFMC move costs of order N operations).

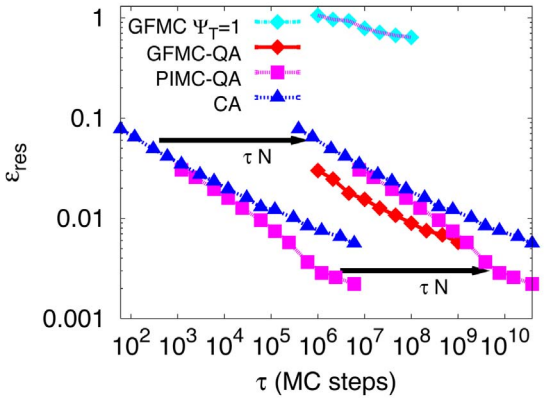

FIG. 2. (Color online) The average best residual energy obtained by GFMC-QA for the random Ising model instance studied in Refs. [7,8], versus the total annealing time τ. Upper rhombi: GFMC-QA results without importance sampling ($\psi_T = 1$). Lower rhombi: GFMC-QA results with importance sampling performed by using the optimal trial wave function $\psi_T^{(\beta)}$ of Sec. III. The GFMC time unit is a single spin flip, while CA and PIMC-QA Monte Carlo time units are sweeps of the entire lattice (see Ref. [8]). The transverse field is linearly reduced down to 10^{-4} in a total annealing time τ, starting from $\Gamma_0 = 2.5$. Previous results obtained by classical simulated annealing (CA) and by path-integral Monte Carlo quantum annealing (PIMC-QA) [7,8] are shown for comparison.

Let us consider the similarity between the CA and the GFMC-QA slopes that Fig. 2 suggests. This similarity must be somehow related to the fact that we have used, as importance wave function for the GFMC, a Boltzmann-type wave function, $\psi_T(x) \propto e^{-(\beta/2)E_{cl}(x)}$, as already observed in Sec. III [27]. More precisely, it is possible to show that by neglecting the weights of the walkers (as well as the associated branching process), GFMC reduces to a VMC sampling of the given trial function [33] (here the Boltzmann-type one). In other words, a GFMC-QA *without weights* would be just a computationally heavy way of doing a classical simulated annealing with a peculiar form of the temperature annealing

schedule $\beta_{opt}(\Gamma)$. (Notice, in passing, that such an optimal effective temperature never gets too low, since β_{opt} saturates to around $\beta_{opt} \approx 2$ for low Γ.)

Since genuine quantum mechanics enters only through the weights that the GFMC carries over, quite evidently, such a weight updating is—in the present disordered case—not sufficiently strong and effective as to make the resulting averages really different from the underlying thermal Markov chain.

V. SUMMARY AND CONCLUSIONS

In this paper we have investigated the practical feasibility of Green's function Monte Carlo (GFMC) as a tool for performing quantum annealing (QA). As a natural test case, we have concentrated our attention on a specific random instance of the two-dimensional Edwards-Anderson Ising model in transverse field, which was studied in Refs. [7,8] using PIMC-QA as well as standard thermal classical annealing (CA). (A more refined ensemble average would certainly be needed if we were to assess the general performance of the algorithm on a typical instance of the problem. A specific instance comparison with the competing algorithms (PIMC-QA and CA), however, turns out to be particularly instructive in the present case, as it shows in the clearest fashion the remarkable and unexpected similarity of our GFMC-QA results with CA.)

We identified the choice of the trial wave function (a necessary ingredient in any GFMC) as the crucial step—as well as the weak point—of a GFMC based QA (GFMC-QA). In particular, we found that the simplest mean-field wave function (analogous in many respects to the Weiss theory of ferromagnetism) is computationally equivalent to the original problem, and then completely useless. Using, instead, a simpler Boltzmann-type trial function (where the pseudotemperature is the only variational parameter), the resulting GFMC is feasible, but the corresponding residual energy results are disappointingly close—in magnitude and in slope, when considered as a function of the annealing time—to those found by a standard classical simulated annealing (which is computationally much cheaper). We can rationalize this finding with the inability of the GFMC algorithm, in the present disordered context, to properly implement the quantum evolution.

Concerning possible improvements in the implementation, we mention that we have not attempted a systematic study of alternative annealing schedules such as power laws. [21] Likewise, we have not experimented with other samplings of the Green's function [34] (for instance, Metropolis acceptance), or using alternative multi-spin-flip Green's functions [21]. Nevertheless, even our limited experimentation with a GFMC-QA algorithm shows that one of the crucial theoretical questions to be addressed is the ability to find a good trial variational wave function describing well enough the small-Γ glassy phase of an Ising spin glass. This is, quite evidently, a highly nontrivial task. Taking inspiration from the existing literature on quantum models without disorder, one might think of introducing pair correlations into the trial wave function—for instance, by means of spin-spin Jastrow factors, either at nearest neighbor or at longer range—as usually done in the framework of correlated lattice models [25], and of electronic structure calculations [17,33]. Unfortunately, for a quantum spin glass, due to frustration and disorder, the form of such pair correlations is far from obvious. Moreover, whenever a large number of variational parameters in the trial function is required, very advanced minimization techniques, such as those discussed in Ref. [35], are mandatory. These kinds of computational schemes, however, have been successfully tested only in equilibrium simulations of ordered systems, while our GFMC-QA should cope with a nonequilibrium dynamics in a disordered system, a highly nontrivial step forward.

We conclude that, at present, without a serious effort in constructing reliable importance sampling variational wave functions for a quantum glass, GFMC-QA is likely not a true competitor of PIMC-QA.

ACKNOWLEDGMENTS

This project was sponsored by MIUR through Grant Nos. FIRB RBAU017S8R004, No. FIRB RBAU01LX5H, PRIN/COFIN2003, and 2004, and by INFM ("Iniziativa trasversale calcolo parallelo"). We acknowledge discussions with Demian Battaglia, Michele Casula, Saverio Moroni, Sandro Sorella, Erio Tosatti, and Osvaldo Zagordi. G.E.S. also thanks H. Nishimori for sending him his recent preprint with results on several convergence theorems on QA.

[1] *Quantum Annealing and Related Optimization Methods*, edited by A. Das and B. K. Chakrabarti, Lecture Notes in Physics (Springer-Verlag, Berlin, 2005).
[2] G. E. Santoro and E. Tosatti, J. Phys. A **39**, R393 (2006).
[3] S. Kirkpatrick, C. D. Gelatt, Jr., and M. P. Vecchi, Science **220**, 671 (1983).
[4] E. Fahri, J. Goldstone, S. Gutmann, and M. Sipser, e-print quant-ph/0001106.
[5] Y. H. Lee and B. J. Berne, J. Phys. Chem. A **104**, 86 (2000).
[6] P. Liu and B. J. Berne, J. Chem. Phys. **118**, 2999 (2003).
[7] G. E. Santoro, R. Martoňák, E. Tosatti, and R. Car, Science **295**, 2427 (2002).
[8] R. Martoňák, G. E. Santoro, and E. Tosatti, Phys. Rev. B **66**, 094203 (2002).
[9] M. Sarjala, V. Petäjä, and M. Alava, J. Stat. Mech.: Theory Exp. (2006) P01008.
[10] Y. H. Lee and B. J. Berne, J. Phys. Chem. A **105**, 459 (2001).
[11] T. Gregor and R. Car, Chem. Phys. Lett. **412**, 125 (2005).
[12] R. Martoňák, G. E. Santoro, and E. Tosatti, Phys. Rev. E **70**, 057701 (2004).
[13] D. A. Battaglia, G. E. Santoro, and E. Tosatti, Phys. Rev. E **71**, 066707 (2005).

[14] C. H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity* (Dover, New York, 1998).
[15] M. R. Garey and D. S. Johnson, *Computers and Intractability* (Freeman, New York, 1979).
[16] L. Stella, G. E. Santoro, and E. Tosatti, Phys. Rev. B **73**, 144302 (2006).
[17] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, Rev. Mod. Phys.0034-6861 **73**, 33 (2001).
[18] F. Barahona, J. Phys. A **15**, 3241 (1982).
[19] C. De Simone, M. Diehl, M. Jünger, P. Mutzel, G. Reinelt, and G. Rinaldi, J. Stat. Phys. **80**, 487 (1995).
[20] A. Messiah, *Quantum Mechanics* (North-Holland, Amsterdam, 1962), Vol. 2.
[21] S. Morita and H. Nishimori, J. Phys. A **39**, 13903 (2006).
[22] J. Brooke, D. Bitko, T. F. Rosenbaum, and G. Aeppli, Science **284**, 779 (1999).
[23] J. Brooke, T. F. Rosenbaum, and G. Aeppli, Nature (London) **413**, 610 (2001).
[24] L. Stella, G. E. Santoro, and E. Tosatti, Phys. Rev. B **72**, 014303 (2005).
[25] M. Calandra Buonaura and S. Sorella, Phys. Rev. B **57**, 11446 (1998).
[26] M. H. Kalos and P. A. Whitlock, *Monte Carlo Methods* (Wiley, New York, 1986), Vol. 1.
[27] L. Stella and G. E. Santoro, e-print cond-mat/0608420.
[28] In particular, if ψ_T coincides with the actual ground-state wave function $\psi_{GS}(x)$, one can show that statistical fluctuations in the algorithm vanish exactly. This is the so-called zero variance property [17]. Note that there is no *sign problem* [17] for the Hamiltonian, Eq. (1).
[29] K. Binder and A. P. Young, Rev. Mod. Phys. **58**, 801 (1986).
[30] T. Castellani and A. Cavagna, J. Stat. Mech.: Theory Exp. (2005) P05012.
[31] N. Trivedi and D. M. Ceperley, Phys. Rev. B **41**, 4552 (1990).
[32] Actually, the final Γ is $\approx 10^{-4}$ and not 0 to prevent severe weight instabilities. Whenever it was possible to perform annealings with smaller cut-offs on Γ we checked that the results obtained are not very sensitive to its value.
[33] B. L. Hammond, W. A. Lester, and P. J. Reynolds, *Monte Carlo Methods in Ab Initio Quantum Chemistry* (World Scientific, Singapore, 1994).
[34] H. Nishimori and J.-I. Inoue, J. Phys. A **31**, 5661 (1998).
[35] S. Sorella, Phys. Rev. B **71**, 241103(R) (2005).