Quantum annealing by the path-integral Monte Carlo method: The two-dimensional random Ising model

Roman Martoňák,[1,*] Giuseppe E. Santoro,[2] and Erio Tosatti[2,3]
[1]*Swiss Center for Scientific Computing, Via Cantonale, CH-6928 Manno, Switzerland and ETH Zurich, Physical Chemistry, Hoenggerberg, CH-8093 Zurich, Switzerland*
[2]*International School for Advanced Studies (SISSA) and INFM (UdR SISSA), Trieste, Italy*
[3]*International Center for Theoretical Physics (ICTP), P.O. Box 586, Trieste, Italy*
(Received 22 March 2002; published 13 September 2002)

Quantum annealing was recently found experimentally in a disordered spin-$\frac{1}{2}$ magnet to be more effective than its classical, thermal counterpart. We use the random two-dimensional Ising model as a test example and perform on it both classical and quantum (path-integral) Monte Carlo annealing. A systematic study of the dependence of the final residual energy on the annealing Monte Carlo time quantitatively demonstrates the superiority of quantum relative to classical annealing in this system. In order to determine the parameter regime for optimal efficiency of the quantum annealing procedure we explore a range of values of Trotter slice number P and temperature T. This identifies two different regimes of freezing with respect to efficiency of the algorithm, and leads to useful guidelines for the optimal choice of quantum annealing parameters.

DOI: 10.1103/PhysRevB.66.094203 PACS number(s): 75.10.Nr, 02.70.Uu, 02.70.Ss, 07.05.Tp

I. INTRODUCTION

Optimization problems encountered in science and technology are not infrequently ''hard'' ones. Well-known hard problem prototypes include, e.g., the traveling salesman problem, finding the ground state of a spin glass, satisfiability problems, etc. For a class of such problems (NP complete problems) the time necessary to find the optimal solution grows faster than any polynomial of N, the size of the system, thus making an exact determination of the overall optimal state practically impossible, except for very small values of N. What therefore acquires particular importance is the possibility to optimize any such system approximately but efficiently. In the search for an approximate optimal state the method of stochastic optimization, or classical simulated annealing (CA),[1,2] is widely used in practice, along with many variants developed over the last two decades. In CA one introduces a real or fictitious temperature T so as to sample stochastically the phase space $\{x\}$ of the problem according to the Boltzmann factor $e^{-H(x)/T}$. Here $H(x)$ is the Hamiltonian of the problem (or more generally, the ''cost function'' associated with a configuration x). Once the thermal energy is sufficient to overcome even the largest barriers surrounding metastable local minima, the system escapes trapping in any specific minimum, and can thus sample the full phase space. By gradually reducing the temperature T during some long annealing time τ, the phase-space sampling concentrates more and more on regions of lower energy, until a final energy $E_{final}(\tau)$ is attained when T hits zero. Generally speaking, for a system which is large and complex, and for a finite annealing rate $1/\tau$, the sampling, unable to negotiate all barriers in the finite time τ, attains at the end of the CA protocol an average energy $E_{final}(\tau)$ strictly larger than the true ground-state (GS) energy E_{GS}. The average residual energy $\epsilon_{res}(\tau)=E_{final}(\tau)-E_{GS}$ is expected to be a slowly decreasing function of the annealing time τ, $\epsilon_{res}(\tau)\sim(\ln\tau)^{-\zeta}$.[3] In the case of a spin glass, general theoretical arguments by Huse and Fisher[3] predict an upper bound $\zeta\leqslant 2$, which implies an exceedingly slow convergence of CA.

Quantum mechanics provides an unexpected alternative to temperature in the annealing problem[4,5]. One introduces in the Hamiltonian a real or fictitious kinetic-energy term, possessing a nonzero commutator with the initial classical Hamiltonian, and characterized by a strength Γ. This strength is initially large, providing the necessary quantum fluctuations and broad wave functions $\Psi(x,t)$ necessary to explore the full phase space. Subsequently Γ is gradually reduced to zero, again in a time τ. Tunneling through barriers, rather than thermal hopping over the barriers, is the mechanism by which quantum mechanics avoids trapping in metastable local minima. As Γ is gradually reduced the system explores regions of increasingly lower potential energy, until, once again, a final energy $E_{final}(\tau)$ is attained. This alternative annealing scheme is known as quantum annealing (QA).[5]

Recently Brooke *et al.* provided an interesting experimental realization of QA on an Ising spin-glass system, specifically the spin-$\frac{1}{2}$ disordered Ising ferromagnet $LiHo_{0.44}Y_{0.56}F_4$.[6] Quantum fluctuations were introduced in the form of an external magnetic field lying in a plane orthogonal to the Ising easy axis, so that Γ is precisely the transverse field strength. Experimental comparison of the properties displayed by the system transported from the same initial state A—a paramagnetic high-T state—to the same nominal final state B—a low-T glassy state—through two different routes in the $[T,\Gamma]$ plane, presents evidence that QA, the ''quantum route'' from A to B, yields for the same ''cooling'' rate $1/\tau$, a better annealed final state B, detected through faster relaxation times in the dynamical susceptibility.[6]

This kind of observation raises a number of questions. First, fundamental, questions are why QA should, at least in this case, work better than CA, and more generally what is the underlying mechanism of QA. A second question, with a

potentially large impact in the field of heuristic optimization algorithms, is if, and how, it could be possible to implement QA on a computer to obtain a more efficient optimization algorithm for more general hard problems. We have recently addressed the first questions,[7] identifying the main element of real-time QA in the form of a cascade of Landau-Zener tunneling events connected with the (avoided) level crossings which the instantaneous ground state undergoes as Γ is reduced. For the spin-glass case, it was moreover suggested that the final energy $E_{final}(\tau)$ in QA should approach the exact ground-state energy E_{GS} in a logarithmically slow manner, $\epsilon_{res}(\tau)=E_{final}(\tau)-E_{GS}\sim(\ln\tau)^{-\zeta_{QA}}$, essentially similar to CA.[7] However, contrary to CA where $\zeta\leqslant 2$, QA may possess a generally larger exponent ζ_{QA} which makes it quantitatively better.[7]

The present paper is devoted to the second class of questions, dealing with the more technical aspects for the efficient implementation of QA as a numerical optimization algorithm, using the two-dimensional (2D) Ising spin glass as a benchmark example, and Monte Carlo (MC) as the technique of choice. In fact, although one could, in principle, conceive a real-time QA algorithm based on numerical integration of the time-dependent Schrödinger equation, such an approach is presently extremely demanding on computer resources and as such it is not feasible except for very small size systems.[8–10] For the larger systems of real interest it is mandatory to resort to stochastic methods, which include a variety of quantum Monte Carlo (QMC) techniques. An early study on the optimal configuration of small Lennard-Jones clusters—where the standard kinetic energy $-(\hbar^2/2M)\nabla^2$ with variable mass M plays the role of the quantum term—used diffusion Monte Carlo to propagate the initial wave function of the system in *imaginary time* (at zero temperature).[5] Another standard QMC technique, the path-integral Monte Carlo (PIMC), allows instead the simulation of a quantum system at a finite temperature. Working at a sufficiently low but finite temperature T, and slowly decreasing the strength of the quantum kinetic term, it is possible to use PIMC as a QA method to approach the ground state.[7,10] Presently, a variant of this method is also being applied to Lennard-Jones clusters.[11] An alternative quantum Monte Carlo scheme, essentially a ground-state path integral, is discussed in Ref. 10, where applications to small 16-city instances of the traveling salesman problem are also provided.

In PIMC, besides the temperature T, one must also specify the imaginary-time slice number (Trotter number) P, counting the coupled classical replicas of the physical system used to represent the quantum system. The optimal choice of T and P for QA is not *a priori* obvious. In some path-integral QA (PIQA) studies conducted so far (applied to a small protein folding problem[12,13]) both parameters were changed during the annealing. In Ref. 12 the strategy adopted was to combine QA with CA by starting at a high temperature and gradually decreasing T to zero, while performing, at each temperature step, a PIMC QA. The Trotter number P was also decreased as T was lowered so that at $T=0$ a classical system with $P=1$ was recovered. In Ref. 13 a Migdal-Kadanoff renormalization scheme was used, consisting of a progressive decimation of Trotter slices (i.e., a decrease of P) while reducing at the same time the quantum fluctuation parameter Γ.

Although many possible variations or hybrid algorithms combining thermal and quantum annealing are conceivable, it appears that a detailed analysis of the ingredients entering a plain PIQA, with constant Trotter number P and constant low temperature T, would be at this point very desirable, especially in the design and control of more sophisticated techniques. The present work aims at filling this gap.

In Ref. 7 we compared the efficiency of such a fixed P,T PIQA with that of standard CA by studying the dependence of the final residual energy $\epsilon_{res}(\tau)$ on the annealing time τ for $L\times L$ 2D samples of an Ising spin glass up to $L=80$. In this work, using the same test system, we explore PIQA in more detail, by studying systematically the annealing performance for different values of P and T. As a result of this analysis we derive useful guidelines for a suitable practical implementation of the PIQA procedure, of hopefully more general validity.

The paper is organized as follows. In Sec. II we describe the model used as our test case, and sketch its path-integral representation. In Sec. III we describe in some detail the MC implementation of the PIQA technique. Section IV is devoted to the discussion of the results of PIQA and CA simulations. Finally, in Sec. V we draw our conclusions and suggest possible directions for future improvements. Technical details about the path-integral representation are included in the Appendix for the reader's convenience.

II. THE MODEL AND THE PATH-INTEGRAL REPRESENTATION OF THE PARTITION FUNCTION

Our test complex system will be the 2D random Ising model, a choice dictated by several reasons. First we were directly inspired by Brooke *et al.*'s experimental system, also a disordered Ising magnet.[6] A second and main reason is that, although technically a polynomial problem[14] and not a spin glass at any $T>0$, the 2D random Ising model is nonetheless of prohibitively large complexity, with a large continuum of metastable minima above the ground state as in a true glass.[15] A final reason is that the exact classical ground state energy of this model is numerically accessible, for an arbitrary realization of random couplings, via the spin-glass server[16] up to sufficiently large lattice sizes $\sim 100\times 100$, permitting an absolutely precise measure of the residual energy after every annealing protocol, which in turn provides the superior accuracy needed to study the asymptotic behavior.

The Edwards-Anderson model Hamiltonian of the Ising spin glass is

$$H_{EA}=-\sum_{\langle ij\rangle} J_{ij}s_i s_j\,. \tag{1}$$

Here Ising spins ($s_i=\pm 1$) occupy the sites of a d-dimensional cubic lattice, and J_{ij} are the random couplings between nearest-neighbor sites drawn from some prescribed distribution. When the couplings J_{ij} fluctuate randomly without a definite sign, the Hamiltonian (1) describes a frustrated and disordered system. The task of finding its ground state is

a hard computational problem. It is known in three dimensions to be *NP* complete:[14,17] the ground state cannot be systematically attained with an algorithm whose time increases as any finite power with the size of the system. While in two dimensions that is strictly speaking not the case—there is also no true spin-glass phase at finite T (Ref. 15)—the annealing problem is in practice still prohibitive, thus providing a suitable testing ground.

Adding a magnetic field orthogonal to the Ising axis one obtains the Hamiltonian of the Ising spin glass in transverse field,

$$H=-\sum_{\langle ij\rangle} J_{ij}\sigma_i^z\sigma_j^z-\Gamma\sum_i \sigma_i^x, \tag{2}$$

where σ_i^x,σ_i^z are Pauli matrices corresponding to the spin on lattice site i. Model (2) is directly inspired by the experimental system of Brooke *et al.*[6] Physically, the transverse field Γ represents a kinetic energy which does not commute with the classical Ising term, inducing transitions between the $\uparrow$ and $\downarrow$ states of each single spin, thus turning the model from classical to quantum.

In order to derive a path-integral representation for the quantum Ising spin-glass model (2) we apply to its canonical partition function the standard path-integral technique.[18] We write

$$H=U+K,$$

$$U=-\sum_{\langle ij\rangle} J_{ij}\sigma_i^z\sigma_j^z,\quad K=-\Gamma\sum_i \sigma_i^x,$$

where, as previously mentioned, the terms U (potential energy) and K (kinetic energy) do not commute, $[K,U]\neq 0$. Defining, as usual, $\beta=1/k_BT$, where k_B is the Boltzmann constant which we set henceforth to 1, the partition function Z at a temperature T can be written as

$$\begin{aligned} Z&=\mathrm{Tr}\, e^{-\beta H}\\ &=\mathrm{Tr}(e^{-\beta H/P})^P\\ &=\mathrm{Tr}(e^{-\beta(K+U)/P})^P=\sum_{s^1}\ \ldots\sum_{s^P}\langle s^1|e^{-\beta(K+U)/P}|s^2\rangle\\ &\quad\times\langle s^2|e^{-\beta(K+U)/P}\ldots|s^P\rangle\langle s^P|e^{-\beta(K+U)/P}|s^1\rangle. \end{aligned} \tag{3}$$

Here $s^k=\{s_i^k\}$ denotes a configuration of all the spins in the kth Trotter slice, and the last equality follows from insertions of the identity operator $\mathbf{1}=\Sigma_{s^k}|s^k\rangle\langle s^k|$. Thus far everything is exact. Now we apply the Trotter breakup formula $e^{-\beta(K+U)}\approx e^{-\beta K}e^{-\beta U}$, which neglects commutators of K and U,[18] obtaining an approximation Z_P to Z whose error is proportional to the square of the Trotter breakup time $\Delta t=\beta/P$. The final expression, after simple algebra which we reproduce in the Appendix for the reader's convenience, is

$$Z\approx Z_P=C^{NP}\sum_{s^1}\ \ldots\sum_{s^P}\ e^{-H_{d+1}/PT}, \tag{4}$$

$$H_{d+1}=-\sum_{k=1}^{P}\left(\sum_{\langle ij\rangle} J_{ij}s_i^ks_j^k+J^\perp\sum_i s_i^ks_i^{k+1}\right), \tag{5}$$

where N is the number of lattice sites in the d-dimensional lattice, and

$$J^\perp=-\frac{PT}{2}\ln\tanh\frac{\Gamma}{PT}>0, \tag{6}$$

$$C=\left(\frac{1}{2}\sinh\frac{2\Gamma}{PT}\right)^{1/2}.$$

Z_P is the partition function of a classical $(d+1)$-dimensional *anisotropic* Ising system at temperature PT, with couplings J_{ij} along the original d-dimensional lattice bonds (same for all Trotter slices), and $J^\perp$ (ferromagnetic and uniform, i.e., the same for all sites i) along the extra dimension, where the system has a finite length P. It follows from the properties of the trace that periodic boundary conditions have to be taken along this dimension, i.e., $s^{P+1}=s^1$. This well-known mapping[18] of the d-dimensional quantum Ising system onto a $(d+1)$-dimensional classical Ising system allows simulation of the quantum system by a standard classical Monte Carlo sampling.

III. MONTE CARLO IMPLEMENTATION

Since we are interested in annealing, that is, approaching as closely as possible the classical ground state, rather than reproducing the exact quantum averages, we do not necessarily have to work in the large Trotter number P limit, or perform extrapolations to $P\to\infty$. For the same reason, we could, in principle, relax the boundary conditions along the imaginary-time direction, required to be periodic by the trace condition. In the present benchmark study, however, we deliberately retained the standard PIMC boundary conditions which are periodic along the imaginary-time direction, as well as along the spatial directions.

The results presented in this paper concern a single realization of the 2D Edwards-Anderson model with random couplings on a 80×80 sample with periodic boundary conditions. The random couplings were distributed according to a uniform distribution on the interval $(-2;2)$. We note that we performed similar simulations with several realizations of random couplings on systems 32×32 and 48×48 and obtained in all cases very similar results. Thus, the results for the 80×80 system can be regarded as typical.

We start by discussing the initialization of the PIQA procedure. It follows from the mapping that at the beginning of PIQA, when the transverse field Γ is large, the coupling $J^\perp$ (6) between the neighboring slices is very small and therefore the $(2+1)D$ system behaves like a collection of noninteracting 2D systems at temperature PT. An appropriate initial configuration is thus to set all Trotter slices equal to one another, with a spin configuration corresponding to the equilibrium Boltzmann distribution at temperature PT. This choice does, however, assume the ability to achieve a thermal equilibrium distribution with temperature PT, which is possible in practice only if PT is not too small, i.e., larger

than the corresponding glass transition temperature T_{glass}. Specifically, one should guarantee that $PT \gtrsim |J|$, where $|J|$ is a typical scale of the couplings. Since our couplings have an average magnitude of $|J| \sim 1$, our PT should not be smaller than 1; we used here values of PT ranging from 1 to 2 ($PT = 1, 1.5, 2$). In all these cases we prepared the initial configuration by performing a classical preannealing that consisted of (i) starting from a completely random spin configuration at a sufficiently large initial temperature of $T_0 = 3$, and (ii) decreasing the temperature in steps of $\Delta T = 0.05$ down to the desired value of PT performing 100 MC steps per spin at each temperature step. This annealing rate turned out to be sufficient for our case.

Next, we need to choose an initial value Γ_0 of the transverse field. In CA the choice of the initial temperature is not critical provided the value is high enough for the system to be able to reach thermal equilibrium. In PIQA, on the other hand, if Γ is too large, the Trotter slices are practically decoupled. Although initialized to the same configuration, after evolution in the decoupled regime for too long they will become uncorrelated, which is bad in view of the subsequent introduction of interslice correlations upon decreasing Γ. Conversely, an excessively small initial Γ_0 value can also have an adverse effect since now the slices do not have enough time to build proper correlations between them. For this reason a moderate initial value of Γ_0 is best. The optimal value of Γ_0 was determined by testing several values for each P and T using a short annealing time τ. For $PT = 1$ the best value turned out to be about 3.0 while for $PT = 1.5$ and $PT = 2$ the value of 2.5 was optimal; Γ_0 found in this way was then used for all values of τ.

We sampled the system by a standard Metropolis algorithm employing both local and global moves as is common in PIMC simulations. The local moves attempt independent spin flips at all sites in all Trotter slices. In a global move, instead, an attempt is made to flip simultaneously all the replicas of the same spin in all Trotter slices. Clearly, the acceptance ratio of the global moves does not depend on the transverse field Γ since both replicas of the same spin in neighboring Trotter slices which interact via $J^{\perp}$ (which depends on Γ) are always flipped at the same time.

The QA annealing protocol consisted of a linear decrease of the transverse field Γ from the initial value of Γ_0 to the final value (close to zero but still finite) of $\Gamma = 10^{-8}$. The temperature T and Trotter number P were kept constant during the annealing. At each value of Γ one MC step per spin (MCS) is performed. Each MCS consists here of a local move followed by a global move and the total number of MCS's performed during the annealing is τ. From this definition it follows that the CPU time required for 1 MCS scales linearly with the Trotter number P. At the end of the quantum annealing we proceeded to identify the Trotter slice with minimum energy and took that energy as E_{final}. The search was always repeated $N = 45$ times and the final energy was averaged over all searches. We note that each search was initialized from a different classically preannealed configuration.

To end this section we briefly mention also the CA annealing protocol used. We started from a spin configuration with randomly chosen spins and the temperature T was linearly decreased from the initial value of $T_0 = 3.0$ to zero, performing at each temperature step one MCS resulting in a total of τ MCS's. Also for CA, we always averaged the final energies over $N = 45$ searches, each search starting from a different random configuration.

IV. RESULTS AND DISCUSSION

We first characterize the efficiency of both the PIQA and CA methods, by studying the cooling rate dependence of $\epsilon_{res}(\tau)$ upon varying the annealing time τ over as many decades as possible.[19] Besides a quantitative comparison between CA and PIQA, which is not our main scope here, this will lead to a discussion of the behavior of PIQA Monte Carlo data as a function of the simulation parameters P and T and provide an interpretation in terms of different regimes of freezing during the PIQA procedure. We stress that, for the large system studied here, it is impossible to characterize the efficiency of an algorithm through the ''probability of getting the actual GS'' during the simulation, $P_{GS}(\tau)$, as done in Refs. 8–10, since $P_{GS}(\tau)$ is vanishingly small.

A. Dependence of quantum annealing on P and T

As previously mentioned, the product PT determines the couplings both between the spin replicas in neighboring Trotter slices (for a given value of Γ) and between the spins within slices [Eqs. (4)–(6)]. For a given value of PT, the Trotter number P itself only determines the size of the lattice along the imaginary-time direction. Therefore it seems natural to study the PIQA efficiency by varying independently PT and P. As mentioned in the previous section we used here three values of $PT = 1$, 1.5, and 2. For $PT = 1.5$ and $PT = 2$ we used the values of $P = 30$ and $P = 40$, respectively, corresponding to $T = 0.05$. For the smallest value $PT = 1$ we performed a more detailed study using several values of $P = 5$, 10, 20, 30, 40, and 50. The values of the total annealing time range from $\tau = 60$ MCS's to $\tau = 300\,000$ MCS's.

The results for the residual energy in QA are shown in a log-log plot in Fig. 1. There are noticeable differences between curves corresponding to different values of PT. The $PT = 2$ curve, obtained for $T = 0.05$ with $P = 40$, starts for low τ at the highest value of ϵ_{res} and initially decreases, on increasing τ, at the fastest rate until $\tau \sim 1000$, where it starts to develop a plateau. Further increase of τ results only in a very slow decrease of $\epsilon_{res}(\tau)$ indicating that PIQA in this regime is not efficient. The $PT = 1.5$ curve, corresponding to the same $T = 0.05$ but now for $P = 30$, looks similar in shape but develops the plateau around a slightly larger value of $\tau \sim 2000$; it also reaches a slightly lower value of ϵ_{res}. For $PT = 1$ and P increasing from 5 to 40 we have a series of curves which also have shapes similar to the previous ones. However, in particular the $P = 20$ and $P = 40$ curves now reach for large τ considerably lower values of ϵ_{res}. Moreover, a comparison of curves with increasing P reveals here interesting convergence properties. The remarkable feature is that the curves corresponding to different values of P, at

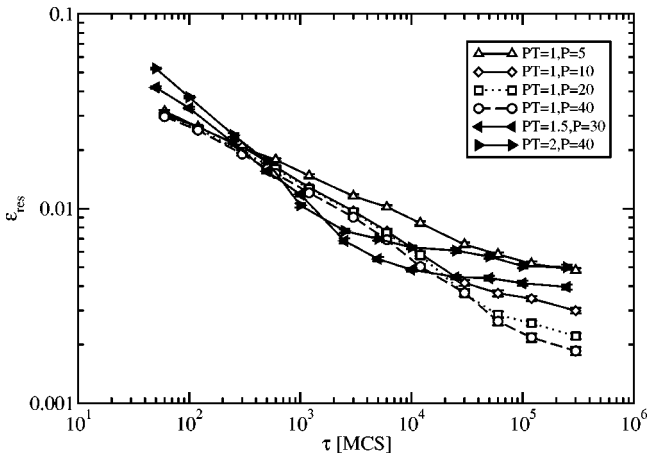

FIG. 1. Residual energy per site obtained by PIQA for an 80 $\times$80 disordered 2D Ising model as a function of total annealing time τ for different values of parameter PT and Trotter number P.

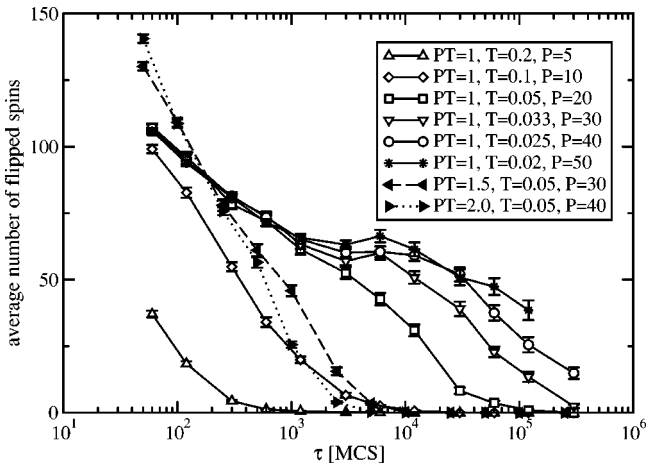

FIG. 2. Average number of flipped spins between neighboring Trotter slices in the final configurations at the end of PIQA as a function of total annealing time τ for different values of parameter PT and Trotter number P. The same symbols as in Fig. 1 have been used for the same runs to make comparison easier.

fixed $PT=1$, coincide up to a certain value of τ. However, for each value of P there is a characteristic value of τ where the corresponding curve splits from the bunch and starts to decrease in a much slower fashion for increasing τ, leveling off to a plateau. This characteristic τ increases with increasing P. While the $P=5$ curve splits already at $\tau\sim 500$, the $P=10$ curve does so at $\tau\sim 10\,000$ and the $P=20$ curve splits only at $\tau\sim 60\,000$. Overall, the $P=40$ curve provides the best results.

These findings can be understood as follows. The PIQA can actually be seen as cooling a classical system from an initial temperature PT to a final low temperature T. However, unlike CA, where the temperature is straightforwardly lowered, in PIQA P decoupled replicas of the system are initially created at a higher temperature PT and subsequently gradually forced into the same configuration by increasing the coupling $J^{\perp}$ as the transverse field Γ is reduced to zero. Doing that slowly enough would indeed lead to a classical equilibrium configuration of the system at a low temperature T; given not enough time, the system will instead freeze in the process. There are two essentially different ways the $(d+1)$-dimensional path-integral system can freeze depending on whether at the end of the quantum annealing (when the couplings between the Trotter slices become very large) all Trotter slices are in the same configuration. The case when this is true can be called ''classical freezing'' and the opposite one perhaps ''quantum freezing.''

In order to check which of the two scenarios applies we calculated the average number of flipped spins between neighboring Trotter slices (transverse field term) in the final configurations at the end of PIQA. These results are shown in Fig. 2 and provide a clue to understanding the data of Fig. 1. Comparison of the two figures shows that a curve splits from the bunch when the number of flipped spins at the end of annealing goes to zero; in such a case the system undergoes a classical freezing at the end of quantum annealing, or even before. The reason for this is that the size P of the lattice along the Trotter dimension is too short for the given PT (which determines couplings among the slices) and τ. Quantum fluctuations thus disappear before or when Γ becomes zero and, not surprisingly, this is not an efficient regime of PIQA. In the alternative case of quantum freezing different Trotter slices explore and tunnel between different energy valleys until the very end of the annealing, which makes the annealing more efficient. With larger values of P and/or a smaller coupling between the slices (smaller PT) the annealing time τ required to correlate completely all Trotter slices, thus reaching the classical freezing regime, grows. This explains the main features of data shown in Fig. 1 and Fig. 2.

B. Comparison between CA and QA

For comparison, we performed also a CA cooling rate study on the same sample[7] for a broad range of τ ranging from 60 to 6×10^6. In Fig. 3 we show the CA results together with the PIQA results. For a fair comparison of different methods we multiplied here the τ values in PIQA by the corresponding Trotter number P, since performing one MCS in PIQA requires updating P replicas of the system, and costs therefore P times more computer time than the corresponding classical sweep of the lattice. The results show that PIQA performs considerably better than CA, reaching, for the largest τ studied, values of ϵ_{res} smaller by a factor of about 3 compared to CA.

To perform a further quantitative comparison between PIQA and CA in terms of CPU time required to reach a given value of ϵ_{res} we can make use of the Huse-Fisher theoretical prediction of a logarithmic dependence $\epsilon_{res}(\tau) = A(\ln \gamma\tau)^{-\zeta}$, where $\zeta\leqslant 2$.[3] In Ref. 7 we showed that the CA data are, for large τ, compatible with this prediction although it is not possible to determine the exact value of the exponent ζ. Nevertheless, we can make an estimate assuming for CA the fastest theoretically allowed decay corresponding to $\zeta=2$.[3] Fitting the CA data in the asymptotic regime with such a form we obtained the values $A=0.9988, \gamma=0.0875$. The lowest value of ϵ_{res} reached by PIQA was $\epsilon_{res}=1.857\times 10^{-3}$, obtained for $PT=1$, $P=40$,

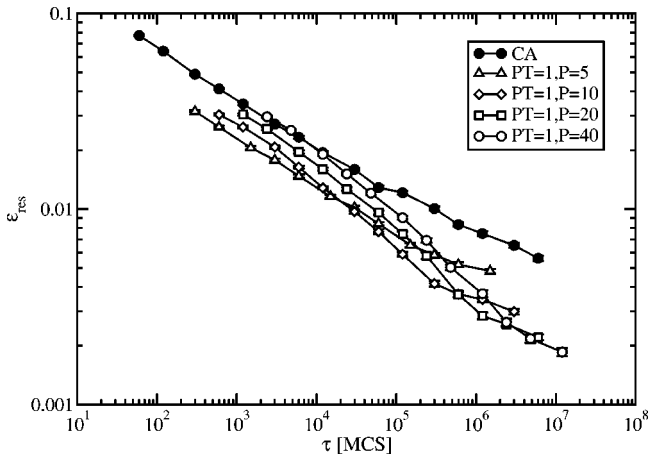

FIG. 3. Comparison of the residual energy per site for an 80 $\times$80 disordered 2D Ising model after CA and PIQA. The PIQA data are shown for $PT=1$ and different values of the Trotter number P. For fair comparison, the actual total annealing time τ used in the PIQA has been multiplied by P so that points at the same τ require roughly the same computer time.

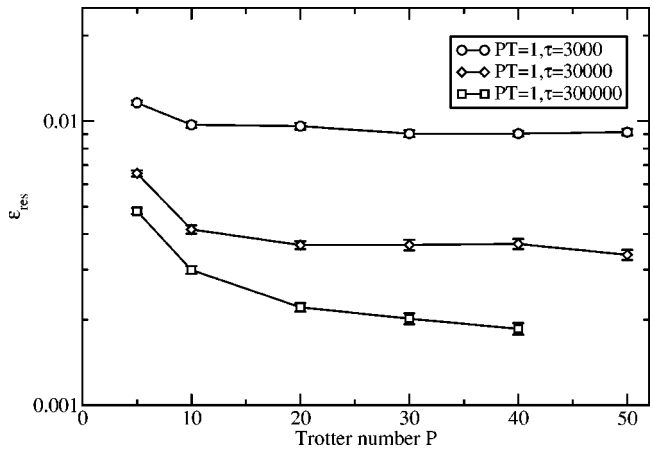

FIG. 4. Residual energy per site obtained by PIQA for an 80 $\times$80 disordered 2D Ising model as a function of the Trotter number P for different values of the total annealing time τ. Lines are just guides for the eye.

and $\tau=300\,000$. According to the fitted logarithmic law, this value would be reached by CA for $\tau=1.35\times10^{11}$ MCS's. In the PIQA simulation we effectively needed a CPU time equivalent to $P\tau=40\times300\,000=1.2\times10^{7}$ MCS's. Such a gain in CPU time with respect to CA of about four orders of magnitude represents a dramatic difference in efficiency: one can calculate using PIQA in one day what would be obtained by plain CA in about 30 yr.

V. GUIDELINES, CONCLUSIONS AND OUTLOOK

A. Guidelines

Finally we can sum up our analysis of path-integral quantum annealing to yield simple guidelines for the choice of parameters P and T. In order to easily prepare the initial configuration for the Trotter slices, one must be able to equilibrate the system at temperature PT. This places a lower bound on PT requiring it to be at least comparable with the couplings J in the system. A reasonable choice thus is to take $PT\sim J$. At fixed PT, QA works progressively better upon increasing P. However, for fixed total annealing time τ the residual energy saturates to a limiting value for large P, as shown in Fig. 4. Therefore for each value of τ there is a value of P beyond which further increasing of P is useless, and merely consumes computer time (Fig. 4). Finally, in order to achieve a better annealing, a longer annealing time τ is needed, but that in turn requires a larger value of P for the best possible ϵ_{res}.

B. Conclusions and outlook

In this paper we have explored the applicability of the plain path-integral Monte Carlo technique with constant Trotter number P and constant low temperature T to quantum annealing using a 2D random Ising model as a test example. We discussed issues related to the efficient implementation of the PIQA procedure and performed a detailed study of the cooling rate dependence of $\epsilon_{res}(\tau)$ for PIQA for different values of P and T. This study enabled us to identify two possible regimes of freezing in PIQA, and to clarify the relationship between the parameters P and T, the freezing regime and efficiency of PIQA. For the system studied, PIQA appears to be a more efficient method than CA since it approaches the ground state considerably faster.

Quite clearly, our path-integral Monte Carlo, as well as the corresponding classical Monte Carlo, was based on the most simple local moves suggested by the problem: single spin flips. It is worth noting that much more efficient global moves for general spin systems have been introduced, in the last decade, through the so-called *loop algorithm*.[20,21] These global moves provide the natural generalization to quantum spin systems of the cluster moves originally introduced by Swendsen and Wang for classical Ising systems.[22] A successful application to the two-dimensional random Ising ferromagnet in a transverse field is reported in Ref. 23, where the interesting possibility of performing a continuous imaginary-time sampling, i.e., without Trotter discretization error, is emphasized. Such approaches are certainly worth pursuing in the attempt to improve on our basic scheme. We mention, however, the following two caveats: (i) the application of the loop algorithm or similar cluster moves to a case with genuine frustration, like a spin glass, is not completely guaranteed to be successful, as the use of Swendsen-Wang moves in classical spin glasses shows; (ii) similar smart global moves are generally not available for a generic optimization problem. In this respect, the simple minded ''single spin flip'' type of moves we have used is, often, all one can afford.

In the future, we plan to apply the methodology we have presented to a 3D spin-glass case where the problem is NP hard. It would also be interesting to elaborate on the optimal annealing schedule; it is possible that decreasing transverse field Γ in a nonlinear way could produce even better results. Our results suggest that PIQA could indeed be a promising optimization technique and it seems worthwhile to continue to study its applicability also to other kinds of problems of physical and general interest. To do so, however, it is neces-

sary to construct for each problem under consideration a suitable kinetic-energy operator and an appropriate Trotter decoupling, which might not be completely trivial. An extra advantage of PIQA is the fact that it can be easily implemented on a parallel computer, each Trotter slice running on one node, providing a parallel optimization algorithm requiring little communication between the nodes. It is also interesting to explore possible connections of the PIQA method to parallel tempering,[24] a project which we are currently pursuing.

ACKNOWLEDGMENTS

This project was sponsored by MIUR Grant No. COFIN2001, and by INFM/F, INFM/G, and INFM Iniziativa Trasversale di Calcolo Parallelo. R.M. would like to acknowledge support from the MINOS project of CINECA, which also provided computer time, as well as the hospitality provided by SISSA and ICTP Trieste during collaborative visits. We are grateful to R. Car, A. Maritan, and R. Zecchina for helpful discussions.

APPENDIX: DETAILS OF THE PATH-INTEGRAL DERIVATION

We include here, for the reader's convenience, a few technical details on the path-integral representation[18] of the Ising spin glass used to perform quantum annealing (QA).

Starting from Eq. (3) one applies the Trotter breakup formula $e^{-(K+U)/PT} \approx e^{-K/PT} e^{-U/PT}$ which neglects commutators of K and U,[18] obtaining

$$Z \approx Z_P = \sum_{s^1} \cdots \sum_{s^P} \langle s^1 | e^{-\beta K/P} e^{-\beta U/P} | s^2 \rangle \times \langle s^2 | \ldots e^{-\beta K/P} e^{-\beta U/P} | s^P \rangle \langle s^P | e^{-\beta K/P} e^{-\beta U/P} | s^1 \rangle,$$

with an error proportional to the square of the Trotter breakup time, $O((\beta/P)^2)$.[18] We need to evaluate the expression $\langle s^k | e^{-\beta K/P} e^{-\beta U/P} | s^{k+1} \rangle$, which is simply reexpressed as

$$\langle s^k | e^{-\beta K/P} e^{-\beta U/P} | s^{k+1} \rangle = \langle s^k | e^{-\beta K/P} | s^{k+1} \rangle e^{-\beta U(s^{k+1})/P} \tag{A1}$$

since the potential energy U is diagonal in the chosen spin basis. The only nontrivial term is therefore the average of the kinetic term between two Trotter slices, $\langle s^k | e^{-\beta K/P} | s^{k+1} \rangle$. Since spin operators corresponding to different sites commute,[25] we can rewrite such a term as

$$\langle s^k | e^{-\beta K/P} | s^{k+1} \rangle = \langle s^k | \exp\left(\frac{\beta\Gamma}{P} \sum_{i=1}^{N} \sigma_i^x \right) | s^{k+1} \rangle = \prod_{i=1}^{N} \langle s_i^k | \exp\left(\frac{\beta\Gamma}{P} \sigma_i^x \right) | s_i^{k+1} \rangle,$$

where N is the number of lattice sites. From simple spin-$\frac{1}{2}$ algebra,[25] it is easy to show that

$$\langle \uparrow | e^{a\sigma_x} | \uparrow \rangle = \langle \downarrow | e^{a\sigma_x} | \downarrow \rangle = \cosh a,$$

$$\langle \uparrow | e^{a\sigma_x} | \downarrow \rangle = \langle \downarrow | e^{a\sigma_x} | \uparrow \rangle = \sinh a,$$

which can be written as an Ising-like interaction (s, s' now mean single spins)

$$\langle s | e^{a\sigma_x} | s' \rangle = C e^{B s s'}$$

with $B = -\frac{1}{2} \ln \tanh a$, and $C^2 = \frac{1}{2} \sinh 2a$. Collecting all pieces together, we get

$$\langle s^k | e^{-\beta K/P} e^{-\beta U/P} | s^{k+1} \rangle = C^N e^{(J^\perp/PT) \Sigma_i s_i^k s_i^{k+1}} e^{(1/PT) \Sigma_{\langle ij \rangle} J_{ij} s_i^k s_j^k},$$

where

$$J^\perp = -\frac{PT}{2} \ln \tanh \frac{\Gamma}{PT} > 0, \tag{A2}$$

$$C^2 = \frac{1}{2} \sinh \frac{2\Gamma}{PT}.$$

The $J^\perp$ term can be seen as a ferromagnetic Ising-like coupling between the Trotter replicas of the same spin which are nearest neighbors (k and $k+1$) along the Trotter dimension.

For the full partition function we thus finally get

$$Z \approx Z_P = C^{NP} \sum_{s^1} \cdots \sum_{s^P} e^{-H_{d+1}/PT}, \tag{A3}$$

$$H_{d+1} = -\sum_{k=1}^{P} \left(\sum_{\langle ij \rangle} J_{ij} s_i^k s_j^k + J^\perp \sum_i s_i^k s_i^{k+1} \right), \tag{A4}$$

which represents the partition function of a classical ($d+1$)-dimensional anisotropic Ising system at temperature PT. The system has couplings J_{ij} along the original d-dimensional lattice bonds (same for all Trotter slices), and $J^\perp$ (same for all sites i) along the extra Trotter dimension, where the system has a finite length P.

*Permanent address: Department of Physics, Faculty of Electrical Engineering and Information Technology, Slovak University of Technology, Ilkovičova 3, 812 19 Bratislava, Slovakia.

[1] S. Kirkpatrick, C. D. Gelatt, Jr., and M. P. Vecchi, Science **220**, 671 (1983).
[2] V. Černý, J. Optim. Theory Appl. **45**, 41 (1985).
[3] D. A. Huse and D. S. Fisher, Phys. Rev. Lett. **57**, 2203 (1986).
[4] P. Amara, D. Hsu, and J. E. Straub, J. Phys. Chem. **97**, 6715 (1993).
[5] A. B. Finnila, M. A. Gomez, C. Sebenik, C. Stenson, and J. D. Doll, Chem. Phys. Lett. **219**, 343 (1994).
[6] J. Brooke, D. Bitko, T. F. Rosenbaum, and G. Aeppli, Science **284**, 779 (1999).
[7] G. E. Santoro, R. Martoňák, E. Tosatti, and R. Car, Science **295**, 2427 (2002).
[8] T. Kadowaki and H. Nishimori, Phys. Rev. E **58**, 5355 (1998).
[9] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, Science **292**, 472 (2001).

[10] T. Kadowaki, Ph.D. thesis, Tokyo Institute of Technology, 1998; quant-ph/0205020 (unpublished).
[11] R. Car (private communication).
[12] Y. Lee and B. J. Berne, J. Phys. Chem. A **104**, 86 (2000).
[13] Y. Lee and B. J. Berne, Ann. Phys. (Leipzig) **9**, 668 (2000); J. Phys. Chem. A **105**, 459 (2001).
[14] F. Barahona, J. Phys. A **15**, 3241 (1982).
[15] D. S. Fisher and D. A. Huse, Phys. Rev. Lett. **56**, 1601 (1986).
[16] Spin glass server, http://www.informatik.uni-koeln.de/ls_juenger/projects/sgs.html
[17] J. C. Angles d'Auriac, M. Preissmann, and R. Rammal, J. Phys. (France) Lett. **46**, L-173 (1985).
[18] M. Suzuki, Prog. Theor. Phys. **56**, 1454 (1976); See also *Quantum Monte Carlo Methods in Equilibrium and Nonequilibrium Systems*, Proceedings of the Ninth Taniguchi International Symposium, Susono, Japan, 1986, edited by M. Suzuki (Springer-Verlag, Berlin, 1987).
[19] G. S. Grest, C. M. Soukoulis, and K. Levin, Phys. Rev. Lett. **56**, 1148 (1986).
[20] H. G. Evertz, G. Lana, and M. Marcu, Phys. Rev. Lett. **70**, 875 (1993). For a review of the loop algorithm see, for instance, H.G. Evertz, cond-mat/9707221 (unpublished).
[21] B. Ammon, H. G. Evertz, N. Kawashima, M. Troyer, and B. Frischmuth, Phys. Rev. B **58**, 4304 (1998).
[22] R. H. Swendsen and J.-S. Wang, Phys. Rev. Lett. **58**, 86 (1987).
[23] H. Rieger and N. Kawashima, Eur. Phys. J. B **9**, 233 (1999). For a different application of continuous time sampling to polaron problems, see P. E. Kornilovitch, Phys. Rev. Lett. **81**, 5382 (1998).
[24] E. Marinari and G. Parisi, Europhys. Lett. **19**, 451 (1992); G. J. Geyer, Stat. Sci. **7**, 437 (1992); K. Hukushima and K. Nemoto, J. Phys. Soc. Jpn. **65**, 1604 (1996); M. C. Tesi, E. J. J. van Rensburg, E. Orlandini, and S. G. Whittington, J. Stat. Phys. **82**, 155 (1996); U. H. E. Hansmann, Chem. Phys. Lett. **281**, 140 (1997).
[25] L.D. Landau and E.M. Lifshitz, *Quantum Mechanics—Non-relativistic Theory* (Pergamon, Oxford, 1977).