

Global Optimization: Quantum Thermal Annealing with Path Integral Monte Carlo

Yong-Han Lee[†] and B. J. Berne*

Department of Chemistry and Center for Biomolecular Simulation, Columbia University, 3000 Broadway, New York, New York 10027

Received: June 8, 1999; In Final Form: October 8, 1999

We investigate a new method (QTA-PIMC) for global optimization on complex potential energy surfaces which combines the path integral Monte Carlo method with quantum and thermal annealing. This method is applied to the BLN protein model (Honeycutt, J. D.; Thirumalai, D. *Biopolymers* **1992**, 32, 695). We show that this new approach outperforms simulated (thermal) annealing (SA) and that in fact SA is a subset of our method. This means that we could invoke the effects of both quantum and/or thermal annealing in the global optimization process, hence the name quantum thermal annealing (QTA-PIMC). The simulation results also suggest that QTA-PIMC scales better than SA in terms of system size in the search for the global minimum.

I. Introduction

Global minimization of arbitrary systems is an important problem in many fields of science. A specific class of problems, the so-called NP-complete systems, presents a particularly difficult challenge since the number of local minima in such systems increases exponentially with system size. Examples of such problems abound: the traveling salesman problem in mathematics, spin glasses in condensed matter physics, the routing of chip circuitry in electrical engineering, and the protein folding problem in biophysics. In general, to locate the global minimum of these systems by an exhaustive search of configuration space will take an astronomical amount of time, even for systems which are of only moderate size. Hence efficient methods have to be developed to circumvent this bleak reality. The most widely used unbiased global minimization algorithms include various incarnations of simulated annealing¹ and genetic algorithms.^{2,3} In addition, some useful and promising algorithms have been developed or enhanced by the chemical-, bio-, and condensed matter physics community. These include the Monte Carlo minimization/“basin-hopping” technique,^{4,5} branch and bound methods,^{6–9} multicanonical methods,^{10,11} potential smoothing^{12,13} and classical density annealing methods like the diffusion equation method (DEM),^{14,15} Gaussian density annealing (GDA),^{16,17} and the packet annealing method (PAM).^{18,19} Quantum annealing methods have also recently been investigated.^{20–22} In this paper, we propose a new way to conduct quantum annealing called the QTA-PIMC method, where we use path integral Monte Carlo (PIMC) as the underlying quantum sampling algorithm. This new method also enables us to include the effects of thermal annealing if desired. We suggest that PIMC presents the most natural way to conduct quantum annealing, with advantages over other quantum methods like diffusion Monte Carlo (DMC) or the solving of the imaginary time Schrödinger equation. In addition, we will discuss the relations between QTA-PIMC and the various nonquantum schemes mentioned above, and how QTA-PIMC actually encompasses many of the important and useful features of these methods. It is interesting to note that there is now experimental verification of quantum annealing.²³

In this paper, we apply the new QTA-PIMC method to the global minimization of a 22-mer and a 46-mer of the Honeycutt and Thirumalai protein model.^{24,25} These off-lattice model proteins are constructed from residues based on a three-letter code. The codes represent hydrophobic (B), hydrophilic (L), and neutral (N) residues found in nature. In this paper, we shall refer to this model as the BLN model.²⁶ These model proteins have been shown to display fertile thermodynamic and kinetic behavior. In addition, they exhibit many similarities with theoretical and experimental studies of real proteins.^{24,25,27,28} In fact, it has been shown²⁶ that the 46-mer is a highly frustrated system. Consequently, the BLN protein model is a good and nontrivial system with which to test our algorithm on. We present detailed simulation results of our method versus the traditional method of simulated annealing. The results show that QTA-PIMC outperforms SA in the global minimization of both the 22- and 46-mer. In addition, the simulation results also suggest that QTA-PIMC scales better than SA in terms of system size in the search for the global minimum. Lastly, we mention preliminary results on the global minimization of Lennard-Jones clusters with the QTA-PIMC method. The global minimum of the hard case of the 38-atom Lennard-Jones cluster has been obtained with our method.

The paper is organized as follows. Section II presents the path integral Monte Carlo (PIMC) method for simulating quantum systems. Section III gives an outline of the underlying concepts behind quantum annealing. Section IV describes the BLN protein model and the energy functional to be minimized. Section V presents the new method of QTA-PIMC. In Section VI we present and analyze the simulation results from QTA-PIMC and SA on the global minimization of the 22- and 46-mer mentioned above. In Section VII, we compare and contrast QTA-PIMC with other established methods of global optimization. The conclusions are given in Section VIII.

II. Path Integral Monte Carlo

Feynman’s path integral formulation of quantum theory^{29,30} forms a powerful basis for the investigation of realistic quantum many-body systems, both analytically and computationally. In this formulation, the quantum canonical partition function is written in terms of a path integral^{31,32}

$$Q(N, V, \beta) = \int d\mathbf{x}_1 \int_{\mathbf{x}_1}^{\mathbf{x}_1} D\mathbf{x}(\tau) \exp\left(-\frac{1}{\hbar} S[\mathbf{x}(\tau)]\right) \quad (1)$$

[†] In partial fulfillment of the Ph.D. in the Department of Physics, Columbia University.

where

$$S[\mathbf{x}(\tau)] = \int_0^{\beta\hbar} d\tau \mathcal{H}(\mathbf{x}(\tau)) \quad (2)$$

is the Euclidean action that corresponds to the path $\mathbf{x}(\tau)$ in Euclidean (imaginary) time τ and $\mathcal{H}(\mathbf{x}(\tau))$ contains the path dependence of the Hamiltonian. In eq 1, the integration (represented by $\int_{\mathbf{x}_0}^{\mathbf{x}_1} D\mathbf{x}(\tau)[\dots]$) is taken over all paths starting at $\mathbf{x}_0 = \mathbf{x}_1$ and ending at $\mathbf{x}(\beta\hbar) = \mathbf{x}_1$.

To be computationally useful, the Euclidean time τ is discretized in units of $\epsilon = \beta\hbar/P$. P is an integer that represents the number of “time” slices used in the discretization; it is also known as the Trotter number. It represents the degree of “quantumness” of a system: a strongly quantum system will need a high value for P and a purely classical system has $P = 1$. With this discretization, the continuous path $\mathbf{x}(\tau)$ is approximated by straight-line paths between neighboring Euclidean times. The above is known as the primitive approximation. Within this scheme, the partition function becomes^{30–34}

$$Q_P(\beta) = \left(\frac{Pm}{2\pi\beta\hbar^2} \right)^{3NP/2} \int d\mathbf{r}_{1,1} \dots d\mathbf{r}_{i,t} \dots d\mathbf{r}_{N,P} \exp \left(- \frac{Pm}{2\beta\hbar^2} \sum_{i=1}^N \sum_{t=1}^P |\mathbf{r}_{i,t} - \mathbf{r}_{i,t+1}|^2 - \frac{\beta}{P} \sum_{t=1}^P V_{\text{cl}}(\{\mathbf{r}_i\}; t) \right) \quad (3)$$

Equation 3 is the discrete path integral expression for the quantum canonical partition function in the primitive approximation. It is written for N distinguishable particles in three dimensions which is relevant in this paper, but can easily be extended to N indistinguishable particles in arbitrary dimensions. Note that we have expressed the time-slice dependence explicitly, so that $\mathbf{r}_{i,t}$ represents the 3-vector coordinates of the i -th particle in the t -th time slice. $V_{\text{cl}}(\{\mathbf{r}_i\}; t)$ is the total classical potential energy evaluated at time slice t . Equation 3 can be written in the form of a classical configurational integral

$$Z_{\text{cl}} = \left[\prod_{i=1}^N \prod_{t=1}^P \int d\mathbf{r}_{i,t} \right] \exp(-\beta_t \Phi_P(\{\mathbf{r}_{i,t}\}; \beta_q, \hbar)) \quad (4)$$

Comparing (3) with (4), one can extract the following effective potential

$$\Phi_P(\{\mathbf{r}_{i,t}\}; \beta_q, \hbar) = \left(\frac{Pm}{2\beta_q^2 \hbar^2} \right) \sum_{i=1}^N \sum_{t=1}^P |\mathbf{r}_{i,t} - \mathbf{r}_{i,t+1}|^2 + \left(\frac{1}{P} \right) \sum_{t=1}^P V_{\text{cl}}(\{\mathbf{r}_i\}; t) \quad (5)$$

Equation 5 provides the basis for the interpretation of a quantum system as a classical system of “ring polymers”. This is the well-known *quantum-classical isomorphism*²⁹ that results from the imaginary time path integral formulation of quantum statistical mechanics. The first term arises from the kinetic energy operator and the second from the potential energy operator in the Hamiltonian operator of the system.

It should be noted that we have introduced two different β 's, namely, β_t and β_q in eqs 4 and 5. In a true quantum simulation, one would require that $\beta_t = \beta_q = 1/(k_B T)$, where T is the thermal temperature. For the purpose of quantum annealing, it is useful to separate out the true inverse thermal temperature β_t from what will in fact be a factor (β_q) of the quantum annealing parameter ($Pm/2\beta_q^2 \hbar^2$). In fact, this differentiation between the β 's enables

one to carry out either quantum or thermal annealing separately or both types of annealing at the same time. If P is set to 1 in the above formalism, the system can only be annealed thermally through β_t for it is now a purely classical system, and thus we recover the method of simulated annealing. Hence, the popular method of simulated annealing can be viewed as a *subset* of QTA-PIMC.

III. Quantum Annealing

The objective is to find the global minimum of a *classical* system. In quantum annealing, this is done by first “quantizing” the classical system and then gradually “annealing” it back to the classical regime. In the process, the classical ground state or global minimum of the system is located, i.e.

classical \rightarrow quantum \rightarrow classical ground state

The approach, outlined below, allows for quantum tunneling in place of purely activated barrier crossings.

A general classical potential in condensed matter physics and theoretical chemistry is

$$\Phi_{\text{cl}} = V_{\text{cl}}(\{\mathbf{r}_i\}), \quad i = 1, 2, \dots, N \quad (6)$$

It would be the potential to be sampled from if one were to do simulated annealing with the Metropolis Monte Carlo algorithm. We “quantize” this system via the imaginary time path integral formalism in a straightforward manner. In doing so, we end up with Φ_P (eq 5), which is the “quantized” version of the classical potential Φ_{cl} (eq 6). By sampling Φ_P , for instance with Metropolis Monte Carlo, or by equivalent molecular dynamics schemes,³⁵ one can determine the equilibrium properties of the quantum system represented by Φ_P . As mentioned before, the degree of “quantumness” of the system is controlled through the Trotter number P and \hbar . [$\hbar = h/2\pi$ where h is Planck's constant. However, in the context of quantum annealing, \hbar is used as an adjustable *parameter*.] By suitably adjusting the values of P and \hbar , it will be possible to bring the quantum system back to the classical regime, thereby locating the global minimum. These are the essential ideas behind QTA-PIMC. The implementation details of our method will be given in Section V.

It is clear from the path integral formulation that each quantum degree of freedom is described by a one-dimensional array of classical degrees of freedom (the size of the array is the Trotter number P). Consequently, the dimensionality goes from $3N$ to $3NP$ when one goes from the quantum to an isomorphic classical picture.^{31,36} Hence quantum annealing methods entail an increase in the dimensionality of the original problem. This increase in dimensionality corresponds to the quantum dispersion associated with the uncertainty principle of quantum mechanics.³⁷ This allows for the *nonlocal* exploration of the original potential energy surface (PES). In addition, it also allows for the effect of quantum tunneling so that classically forbidden energy barriers can be overcome.³⁷

It is important to consider the extra CPU cost that is associated with this increase in dimensionality. It is useful to keep in mind that we are interested in the global minimum of *classical* systems. Since these systems are not quantum in the first place, it may not be necessary to choose large Trotter numbers P . We note that if the original classical system has N particles, the use of the path integral formulation means that there will now be NP particles. The CPU time for Monte Carlo simulations on classical systems with untruncated nonbonded potentials scales as N^2 . Since Trotter beads from different time slices do not interact through the external potential, the CPU time required

for the corresponding potential calculations in PIMC would scale as N^2P . However, in our QTA-PIMC scheme, the number of Trotter time slices P is actually systematically reduced as we proceed with the simulation, and consequently the CPU cost is kept under control as we go from the quantum to the classical regime. The above is reminiscent of the “relaxation of dimensionality” method of Scheraga and co-workers.^{38–40} However, QTA-PIMC allows for a more general, systematic, and easier scheme for such a purpose. The increase in particle number is being offset by the increased search efficiency on the PES, and ultimately QTA-PIMC outperforms a purely classical method like simulated annealing, as we shall see in Section VI.

IV. Model Protein Potential

The system used in our simulations is the BLN protein model of Honeycutt and Thirumalai.^{17,24,25,41} Protein sequences are built from residues of three types, hydrophobic (B), hydrophilic (L), and neutral (N). In this paper, we shall study in detail the global minimization of a 22-mer $(\text{LB})_5\text{N}_2(\text{LB})_5$, and that of a longer protein chain, the 46-mer $\text{B}_9\text{N}_3(\text{LB})_4\text{N}_3\text{B}_9\text{N}_3(\text{LB})_5\text{L}$.

The potential energy of the BLN protein model used here is made up of the following terms:¹⁷

$$E_p(\{\mathbf{r}_i\}) = V_{\text{bond}}(\{\mathbf{r}_i\}) + V_{\text{bend}}(\{\theta_i\}) + V_{\text{tors}}(\{\phi_i\}) + V_{\text{non}}(\{r_{ij}\}) \quad (7)$$

with

- bond-length potential

$$V_{\text{bond}}(\{\mathbf{r}_i\}) = \sum_{i=1}^{N-1} \frac{k_r}{2} (|\mathbf{r}_{i+1} - \mathbf{r}_i| - a)^2 \quad (8)$$

where $k_r = 400\epsilon_h/a^2$, a is the average bond length between two residues, and ϵ_h , the average strength of the hydrophobic interaction, serves as the unit of energy in this model.

- bond-angle potential

$$V_{\text{bend}}(\{\theta_i\}) = \sum_{i=1}^{N-2} \frac{k_\theta}{2} (\theta_i - \theta_0)^2 \quad (9)$$

where θ_i is defined by residues i , $i+1$ and $i+2$, $k_\theta = 20\epsilon_h/(\text{rad})^2$, and $\theta_0 = 1.8326$ rad or 105° .

- dihedral-angle potential

$$V_{\text{tors}}(\{\phi_i\}) = \sum_{i=1}^{N-3} [A_i(1 + \cos\phi_i) + B_i(1 + \cos 3\phi_i)] \quad (10)$$

where ϕ_i is defined by residues i , $i+1$, $i+2$, and $i+3$. If two or more of the four defining residues of ϕ_i are neutral (N), then $A_i = 0\epsilon_h$, $B_i = 0.2\epsilon_h$, otherwise $A_i = B_i = 1.2\epsilon_h$.

- nonbonded potential

$$V_{\text{non}}(\{r_{ij}\}) = \sum_{i=1}^{N-3} \sum_{j=i+3}^N V_{\alpha\beta}(r_{ij}) \quad (11)$$

where $r_{ij} = |\mathbf{r}_i - \mathbf{r}_j|$ and $\alpha, \beta = \text{B, L, or N}$. $V_{\alpha\beta}$ represents soft-sphere type repulsions or Lennard-Jones interaction between residues α and β :

$$V_{\text{L}\beta}(r_{ij}) = 4\epsilon_L \left[\left(\frac{\sigma}{r_{ij}} \right)^{12} + \left(\frac{\sigma}{r_{ij}} \right)^6 \right] \quad (\beta = \text{B or L, } \epsilon_L = 2/3\epsilon_h) \quad (12)$$

$$V_{\text{N}\beta}(r_{ij}) = 4\epsilon_h \left[\left(\frac{\sigma}{r_{ij}} \right)^{12} \right] \quad (\beta = \text{B, N, or L}) \quad (13)$$

$$V_{\text{BB}}(r_{ij}) = 4\epsilon_h \left[\left(\frac{\sigma}{r_{ij}} \right)^{12} - \left(\frac{\sigma}{r_{ij}} \right)^6 \right] \quad (14)$$

In addition, we employ a weak boundary potential¹⁷

$$V_{\text{bp}}(\{\mathbf{r}_i\}) = \sum_{i=1}^N \frac{k_b}{2} (\mathbf{r}_i - \mathbf{r}_{\text{com}})^2 \quad (15)$$

This potential prevents the protein chain from dissociating at high temperatures, and it also encourages folding of the protein. \mathbf{r}_{com} is the center of mass position of the protein chain and k_b is reduced from 0.05 to 0.005 during the course of our simulations. At the end of each simulation, k_b is set to 0.0 before we refine the energy of the resulting configuration(s) with the conjugate gradient algorithm.

Unless otherwise stated, all physical quantities are in reduced units, i.e., the thermal temperature T is in units of ϵ_h/k_B and energies are in units of ϵ_h . In our simulations, the mass of each residue m , the bond length a , the hydrophobic energy constant ϵ_h , the Lennard-Jones parameter σ , and the Boltzmann constant k_B are set to unity. The value of the *true* \hbar is unity as well. However, since \hbar is used as an adjustable parameter in QTA-PIMC, it can take on other values.

It should be noted that the Hamiltonian of the BLN protein model is invariant under inversion $\mathbf{r}_i \rightarrow -\mathbf{r}_i$. Consequently, there are two configurations corresponding to each energy, including the global minimum energy; see Table 1. Hence the global minimum is twofold degenerate for each protein.

V. Quantum Thermal Annealing with Path Integral Monte Carlo

A. Quantum-Classical Isomorphism. In Figure 1, the “quantization” of a protein chain via the path integral formulation is shown. We see that each residue, in its quantum form, is now represented by P Trotter beads. The lines joining the beads on each residue signify the harmonic potential coupling between them.

B. Path Sampling Methods. We use Monte Carlo sampling of the integrands in eqs 3 or 4 as follows:^{31,32}

1. *Local Moves.* A single Trotter bead on each residue is displaced randomly within a cube of size Δ_{local} and the Metropolis accept/reject criterion is applied to this trial move. During the simulation, Δ_{local} is adjusted so as to maintain a 50% acceptance rate. Each set of local moves consists of a sequential pass through the Trotter beads of all the residues.

2. *Global Moves.* The entire chain of P beads representing each residue is translated by the same amount δ and the Metropolis accept/reject criterion is applied to this trial move. The size of the displacement δ is sampled randomly from within a cube of size Δ_{global} . During the simulation, Δ_{global} is adjusted so as to maintain a 50% acceptance rate. Each set of global moves consists of a sequential pass through the Trotter chains of all the residues.

In regular Monte Carlo simulations, one Monte Carlo sweep is defined as the attempted moves of N particles, where N is the size of the system. In QTA-PIMC, since Trotter beads from different time slices do not interact through the external classical potential, the attempted moves of all $N \times P$ beads once for a set of global moves would involve P Monte Carlo sweeps. The attempted moves of all $N \times P$ beads once for a set of local moves would involve P Monte Carlo sweeps as well, although

TABLE 1: (a) Global Minimum and the Next Three Lowest Energy Structures of the 22-mer $(\text{LB})_5\text{N}_2(\text{LB})_5^a$ and (b) Corresponding Data for the 46-mer $\text{B}_9\text{N}_3(\text{LB})_4\text{N}_3\text{B}_9\text{N}_3(\text{LB})_5\text{L}^b$

(a) 22-mer $(\text{LB})_5\text{N}_2(\text{LB})_5$		
name	energy	rmsd from ref struct 22.1A
22.1A	-3.060	0.000
22.1B		0.579
22.2A	-2.728	0.392
22.2B		0.581
22.3A	-2.633	0.624
22.3B		0.745
22.4A	-1.651	0.823
22.4B		N/A

(b) 46-mer $\text{B}_9\text{N}_3(\text{LB})_4\text{N}_3\text{B}_9\text{N}_3(\text{LB})_5\text{L}$		
name	energy	rmsd from ref struct 46.1A
46.1A	-49.264	0.000
46.1B		0.862
46.2A	-49.186	0.424
46.2B		0.946
46.3A	-49.149	0.450
46.3B		0.960
46.4A	-49.063	0.577
46.4B		1.050

^a The twofold degeneracy is due to the fact that the BLN Hamiltonian is invariant under inversion $r_i \rightarrow -r_i$. The global minimum structures are labeled 22.1A and 22.1B, and the next lowest 22.2A, 22.2B, and so on. The rmsd of these structures from 22.1A are listed in the last column. ^b Note that numerically the energy of the minima differ in the first decimal place, suggesting that the 46-mer has a glassy potential energy surface.




Figure 1. Quantum-classical isomorphism. A classical molecule is “quantized” with the path integral formalism into cyclic chain polymers. Each Trotter bead on the chains is labeled by a number that indicates the time-slice that it belongs to.

in addition there is a slight overhead due to the calculation of the harmonic potential between adjacent beads. In this paper, we define the attempted moves of all $N \times P$ Trotter beads once as a PIMC pass. Consequently, one PIMC pass of either local or global moves entails P Monte Carlo sweeps. In order to compare the efficiency of QTA-PIMC versus SA, we use the same number of MC sweeps (thus essentially the same CPU time) for each method.

Note that there are more advanced path sampling techniques like normal-mode sampling,^{42–45} the staging algorithm,^{35,46} and others.^{31,32} However, due to the relatively small number of Trotter beads P used in this study, the use of just local and global moves are sufficient.

C. Simulation Scheme. An initial protein configuration is constructed by assigning residue sites randomly on an extended zig-zagged line (Figure 2a). The distance between each residue site is approximately a . P_0 Trotter beads representing each residue are then placed overlapping one another on these sites. Alternatively, we run a high-temperature walker with Langevin




Figure 2. Initial configurations used in our simulations. Each filled circle represents P_0 overlapping Trotter beads at each protein residue site. (a) A random extended initial configuration. (b) A nonextended initial configuration generated using Langevin dynamics.

dynamics at $T = 5.0$ and keep configurations every τ time steps. Each of these configurations is then quenched down to $T = 1.0$ and equilibrated for another τ time steps before the final configuration is stored. For the 22-mer, τ was 20 000, and for the longer 46-mer, τ was 50 000. For each of these final configurations, we place P_0 overlapping Trotter beads at each residue site. These configurations serve as nonextended initial configurations (Figure 2b) for our simulations. [We find that both sets of initial configurations give comparable results in terms of average success rate for locating the global minimum of the 22-mer. Either set should serve as good initial configurations for an arbitrary BLN protein.]

Starting with one of the above initial configurations, the simulation is carried out as follows:

1. Set the initial value of the Trotter number P to P_0 , \hbar to \hbar_i , and the thermal temperature T to T_i .
2. Perform quantum annealing as follows: Keeping the number of Trotter beads P and the thermal temperature T fixed, reduce \hbar from an initial value \hbar_i to a final value \hbar_f linearly in n_{dh} steps. [The value of \hbar_f is typically 0.0001, i.e., close to 0.0.] At each value of \hbar , perform n_{global} global and n_{local} local PIMC passes. A fraction of these passes are used for equilibration during which Δ_{global} and Δ_{local} are adjusted so as to maintain a 50% acceptance rate. At the end of this \hbar annealing stage, when $\hbar = \hbar_f$, all P Trotter beads representing each residue would have converged back toward a single point (since the classical regime has been reached). We refer to this whole \hbar annealing cycle (where \hbar is reduced from \hbar_i to \hbar_f) as a *quantum annealing stage*, and the configuration produced at the end of this cycle an *intermediate* classical configuration. Keep this configuration for later.
3. Reduce P by ΔP , T by dT , k_b by dk_b and reset \hbar to \hbar_i , then go back to step 2. This process is repeated until $P = 1$ (classical regime) and a final temperature of T_f . The total MC sweeps in step 2 is kept constant by adjusting n_{global} and n_{local} appropriately. [If $dT = 0$, we have plain quantum annealing. If $dT \neq 0$, we

introduce thermal annealing into our method. Hence QTA-PIMC enables us to utilize the effects of quantum and/or thermal annealing.]

4. Apply the conjugate gradient algorithm to the stored configurations from step 2 and the final configuration obtained at the end of the whole simulation from step 3 to refine the energies. The binding potential V_{bp} is not needed at this point, so it is switched off (by setting $k_b = 0.0$) before energy refinement. The global minimum of the protein would be obtained at this point.

In the description above, we see that the molecule is brought back to the classical regime in the QTA-PIMC method in *two* ways. In step 2, this is done by having \hbar being gradually reduced to ≈ 0 even though $P \neq 1$. As mentioned earlier, we call the configurations so produced the *intermediate* classical configurations. At the end of step 3, the classical regime is attained by having both $\hbar \rightarrow 0$ and $P \rightarrow 1$. We call this configuration the *final* classical configuration. The reason for keeping the intermediate classical configurations in step 2 is that the global minimum is often already found at these stages (where $P \neq 1$) instead of at the very end of the simulation after step 3. This possibility of locating the global minimum early in the simulation is another advantage of using QTA-PIMC over traditional methods like simulated annealing. A nice illustration and discussion of this is given in Section VIB.

The reason for performing energy refinement with the conjugate gradient algorithm in step 4 is to pinpoint the global minimum out of the set of local minima which have energies very close in value (see Table 1) to that of the true global minimum.

A schematic diagram of the above procedure is shown in Figure 3. Actual numerical values used for the parameters are given in Section VI.

VI. Results and Analysis

In a previous study¹⁷ of classical density annealing optimization algorithms (Gaussian phase packet kinetic annealing and adiabatic Gaussian density annealing) using the BLN model, good results were obtained with 22-mers. However, these methods apparently do not do very well with the longer 46-mer. In this paper, in addition to the 22-mer $(LB)_5N_2(LB)_5$ we shall also test our method on the more difficult 46-mer $B_9N_3(LB)_4N_3B_9N_3(LB)_5L$. From now on whenever we mention the 22-mer and 46-mer we will be referring to the aforementioned sequences.

A. The Global Minimum of the 22-mer and 46-mer. In Table 1, we list the (twofold degenerate) global minimum and the next three lowest energy structures of the 22-mer and 46-mer, respectively. As mentioned before, since the BLN model is invariant under $\mathbf{r}_i \rightarrow -\mathbf{r}_i$, there are two configurations corresponding to the same energy value, hence the twofold degeneracy. These are the lowest energy structures that we have obtained out of all of our simulation runs. They were first found by conducting 10 long simulation runs of QTA-PIMC and SA. No lower energy structures were found in subsequent runs using either methods.

The global minimum energy found here for the 22-mer is in fact lower than that found by Amara et al.¹⁷ A comparison of the structure in their Figure 1 with our structures indicates that their global minimum is in fact our second lowest minimum. We cannot tell if we have found a lower energy for the global minimum of the 46-mer since they did not quote a numerical value for the global minimum energy found.

From Table 1, it is clear that the energy differences between the local minima of the 22-mer are more pronounced than that




Figure 3. QTA-PIMC scheme used in this paper, illustrated with a single classical particle. The classical particle is quantized into P Trotter beads (shaded circles). The quantized particle is annealed back to the classical regime in two ways: by decreasing \hbar , and by decreasing P . For a fixed P , the quantized particle is annealed by decreasing the value of \hbar . When \hbar reaches 0, all the Trotter beads would converge back to a single point which corresponds to an intermediate classical configuration (solid circle). At this juncture, we remove one or more of the overlapping Trotter beads. We reset the value of \hbar to \hbar_i , reduce T by dT (if thermal annealing is desired), and repeat the whole process again. The final classical configuration is obtained when both P and \hbar are annealed to the classical values of 1 and 0, respectively. This configuration would be the global minimum of the system, although often the global minimum would already have been found in one of the intermediate stages.

of the 46-mer. It is known^{17,25} that the 46-mer has a huge number of local minima at low energy as one approaches the global energy minimum. In fact, a recent study²⁶ shows that the 46-mer is a highly frustrated system. In our investigation, we found that the energies of the 4 lowest energy configurations differ essentially in the first decimal place. In addition, while these local minimum structures do take on the form of barrel-like structures and look very similar, there is no hard evidence that they are actually in the same big funnel on the potential energy surface. This supports the notion that the 46-mer has a much more glassy (or rougher) potential energy surface than the 22-mer. Consequently, this makes consistent location of the global minimum of the 46-mer more challenging.

B. Results from a Simulation Run of QTA-PIMC. In this section, we present results of using QTA-PIMC in the global minimization of the 46-mer in a typical simulation run. It is useful at this point to refer back to the simulation scheme given in Section VC. In this run, we started with 20 Trotter beads (i.e., $P_0 = 20$) and reduced P by 1 (i.e., $\Delta P = 1$) at the end of each quantum annealing stage until we were back to the classical regime ($P = 1$). The numerical values of the other parameters used were $T_i = 0.1$, $T_f = 0.01$, $dT = 0.005$, $\hbar_i = 2.0$, $\hbar_f = 0.0001$, $n_{dh} = 200$, $n_{global} = n_{local} = 200$. We monitored




Figure 4. Potential energy of the 46-mer versus MC sweeps. The solid line corresponds to the energy of the configuration represented by time slice 0 of the path integral. The dotted line corresponds to the lowest energy configuration among all the time slices.




Figure 5. Root-mean-square deviation (rmsd) of the 46-mer versus MC sweeps. The configuration being monitored is the one corresponding to the 0th time slice of the path integral. The rmsd is calculated with respect to the global minimum structure 46.1A. The dotted line shows the reduction in Trotter number P as the simulation progresses. See text for explanation of the labels in the figure.

the progress of the protein configuration corresponding to Trotter time-slice 0 (cf Figure 1), unless otherwise stated.

In Figure 4, we show the energy of the 46-mer as the simulation proceeded. At the end of each quantum annealing stage, when the 46-mer was back in the classical regime, it was also in a local minimum on the potential energy surface. The repeated returns to the classical regime at the end of each quantum annealing stage allowed for multiple attacks on the PES in locating the global minimum. In a way, this is reminiscent of the “basin-hopping” method of Wales and Doye.⁵ We will look at this in more detail in Section VII.

In Figure 5, we show the RMSD (root-mean-square deviation from structure 46.1A) of the 46-mer configuration as the simulation progressed. In the same figure, we show the reduction schedule for the number of Trotter time slices P . We see that the 46-mer reached the global minimum at the end of the $P = 16$ annealing stage, this is indicated by arrow (a) in Figure 5. Thus, QTA-PIMC was able to locate the global minimum of the 46-mer with just *one-quarter* of the target total MC sweeps. It is interesting to note that in the next two quantum annealing stages, where $P = 15$ and $P = 14$, respectively, the 46-mer actually tunneled out of its global minimum before going back




Figure 6. Root-mean-square “size” of the classical ring polymers representing the quantum cloud of the residues in the path integral formulation versus MC sweeps.

to intermediate classical regimes (where $\hbar \approx 0.0$) again. This is indicated by arrows (b) and (c) in Figure 5. However, in the next annealing stage where $P = 13$ (indicated by arrow (d)), and beyond, the 46-mer stayed in the global minimum basin. We note that this “early detection” power of QTA-PIMC was not unique to this particular run, but is a feature observed in most of the runs that managed to find the global minimum.

In Figure 6, we show the root-mean-square “size” of the classical ring polymers as the simulation progressed. This was calculated with^{47,48}

$$r_{\text{size}}^2 = \frac{1}{N} \sum_{i=1}^N \left(\frac{1}{P} \sum_{t=1}^P |\mathbf{r}_{i,t} - \mathbf{r}_{i,\text{com}}|^2 \right) \quad (16)$$

where

$$\mathbf{r}_{i,\text{com}} = \frac{1}{P} \sum_{t=1}^P \mathbf{r}_{i,t} \quad (17)$$

is the center of mass (or centroid) of the Trotter chain of the i -th residue. r_{size} is a measure of the size of the quantum cloud associated with each quantized residue. The simulation results show that overall, as P was reduced, the size of the cloud also decreased, as expected. During each quantum annealing stage, when P was fixed, but \hbar was decreased from \hbar_i to \hbar_f , the size of each cloud also decreased, all the way to ≈ 0.0 at the end of the stage. The above observations are as expected from eq 5. Every time the quantum cloud collapsed back to a single point, we were back in an intermediate classical regime. The above clearly illustrates the ability of QTA-PIMC to explore the potential energy surface non-locally.

Finally, in Figure 7, we show the radius of gyration of the 46-mer as the simulation proceeded. We observe that the 46-mer was in compact configurations for a large part of the simulation. This indicates that QTA-PIMC did not waste unnecessary time in exploring noncompact configurations but was actually sorting out the global minimum from the set of nearly folded structures.

C. Comparison of Results: QTA-PIMC versus SA. In this section, we compare QTA-PIMC with simulated annealing in the global optimization of the 22-mer and 46-mer. The annealing scheme for QTA-PIMC is given in Section VC and will not be repeated here. For SA, we follow the annealing scheme of




Figure 7. Radius of gyration R_g of the 46-mer versus MC sweeps. The configuration being monitored is the one corresponding to the 0-th time slice of the path integral.

Thirumalai and co-workers.^{25,41} They used a linear annealing schedule where the thermal temperature is decreased linearly

$$T_k = T_0 - k\Delta T, \quad k = 0, 1, 2, \dots, (n_{dT} - 1) \quad (18)$$

where T_0 is the initial temperature, ΔT is the change in temperature for each annealing step, and n_{dT} is the total number of annealing steps. (We adopted the same thermal annealing scheme in QTA-PIMC.) We conducted n_T Monte Carlo sweeps at each temperature, and 20% of these were used for equilibration after each temperature change. The total MC sweeps used for each run was thus $n_{dT} \times n_T$. During the equilibration stage at each temperature, the maximum displacement Δ for Monte Carlo moves was adjusted so as to maintain a 50% acceptance rate. The initial configurations were produced using Langevin dynamics as described in Section VC, but here of course there was only 1 particle per residue site (i.e., $P = 1$).

1. 22-mer $(\text{LB})_5\text{N}_2(\text{LB})_5$

In Figure 8, we present results for 20 independent simulation runs of QTA-PIMC with the 22-mer. The total number of MC sweeps used in each simulation run was 1.6×10^6 . In each of these runs, $P_0 = 7$, $\Delta P = 2$, i.e., P was reduced from $7 \rightarrow 5 \rightarrow 3 \rightarrow 1$. The temperature T was fixed at 0.1, $\hbar_i = 3.0$, $\hbar_f = 0.0001$, $n_{dh} = 650$, and $n_{\text{global}} = n_{\text{local}} = 50$. We also did another set of 20 runs where all the parameters remained the same except \hbar_i was changed to 4.0. QTA-PIMC was able to achieve 100% success rate in locating the global minimum of the 22-mer with either of these schedules, as shown in Figure 8a. Note that the temperature T was held fixed. We did this to show that the effects of quantum annealing alone (without thermal annealing) was sufficient to obtain perfect results in the global optimization of the 22-mer. In Figure 8a, we also show results for the global minimization of the 22-mer using simulated annealing. As in QTA-PIMC, we conducted 20 independent runs for each annealing schedule and used 1.6×10^6 MC sweeps for each of these runs. The initial and final temperatures T_0 and T_f were set to 1.0 and 0.0001, respectively. We tried four different annealing schedules where we varied n_{dT} and n_T while keeping the total number of MC sweeps fixed. The 4 schedules used had $n_{dT} \times n_T = 32 \times 50$, $000,320 \times 5000$, 3200×500 , and $32\,000 \times 50$. They are labeled as SA1, SA2, SA3, and SA4, respectively in Figure 8a, SA(avg) is the average of SA1 to




Figure 8. (a) Success rates of locating the global minimum of the 22-mer with QTA-PIMC and SA. 20 independent trials are performed to produce each set of results. The filled bar of SA(avg) is obtained by averaging over the success rates from SA runs using four different schedules (SA1 to SA4). (b) Minimum energies obtained from these simulation trials. The dotted line indicates the global minimum energy value of the 22-mer. Note that some of the trials produce results with the same energy values; these overlap one another on the graph. QTA-PIMC is able to locate the global minimum with a 100% success rate, hence all the energy values fall on the same point (the filled circle).

SA4. We increased n_{dT} by an order of magnitude from one schedule to the next in order to see the dependence of the optimization results on the annealing parameters. From Figure 8a, we see that all four annealing schedules gave comparable results (i.e., there was no order-of-magnitude difference) in the success rate. The average success rate was 68%. In Figure 8b, we show the spectrum of minimum energies obtained by QTA-PIMC and SA. QTA-PIMC was able to locate the global minimum consistently. SA, on the other hand, was prone to getting stuck at higher energy metastable states. Consequently, there is a large spread of energy values for SA, as shown in Figure 8b.

We also increased the number of Monte Carlo sweeps in order to find out how many more sweeps are needed before we can achieve 100% success rate with simulated annealing. As before, we tried four different annealing schedules. In Table 2, we summarize our findings. A graphical counterpart of the table is given in Figure 9. We observe that the average success rate increased, as expected, if we increased the number of MC sweeps used. However, we note that there was no “best” schedule which consistently gave the highest success rate, therefore we averaged over the success rates for the four schedules. From Table 2, we see that a 100% success rate was first obtained by SA if 7×10^6 MC sweeps were used. With QTA-PIMC, a 100% success rate was achieved with just $1.6 \times$

TABLE 2: Results of SA for the 22-mer Using Different Number of Total MC Sweeps and Different Annealing Schedules^a

total MC sweeps	$n_{dT} \times n_T$	success percentage	av success rate (%)
1.0×10^6	20000×50	45	50
	2000×5000	65	
	200×50000	45	
2.0×10^6	40000×50	80	70
	4000×500	65	
	400×50000	70	
3.0×10^6	60000×50	90	75
	6000×500	75	
	600×50000	65	
4.0×10^6	80000×50	90	85
	8000×500	95	
	800×50000	80	
5.0×10^6	100000×50	90	89
	10000×500	90	
	1000×5000	80	
6.0×10^6	120000×50	95	86
	12000×500	85	
	1200×5000	95	
7.0×10^6	140000×50	75	86
	14000×500	95	
	1400×5000	75	
8.0×10^6	140×50000	100	89
	160000×50	85	
	16000×500	90	
9.0×10^6	1600×5000	90	90
	160×50000	90	
	18000×500	95	
10.0×10^6	180×50000	80	91
	200000×50	85	
	20000×500	85	
	2000×5000	100	
	200×50000	95	

^a Each line of results is obtained from 20 independent SA runs.

10^6 MC sweeps. Since the number of MC sweeps used is proportional to the CPU time utilized, we estimate that QTA-PIMC is over 4 times more efficient than SA in the global minimization of the 22-mer. In addition, we observe that with SA, the average success rate reached a maximum of about 90% with 5×10^6 MC sweeps. Beyond this, even though we increased the number of MC sweeps, the average success rate did not increase accordingly but actually plateaued out, staying approximately constant at about 90%. This is illustrated in Figure 9.

2. 46-mer $B_9N_3(LB)_4N_3B_9N_3(LB)_5L$

Next, we apply QTA-PIMC and SA to the global minimization of the 46-mer. The total number of MC sweeps used in each simulation run for either method was 32×10^6 . The strategy used is similar to that used for the 22-mer.

In Figure 10, we present results obtained from 20 simulation runs of QTA-PIMC. In each of these runs, $P_0 = 20$, $\Delta P = 1$, $\hbar_i = 4.0$, $\hbar_f = 0.0001$, $n_{d\hbar} = 800$, and $n_{\text{global}} = n_{\text{local}} = 50$. The initial and final temperatures were $T_i = 0.2$ and $T_f = 0.02$, respectively. The temperature was reduced once after each quantum annealing stage until it reached T_f . We see from Figure 10a that QTA-PIMC was able to achieve a 60% success rate in




Figure 9. Success rates in locating the global minimum of the 22-mer using QTA-PIMC and SA. The results are plotted with respect to the number of MC sweeps used. The error bar on each SA data point is obtained by averaging over the results of four SA annealing schedules. The filled circle indicates that QTA-PIMC is able to locate the global minimum with 100% certainty using just 1.6×10^6 MC sweeps. SA, on the other hand, is not able to achieve that, up to the usage of 10×10^6 MC sweeps.

locating the global minimum of the 46-mer using 32×10^6 MC sweeps. In this case, we utilized the power of both quantum and thermal annealing. Notice that the temperature used ranged from 0.2 to 0.02. These were lower than the unfolding–folding transition temperature T_F of the 46-mer, which is approximately 0.65.²⁵ The possibility of using low thermal temperatures for QTA-PIMC simulations means that low-temperature structures could be probed more effectively throughout the simulations. In other words, the algorithm spent most of its time probing near the local minima of the PES rather than wasting too much time searching through the high-energy regions. This increased the chances of locating the global minimum. In Figure 10a, we also show results for the global minimization of the 46-mer using simulated annealing. As before, we conducted 20 runs for each annealing schedule. The initial and final temperatures T_0 and T_f were set to 1.0 and 0.0001, respectively. We tried five different annealing schedules where we varied n_{dT} and n_T while keeping the total number of MC sweeps fixed at 32×10^6 . The five schedules used had $n_{dT} \times n_T = 64 \times 50\,000$, $640 \times 50\,000$, 6400×5000 , $64\,000 \times 500$, and $640\,000 \times 50$. They are labeled as SA1, SA2, SA3, SA4, and SA5, respectively, in Figure 10a, SA(avg) is the average of SA1 to SA5. We increased n_{dT} by an order of magnitude from one schedule to the next in order to see the dependence of the optimization results on the annealing parameters. As in the case of the 22-mer, we see that all of the annealing schedules gave comparable results in the success rate. We take the average of these to find an average success rate of 9%. Consequently, QTA-PIMC on average is about 6.7 times more effective than simulated annealing in terms of the *success rate* in locating the global minimum of the 46-mer. In Figure 10b, we show the spectrum of minimum energies obtained by QTA-PIMC and SA. As in the case of the 22-mer, SA was prone to getting stuck at higher energy metastable states. This results in the larger spread of energy values in Figure 10b.

Note that the success rate is a quite stringent test of the effectiveness of a global minimization algorithm. For the 22-mer, we fixed the total number of MC sweeps used at 1.6×10^6 and found that the average success rates for QTA-PIMC and SA were 100% and 68%, respectively. This means that QTA-PIMC is about 1.5 times more effective than SA in terms




Figure 10. (a) Success rates of locating the global minimum of the 46-mer with QTA-PIMC and SA. 20 independent trials are performed to produce each set of results. The filled bar of SA (avg) is obtained by averaging over the success rates from SA runs using 5 different schedules (SA1 to SA5). (b) Minimum energies obtained from these simulation trials. The dotted line indicates the global minimum energy value of the 46-mer. Note that some of the trials produce results with the same energy values; these overlap one another on the graph. Since SA tends to produce structures that get stuck at higher energy metastable states, we see a larger spread in the SA minimum energy spectrum.

of average success rates for the global minimization of the 22-mer. Doing the same type of investigation for the case of the 46-mer, we found that if the total MC sweeps were held fixed at 32×10^6 , the average success rates for QTA-PIMC and SA were 60% and 9%, respectively. This means that QTA-PIMC is over 6 times more effective than SA in terms of average success rates for the global minimization of the 46-mer. The efficiency of QTA-PIMC over simulated annealing measured in terms of the amount of MC sweeps (and hence CPU time) required to achieve the same success rate is expected to be much larger, since the success rate scales nonlinearly with the number of MC sweeps used.

Thus, by comparing the relative success rates of QTA-PIMC and SA in locating the global minimum of the 22-mer and 46-mer, we see that QTA-PIMC does much better at a higher residue number N . This suggests that QTA-PIMC scales better than SA with respect to N and is consequently more effective in dealing with larger systems.

VII. Discussion

It is useful to compare and contrast QTA-PIMC with some of the other established methods of global minimization.

Simulated annealing (SA) is one of the most popular and powerful methods for performing global optimization. However, it allows only thermal annealing while QTA-PIMC encompasses

both quantum and thermal annealing and is therefore a stronger method. We have shown in this paper that QTA-PIMC outperforms SA in the global minimization of two model protein sequences. QTA-PIMC can be viewed as an extension of SA into the quantum regime.

QTA-PIMC is relatively easy to implement in comparison to certain other quantum annealing methods like those that involve solving of the imaginary time Schrödinger equation. In that approach, an approximate solution to the equation involves the extra step of finding the effective potential of the system with a set of Gaussian functions. In addition to being approximate, this may be tedious to do for a general complex potential. A better way to conduct pure quantum annealing which does not suffer from the aforementioned difficulties is with diffusion Monte Carlo (DMC).²¹ With pure quantum annealing, the quantum ground state wave function at $T = 0$ is first obtained. This is then annealed to the classical regime by letting $\hbar \rightarrow 0$. However, there are certain problems associated with such a scheme. We will describe two scenarios where pure quantum annealing will likely have difficulties locating the correct classical ground energy minimum. As a first example, we consider the quantum anisotropic rotor potential. For $B > 0.4$, we know that the system goes from an orientationally disordered state at high temperatures to an ordered state below a certain temperature. However, as the system is cooled further it undergoes another phase transition to a disordered state. This is called *reentrance*.⁴⁹ Thus the quantum ground state is disordered. Presumably this is what a quantum algorithm like DMC would find. The classical ground state is, however, ordered. As a second example, we consider a one-dimensional potential with a very broad well separated from a very narrow well by an energy barrier. If the narrow well is slightly lower in energy than the broad well the ground state will have a very large density in the broad well. Pure quantum annealing will then not easily locate the global minimum. A series of DMC simulations with such a potential verified this difficulty: the annealed distribution consistently ended up in the broad well rather than the global minimum which is the narrow well. QTA-PIMC does not suffer from such a problem and was able to locate the narrow well consistently.

In Section VIB, we mentioned that QTA-PIMC is similar to the “basin-hopping” method (BH) in that local minima are probed at intermediate stages of the annealing run. However, “basin-hopping” clearly entails large computational overhead since the local minimization has to be done after each Monte Carlo step.⁵ Since we are simulating at a relatively low temperature with QTA-PIMC, we expect the sampled configurations to be close to local minima. As a result, the classical global minimum could be found more readily. This was illustrated in Section VIB. In BH, the search is essentially a series of random quenches to local potential minima and one hopes that one of these would correspond to the global minimum. In the QTA-PIMC method, we can envision the quantum clouds of the particles probing the PES in an amoeba-like manner and tunneling through energy barriers if need be. Consequently, QTA-PIMC is able to search the PES more intelligently.

Multicanonical algorithms (MUCA) have become an useful method for calculating thermodynamic properties and the identification of the global minimum of physical and chemical systems. MUCA’s advantage is that a random walk in energy space means that energy barriers between local minima do not present a problem, unlike simulated annealing in which one relies solely on thermal processes to overcome potential barriers.

QTA-PIMC is also able to go through energy barriers via quantum tunneling.

In classical density annealing methods, instead of moving what is essentially a point through phase space (MD methods) or configuration space (MC methods), a density distribution is used instead. This “blob” is able to sample phase or configuration space nonlocally. For quantum annealing with QTA-PIMC, where we have Trotter beads representing the quantum clouds of the quantized system, it is clear that we have the advantage of nonlocal sampling as well.

VIII. Conclusions

In this paper, we have presented a new technique (QTA-PIMC) for performing global optimization on complex potential energy surfaces. This approach makes use of path integral Monte Carlo, a powerful quantum sampling algorithm, to conduct quantum and thermal annealing. While traditional simulated annealing uses only thermal effects in the search for the global minimum, our method makes use of both quantum and thermal processes. We have shown that SA is in fact a subset of this new technique. Instead of having just one point wandering through configuration space, our method allows for nonlocal search on the potential energy surface. In addition, while SA might allow a system to get trapped in metastable states, QTA-PIMC enables the system to tunnel through these states, thus ensuring a higher chance of locating the global minimum.

We compared the performance of our method with that of SA by conducting global optimization of a 22- and a 46-mer of the BLN protein model. We found that our method outperformed SA in both cases. In addition, our results also suggest that QTA-PIMC actually scales better with system size than SA in the search for the global minimum. It is expected that even better annealing schemes for QTA-PIMC could be devised that would outperform the ones that have been used in this paper.

The efficiency of our new method may potentially be improved further by combining it with classical methods such as the fluctuating potential method,⁵⁰ the Smart-Walking method,⁵¹ multicanonical methods,^{10,11} the Jump-Walking method,⁵² and the recently proposed multicanonical Jump-Walking method.⁵³

Finally, it would be useful to test QTA-PIMC on harder optimization problems. In this paper, we have tested the efficiency of QTA-PIMC, but not its robustness. It would be interesting to see how QTA-PIMC performs on potential energy surfaces where most optimization methods have failed, for example, Lennard-Jones clusters consisting of $N = 38, 75-77, 102-104$ atoms,⁵ or low-autocorrelation binary sequences (where the PES has a “golf-course” character).⁵⁴⁻⁵⁷ With our current QTA-PIMC annealing scheme, we have located the global minimum of the 38-atom Lennard-Jones cluster; investigation of the other difficult cases of Lennard-Jones clusters are currently being carried out. For these potentially complex energy surfaces, a larger number of Trotter beads than was used here might be needed. Consequently, more sophisticated sampling techniques from the PIMC arsenal may be required.^{31,32,35}

Acknowledgment. Y.L. thanks Dr. J. Daniel Gezelter for useful discussions. This work has been supported by a grant from the NIH (5-RO1-GM43340). The computational resources of the NIH Biotechnology Resource Center for Biomolecular Simulations at Columbia University (RR-06892) are gratefully acknowledged.

References and Notes

(1) Kirkpatrick, S.; Gelatt, C. D., Jr.; Vecchi, M. P. *Science* **1983**, *220*, 671.

- (2) Holland, J. *SIAM J. Comput.* **1973**, *2*, 88.
 (3) Davis, L., Ed. *Handbook of Genetic Algorithms*; van Nostrand Reinhold: New York, 1991.
 (4) Li, Z.; Scheraga, H. A. *Proc. Natl. Acad. Sci. U.S.A.* **1987**, *84*, 6611.
 (5) Wales, D. J.; Doye, J. P. K. *J. Phys. Chem. A* **1997**, *101*, 5111.
 (6) Maranas, C. D.; Floudas, C. A. *J. Chem. Phys.* **1994**, *100*, 1247.
 (7) Maranas, C. D.; Floudas, C. A. *J. Global Optim.* **1994**, *4*, 135.
 (8) Standley, D. M.; Eyrych, V. A.; Felts, A. K.; Friesner, R. A.; McDermott, A. E. *J. Mol. Biol.* **1999**, *285*, 1691.
 (9) Eyrych, V. A.; Standley, D. M.; Felts, A. K.; Friesner, R. A. *Proteins: Struct. Funct. Genet.* **1999**, *35*, 41.
 (10) Berg, B. A.; Neuhaus, T. *Phys. Lett. B* **1991**, *267*, 249.
 (11) Hansmann, U. H. E.; Okamoto, Y. *Physica A* **1994**, *212*, 415.
 (12) Stillinger, F. H. *Phys. Rev. B* **1985**, *32*, 3134.
 (13) Stillinger, F. H.; Weber, T. A. *J. Stat. Phys.* **1988**, *52*, 1429.
 (14) Piel, L.; Kostrowicki, J.; Scheraga, H. A. *J. Phys. Chem.* **1989**, *93*, 3339.
 (15) Li, Z.; Scheraga, H. A. *Proc. Natl. Acad. Sci. U.S.A.* **1987**, *84*, 15.
 (16) Ma, J.; Straub, J. E. *J. Chem. Phys.* **1994**, *101*, 533.
 (17) Amara, P.; Straub, J. E. *J. Phys. Chem.* **1995**, *99*, 14840.
 (18) Shalloway, D. *Global Optim.* **1992**, *2*, 281.
 (19) Shalloway, D. Variable-scale coarse-graining in macromolecular global optimization. In *Large-Scale Optimization with Applications, Part III: Molecular Structure and Optimization*; Biegler, L. T., Coleman, T. F., Conn, A. R., Santosa, F. N., Eds.; Springer Verlag: New York, 1997.
 (20) Amara, P.; Hsu, D.; Straub, J. E. *J. Phys. Chem.* **1993**, *97*, 6715.
 (21) Finnila, A. B.; Gomez, M. A.; Sebenik, C.; Stenson, C.; Doll, J. D. *Chem. Phys. Lett.* **1994**, *219*, 343.
 (22) Kadowaki, T.; Nishimori, H. *Phys. Rev. E* **1998**, *58*, 5355.
 (23) Brooke, J.; Bitko, D.; Rosenbaum, T. F.; Aeppli, G. *Science* **1999**, *284*, 779.
 (24) Honeycutt, J. D.; Thirumalai, D. *Proc. Natl. Acad. Sci. U.S.A.* **1990**, *87*, 3526.
 (25) Honeycutt, J. D.; Thirumalai, D. *Biopolymers* **1992**, *32*, 695.
 (26) Nymeyer, H.; García, A. E.; Onuchic, J. N. *Proc. Natl. Acad. Sci. U.S.A.* **1998**, *95*, 5921.
 (27) Guo, Z.; Thirumalai, D.; Honeycutt, J. D. *J. Chem. Phys.* **1992**, *97*, 525.
 (28) Guo, Z.; Thirumalai, D. *Biopolymers* **1994**, *35*, 137.
 (29) Feynman, R. P.; Hibbs, A. R. *Quantum mechanics and path integrals*; McGraw-Hill: New York, 1965.
 (30) Feynman, R. P. *Statistical mechanics*; Benjamin: New York, 1972.
 (31) Berne, B. J.; Thirumalai, D. *Annu. Rev. Phys. Chem.* **1986**, *37*, 401.
 (32) Ceperley, D. M. *Rev. Mod. Phys.* **1995**, *67*, 279.
 (33) Jordon, H. F.; Fosdick, L. D. *Phys. Rev.* **1968**, *171*, 128.
 (34) Barker, J. A. *J. Chem. Phys.* **1979**, *70*, 2914.
 (35) Tuckerman, M. E.; Berne, B. J.; Martyna, G. J.; Klein, M. L. *J. Chem. Phys.* **1993**, *99*, 2796.
 (36) Shankar, R. *Principles of quantum mechanics*, 2nd ed.; Plenum Press: New York, 1994.
 (37) Chandler, D.; Wolynes, P. G. *J. Chem. Phys.* **1981**, *74*, 4078.
 (38) Scheraga, H. A. *Biophys. Chem.* **1996**, *59*, 329.
 (39) Purisima, E. O.; Scheraga, H. A. *Proc. Natl. Acad. Sci. U.S.A.* **1986**, *83*, 2782.
 (40) Purisima, E. O.; Scheraga, H. A. *J. Mol. Biol.* **1987**, *196*, 697.
 (41) Veitshans, T.; Klimov, D.; Thirumalai, D. *Folding Des.* **1996**, *2*, 1.
 (42) Herman, M. F.; Bruskin, E. J.; Berne, B. J. *J. Chem. Phys.* **1982**, *76*, 5150.
 (43) Thirumalai, D.; Berne, B. J. *J. Chem. Phys.* **1983**, *79*, 5029.
 (44) Doll, J. D.; Freeman, D. L. *J. Chem. Phys.* **1984**, *80*, 2239.
 (45) Freeman, D. L.; Doll, J. D. *J. Chem. Phys.* **1984**, *80*, 5709.
 (46) Pollock, E. L.; Ceperley, D. M. *Phys. Rev. B* **1984**, *30*, 2555.
 (47) Raedt, B. D.; Sprik, M.; Klein, M. L. *J. Chem. Phys.* **1984**, *80*, 5719.
 (48) Parrinello, M.; Rahman, A. *J. Chem. Phys.* **1984**, *80*, 860.
 (49) Hetényi, B.; Müser, M. H.; Berne, B. J. *Phys. Rev. Lett.*, submitted for publication.
 (50) Liu, Z.; Berne, B. J. *J. Chem. Phys.* **1993**, *99*, 6071.
 (51) Zhou, R.; Berne, B. J. *J. Chem. Phys.* **1997**, *107*, 9185.
 (52) Frantz, D. D.; Freeman, D. L.; Doll, J. D. *J. Chem. Phys.* **1990**, *93*, 2769.
 (53) Xu, H.; Berne, B. J. *J. Chem. Phys.* **1999**, *110*, 10299.
 (54) Bernasconi, J. *J. Physique* **1987**, *48*, 559.
 (55) Krauth, W.; Mezard, M. Z. *Phys. B* **1995**, *97*, 127.
 (56) Dittes, F.-M. *Phys. Rev. Lett.* **1996**, *76*, 4651.
 (57) Wenzel, W.; Hamacher, K. *Phys. Rev. Lett.* **1999**, *82*, 3003.