

Quantum Thermal Annealing with Renormalization: Application to a Frustrated Model Protein

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: July 20, 2000; In Final Form: October 12, 2000

A renormalization approach is introduced in quantum thermal annealing. This new global optimization algorithm (QTAR) is applied to a highly frustrated BLN model protein with 46 residues. The underlying quantum sampling algorithms utilized are primitive and staging path integral Monte Carlo. The current method enables us to achieve significant improvement in the success rate for locating the global minimum of this protein, while using much less computational effort, compared to previous schemes used. Further applications and possible enhancements of the algorithm are also discussed.

I. Introduction

In many fields of science and technology, global optimization represents an important yet often very difficult challenge. This is especially true of the so-called NP-complete problems, where the number of local minima for the underlying system increases exponentially with its size. The protein folding and traveling salesman problems are well-known examples of such cases. It is impossible to exhaustively search for the global minimum of these systems once they reach a large enough size. Instead, one has to rely on existing optimization methods and hope that they will be able to sample these huge search spaces efficiently and reliably.

Many good optimization methods exist, for example the widely used simulated annealing algorithm,¹ pure quantum annealing methods,^{2–5} quantum thermal annealing,^{6,7} (a superset of the aforementioned methods), various potential smoothing^{8,9} and classical density annealing^{10–14} techniques, Monte Carlo minimization,^{15,16} multicanonical algorithms,^{17–19} dimensional strategies,^{20–24} and others. However, the presence of a multitude of energy scales on the energy landscapes of certain classes of systems might present difficulties for some optimization algorithms in a practical sense. A search method might get stuck in a metastable energy basin and subsequently be unable to overcome energy barriers to explore other parts of the energy landscape. An even more dire scenario is the presence of a huge number of similar local minima which have energy values very similar to that of the global minimum but are separated by energy barriers of disparate scales. Most optimization algorithms would have difficulty picking out the correct global minimum readily out of the many comparable local minima in these so-called frustrated systems. In this paper, we investigate the use of a new algorithm, quantum thermal annealing with renormalization (QTAR), to effectively overcome the aforementioned issues in global optimization.

It has been shown computationally that quantization of a system moving on a rough energy landscape softens the potential, thereby hastening convergence to the global minimum in quantum annealing.⁶ In fact, quantum annealing has been

achieved experimentally and it proves to be much more efficient than classical thermal annealing in the determination of the ground state of a disordered magnet.²⁵ However, to be an effective computational tool, one needs to address the overhead associated with the additional degrees of freedom that are introduced when a classical system is quantized. This is especially apparent in quantum thermal annealing with a path integral Monte Carlo approach. In this study, we discuss how one can systematically reduce the number of degrees of freedom (and hence the computational cost) as one anneals the system from the quantum regime back to the classical realm, through the use of a *renormalization* approach. In addition to cost saving, this approach also allows the isomorphic classical system to sample a hierarchy of energy and length scales during the search for its global minimum. The particle-based character of the path integral Monte Carlo method makes renormalization a natural addition to our quantum thermal annealing scheme. The current QTAR method contrasts with our previous QTA schedule,⁶ where we annealed the system quantum mechanically by decreasing the value of \hbar within a quantum-to-classical cycle. In the present implementation, we anneal the system by decimating P , the number of Trotter time slices as well. This leads to an efficient systematic reduction in the number of degrees of freedom.

The 46-residue BLN model protein of Honeycutt and Thirumalai^{26,27} is an interesting and challenging case for global optimization methods. It has been shown through thermodynamics and kinetics²⁸ that this 46-mer is a highly frustrated system. This fact is confirmed by further studies of its potential energy surface (PES) through the use of disconnectivity graph analysis and examination of pathways on the energy landscape.²⁹ Both studies show that there is no single dominant funnel on the PES that would lead the 46-mer to its global minimum consistently. In fact, the PES is a very rugged, “glassy” one which is made up of many minima with energies close to that of the global minimum. These minima are separated by a diverse range of energy barriers, which makes the global minimum particularly difficult to locate. Structurally, this corresponds to having many general β -barrel-like low energy structures, out of which only one is the true global minimum (strictly speaking there are two degenerate global minima here because the Hamiltonian of the

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

BLN model is invariant under inversion $\mathbf{r}_i \rightarrow -\mathbf{r}_i$). The ability for any one of these structures to transform to another is greatly hindered by the need for cooperative and collective rearrangements of the residues.

In this paper, we shall be investigating the frustrated BLN 46-mer. Previous studies^{6,19} have shown that a standard method like simulated annealing (SA) is very ineffective in dealing with such a system. Its success rates in locating the correct lowest energy structure of the protein are discouragingly small even though long simulation times are used. In the following, we present the QTAR algorithm, and apply it to the BLN protein. Simulation studies with this new method show that it is able to efficiently and very reliably locate the global minimum of this highly frustrated system.

II. Method

In Feynman's path integral formulation of quantum statistical mechanics,³⁰ the quantum canonical partition function is written in terms of a path integral. For the formalism to be computationally useful, one could discretize the path integral in various ways.^{31,32} Two such discrete expressions are results of the primitive approximation^{33,34} and staging transformation.^{35,36} The partition function with the primitive Hamiltonian for a system of size N in three dimensions is

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

where P , the Trotter number, is an integer that denotes the number of "time" slices used in the discretization, and $\omega_p \equiv (\beta\hbar)^{-1}\sqrt{P}$. Accurate treatment of a highly quantum system requires a large value for P , while a purely classical system has $P = 1$. Hence, P is a measure of the "quantumness" of a system. In the above equation, $\mathbf{r}_{i,t}$ is the 3-vector position of the i -th particle in the t -th time slice, and $V_{\text{cl}}(\{\mathbf{r}_i\}; t)$ represents the total classical potential energy evaluated at time slice t . For a strongly quantum system where P is large, one could sample its equilibrium properties more efficiently with the staging Hamiltonian. The corresponding partition function³⁶ is

$$Q_P^{\text{stag}}(\beta) = \left\{ \frac{\beta m \omega_j^2}{2\pi} \prod_{k=2}^j \left(\frac{\beta m_k \omega_p^2}{2\pi} \right) \right\}^{3Nn/2} \int d\mathbf{u}_{1,1} \dots d\mathbf{u}_{i,t} \dots d\mathbf{u}_{N,P} \exp \left[-\beta \left(\sum_{i=1}^N \sum_{s=0}^{n-1} \frac{1}{2} m\omega_j^2 |\mathbf{u}_{i,sj+1} - \mathbf{u}_{i,(s+1)j+1}|^2 + \sum_{i=1}^N \sum_{s=0}^{n-1} \sum_{k=2}^j \frac{1}{2} m_k \omega_p^2 \mathbf{u}_{i,sj+k}^2 + \frac{1}{P} \sum_{t=1}^P V_{\text{cl}}(\{\mathbf{r}_i(\mathbf{u})\}; t) \right) \right] \quad (2)$$

with $nj = P$, where n and j are the number of end-point and staging Trotter beads, respectively. The staging coordinates are $\mathbf{u}_{i,t}$, with $m_k = mk/(k-1)$, and $\omega_j \equiv (\beta\hbar)^{-1}\sqrt{P/j}$.

For the present scheme, quantum thermal annealing is achieved by methodically reducing both P and \hbar . In each quantum-to-classical cycle, we wish to systematically remove half of the total number of Trotter time slices of the primitive and/or staging Hamiltonians in stages until we reach $P = 1$ (classical regime):

$$P_0 \rightarrow \frac{P_0}{2^1} \rightarrow \frac{P_0}{2^2} \rightarrow \frac{P_0}{2^3} \rightarrow \dots \rightarrow 2 \rightarrow 1 \quad (3)$$

where $P_0 = 2^\alpha$ ($\alpha \geq 1$) is the initial number of Trotter time slices used. The reduction in P is accomplished through the use of a *renormalization* approach for both Hamiltonians. Between renormalizations, in each P stage where the number of Trotter beads is held constant, the system is allowed to explore configuration space via PIMC moves. Different types of renormalization schemes are possible. The one used here is that due to Migdal and Kadanoff (MK).³⁷⁻³⁹ We choose the MK approach because it provides a simple way to incorporate renormalization in QTAR. First, MK bond moving operations are performed whereby all bonds representing $V_{\text{cl}}(\{\mathbf{r}_i\}; t)$ with odd-numbered Trotter time slices are moved to their *adjacent* even-numbered sites (the designation of odd and even is arbitrary). Consequently, instead of having $V_{\text{cl}}(\{\mathbf{r}_i\}; t = a + 1)$ at a particular time slice $t = a + 1$ (which is even), we now have 2 sets of bonds $V_{\text{cl}}(\{\mathbf{r}_i\}; t = a + 1) + V_{\text{cl}}(\{\mathbf{r}_i\}; t = a)$ at $t = a + 1$. Unlike lattice systems (e.g., the Ising model) for which the MK transformation was originally used, the BLN 46-mer (and other chemical and biological molecules of interest here) is off-lattice. For such systems, the bond-moving operations do not in general result in bonds that fall exactly on top of their targets. However, since the configurations corresponding to time slices $t = a + 1$ and $t = a$ are adjacent, they are expected to be quite similar to each other in terms of configuration and thus energy. As a result, we take

$$\begin{aligned} & V_{\text{cl}}(\{\mathbf{r}_i\}; t = a + 1) + V_{\text{cl}}(\{\mathbf{r}_i\}; t = a) \\ & \approx V_{\text{cl}}(\{\mathbf{r}_i\}; t = a + 1) + V_{\text{cl}}(\{\mathbf{r}_i\}; t = a + 1) \\ & = 2V_{\text{cl}}(\{\mathbf{r}_i\}; t = a + 1) \end{aligned} \quad (4)$$

If more than one set of bonds (say p of them) are moved instead of just the nearest-neighbor set, we expect the approximation above to be less valid because it is not likely that all p adjacent configurations would be similar to one another. In this case, one might consider using other renormalization schemes such as those that involve potential averaging. However, that would add additional computational costs to the scheme. Hence, the MK approach serves us better for quantum thermal annealing purposes. Upon performing the MK bond-moving operations, all the odd-numbered Trotter beads are now free from the external potential $V_{\text{cl}}(\{\mathbf{r}_i\}; t)$. These odd-numbered Trotter beads can now be integrated (decimated) out. The end result is rather simple. For the primitive Hamiltonian, the functional form remains the same as in eq 1, but with P replaced by $P' = P/2$. The MK renormalization of the staging Hamiltonian,⁷ while more involved, gives an analogous result. Its functional form is also conserved as in eq 2, but with P replaced by $P' = P/2$, n by n' , and j by j' , with the condition that $n'j' = P'$. Hence, upon renormalization, we have half as many degrees of freedom as before in both the primitive and staging Hamiltonians. This process could then be repeated until the classical regime ($P = 1$) is reached, as in eq 3.

In addition to P , the other parameter that controls the degree of "quantumness" of the system is $\hbar = h/2\pi$, where h is Planck's constant. For quantum annealing purposes, \hbar is used as an adjustable *parameter*. By suitably controlling this parameter, or more generally, the force constant $k_p \equiv m\omega_p^2$, we could adjust how free or bounded each Trotter bead is with respect to

their neighbors in the same quantum polymer chain. This allows us to control the size of the quantum clouds during annealing.

III. Implementation

Writing eqs 1 and 2 in the form of a classical configurational integral,

$$Z_{\text{cl}} = \left[\prod_{i=1}^N \prod_{t=1}^P \int d\mathbf{r}_{i,t} \right] \exp\left(-\beta_t \Phi(\{\mathbf{r}_{i,t}\}; \beta_q, \hbar)\right) \quad (5)$$

one can extract the following effective potentials

$$\Phi_{\text{prim}} = \sum_{i=1}^N \sum_{t=1}^P \frac{1}{2} m \omega_p^2(\beta_q) |\mathbf{r}_{i,t} - \mathbf{r}_{i,t+1}|^2 + \frac{1}{P} \sum_{t=1}^P V_{\text{cl}}(\{\mathbf{r}_i\}; t) \quad (6)$$

and

$$\Phi_{\text{stag}} = \sum_{i=1}^N \sum_{s=0}^{n-1} \frac{1}{2} m \omega_j^2(\beta_q) |\mathbf{u}_{i,sj+1} - \mathbf{u}_{i,(s+1)j+1}|^2 + \sum_{i=1}^N \sum_{s=0}^{n-1} \sum_{k=2}^j \frac{1}{2} m_k \omega_p^2(\beta_q) \mathbf{u}_{i,sj+k}^2 + \frac{1}{P} \sum_{t=1}^P V_{\text{cl}}(\{\mathbf{r}_i(\mathbf{u})\}; t) \quad (7)$$

Equations 6 and 7 are associated with the primitive and staging Hamiltonians, respectively. Sampling these effective potentials with the Monte Carlo algorithm constitutes the path integral Monte Carlo approach. In general, one samples Φ_{prim} using *local* and *global* MC moves and Φ_{stag} with *staging* MC moves. For each local move, a single Trotter bead on each residue is displaced randomly within a cube of width Δ_{local} and the Metropolis criterion is then used to determine the acceptance or rejection of this trial move. Similarly, a global and staging move is made up of trial displacements of the entire Trotter chain of P beads and a partial chain segment of j beads of a residue, respectively.^{6,36} In the series of equations above, we have introduced two different β 's, specifically β_t and β_q . The former is the true inverse thermal temperature, while the latter is a factor in the quantum annealing parameters $\omega_p(\beta_q)$ and $\omega_j(\beta_q)$. This differentiation allows one to perform either quantum or thermal annealing individually or both types of annealing concurrently. In what follows, we define one Monte Carlo sweep (MCS) to be the attempted displacements of N particles, where N is the system size. The attempted moves of all $N \times P$ Trotter beads once is called a PIMC pass. There is a slight overhead associated with the calculation of the harmonic potential when making local moves with the primitive Hamiltonian or doing the coordinate transformation in staging PIMC. However, this is not significant compared to the dominant calculational demands of the nonbonded interaction potentials which scale as N^2 . The actual CPU overhead required will be presented later in the paper. Therefore, one PIMC pass of either local, global, or staging moves involves approximately P MC sweeps.

A brief description of the algorithm follows (details of the implementation will be given later):

1. Generate an initial random configuration for the system.
2. Set the initial value of the Trotter number P to P_0 , k_p to k_{P_0} , and the thermal temperature T to T_0 . Quantize the classical configuration by going to step 3.
3. Perform n_{staging} staging PIMC passes with eq 7. If P gets relatively smaller (typically ≤ 64 here), perform n_{local} local and n_{global} global PIMC passes with eq 6 instead. After each PIMC pass, T is reduced linearly by ΔT and k_p is increased linearly by Δk_p .

4. Reduce P to $P' = P/2$ through renormalization. The number of PIMC passes for this new P' -stage is doubled so that the total number of MC sweeps used in each P -stage remains constant. Go back to step 3. This process is repeated until $P = 1$, when an *intermediate* classical configuration is obtained.

5. Terminate if the stop criterion is met. If not, repeat the quantum thermal annealing process with the intermediate classical configuration by going to step 2.

Steps 2–4 make up a QTAR *cycle*. A graphical depiction of the above algorithm is given in Figure 1.

The 46-residue BLN model protein studied here has the sequence B₉N₃(LB)₄N₃B₉N₃(LB)₅L, where the letter codes B, L, and N represent hydrophobic, hydrophilic and neutral residues, respectively. The potential energy of the BLN system is given by

$$V_{\text{cl}}(\{\mathbf{r}_i\}) = \sum_{i=1}^{N-1} \frac{k_r}{2} (|\mathbf{r}_{i+1} - \mathbf{r}_i| - a)^2 + \sum_{i=1}^{N-2} \frac{k_\theta}{2} (\theta_i - \theta_0)^2 + \sum_{i=1}^{N-3} [A_i(1 + \cos \phi_i) + B_i(1 + \cos 3\phi_i)] + 4\epsilon \sum_{i=1}^{N-3} \sum_{j=i+3}^N C_{ij} \left[\left(\frac{\sigma}{r_{ij}} \right)^{12} - D_{ij} \left(\frac{\sigma}{r_{ij}} \right)^6 \right] \quad (8)$$

where

$$C_{ij} = 1, D_{ij} = 1 \quad \text{if } i, j \in B$$

$$C_{ij} = \frac{2}{3}, D_{ij} = -1 \quad \text{if } i \in L, j \in B, L$$

$$C_{ij} = 1, D_{ij} = 0 \quad \text{if } i \in N, j \in B, L, N$$

In eq 8 above, N is the number of residues and $r_{ij} = |\mathbf{r}_i - \mathbf{r}_j|$. All physical quantities are in reduced units. In our simulations, the mass of each residue m , the bond length a , the energy constant ϵ , the Lennard-Jones parameter σ , and the Boltzmann constant k_B are set to unity. The other constants are $k_r = 400\epsilon/a^2$, $k_\theta = 20\epsilon/(\text{rad})^2$, and $\theta_0 = 1.8326$ rad. For the dihedral-angle potential term, if two or more of the four defining residues of ϕ_i are neutral (N), then $A_i = 0\epsilon$, $B_i = 0.2\epsilon$, otherwise $A_i = B_i = 1.2\epsilon$. Following previous studies,^{6,13} we also utilize a weak boundary potential $V_{bp}(\{\mathbf{r}_i\}) = \sum_{i=1}^N (k_b/2) |\mathbf{r}_i - \mathbf{r}_{\text{com}}|^2$ to prevent the protein from dissociation and also to encourage folding. \mathbf{r}_{com} is the center of mass of the protein chain and k_b is reduced from 0.05 to 0.005 in each QTAR cycle.

IV. Results

We conduct 20 simulation trials for QTAR. Each run starts with a random extended protein configuration as in ref 6. Every residue of the starting configuration of the 46-mer is endowed with $P_0 = 256$ Trotter beads. The system is then “quantized” through the use of staging moves, which are more efficient than either local or global moves for large P . The number of staging passes used is $n_{\text{staging}} = 15$. At the end of these moves, renormalization of Trotter time slices is carried out to leave us with $P = 128$. The number of staging passes is now doubled to 30 such that the total number of MC sweeps used in each P stage remains constant. For $P = 256$, staging moves of length $j = 24$ are used, and the Gaussian widths of such moves are approximately 0.12. For $P = 128$, we use $j = 8$, and the widths of staging moves are 0.08. That the use of a large P_0 together with the staging Hamiltonian are effective for “quantization”

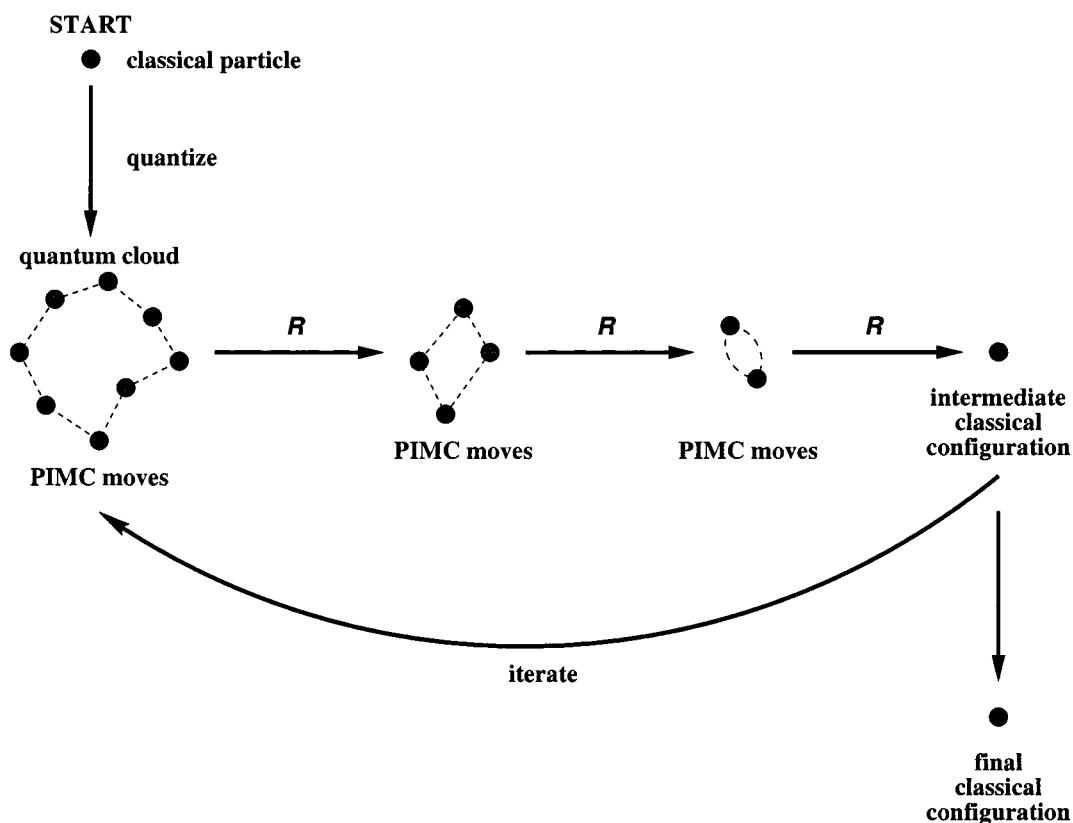


Figure 1. The QTAR (quantum thermal annealing with renormalization) algorithm illustrated schematically with a single classical particle. The classical particle is repeatedly quantized and annealed back to the classical regime using PIMC moves and the Migdal-Kadanoff renormalization operation R . At the end of the quantum thermal annealing cycles, the global minimum of the classical system would be obtained. Note that for illustration purposes, P_0 is set to 8 in this diagram. The actual P_0 used in our simulations is 256.

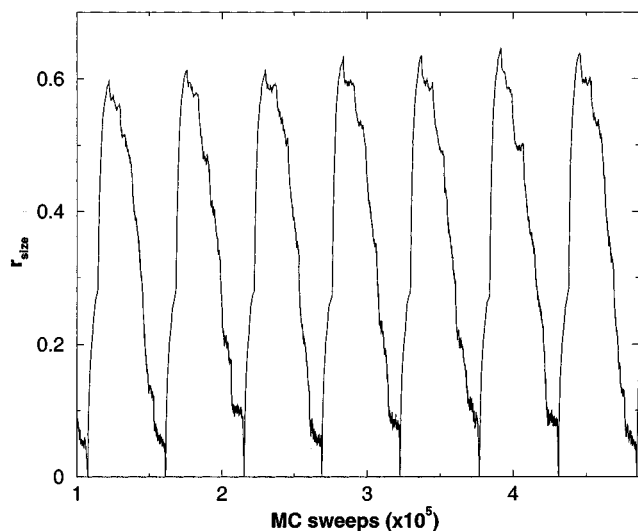


Figure 2. A snapshot of the root-mean-square size (r_{size}) of the classical ring polymers representing the quantum clouds of the residues versus MC sweeps, for 7 QTAR cycles. The use of a relatively large number of Trotter beads ($P_0 = 256$) at the beginning of each cycle enables the system to be “quantized” efficiently and facilitates strong tunneling events, resulting in better searches on the PES.

is evident in the rapid rise in the average size of the quantum clouds at the beginning of each QTAR cycle, as shown in Figure 2. For $P \leq 64$, local and global moves with the primitive Hamiltonian are used instead of staging moves. At these P values, j drops too low if a reasonable acceptance ratio is to be maintained for the set of simulation parameters used here. This is not unexpected since the staging algorithm has been designed for strongly quantum systems where P is huge, typically in the

hundreds or more. Hence, it is very useful at the beginning of a QTAR cycle. The number of MC sweeps for these local and global moves in each P stage is also kept constant by doubling n_{local} and n_{global} every time P is halved. For these moves, $k_{P_0} = 0.24$ and $k_{P_f} = 1.0$. The thermal temperature is annealed linearly from $T_0 = 0.2$ to $T_f = 0.02$ in each QTAR cycle. PIMC moves and renormalization are thus carried out repeatedly in the aforementioned manner until we reach $P = 1$, where the quantum cloud of each residue would collapse back to a point, signified by $r_{size} = 0$ in Figure 2. The intermediate classical configuration so obtained is subjected to a conjugate gradient refinement of its energy. The above completes one QTAR cycle. This whole process is iterated until the currently known global minimum of the 46-mer is located, which is the stop criterion for the QTAR schedule here. Figure 3 plots the potential energy attained at the end of each QTAR cycle versus MC sweeps for a typical run. Finally, we note that 10 QTAR cycles are used during each trial for equilibration purposes whereby the number of staging beads $\{j\}_{P_i}$, the set of maximum displacements at different values of P for local moves $\{\Delta_{local}\}_{P_i}$, and the corresponding displacements for global moves $\{\Delta_{global}\}_{P_i}$, are adjusted such that the acceptance ratio for each type of move is kept approximately constant at 40%.

The results of the simulation runs are presented in Figure 4. Alongside these, we place corresponding data obtained with SA from ref 6. In Figure 4a, we see that QTAR is able to locate the global minimum of the 46-mer with a perfect success rate of 100%, versus 9% for SA. How much computational effort is used to obtain either set of results? For SA, 32×10^6 MC sweeps are used per simulation trial. For QTAR, the average number of MC sweeps utilized in each run is 3.8×10^6 . This is shown in Figure 4b. However, 1 MC sweep for QTAR costs

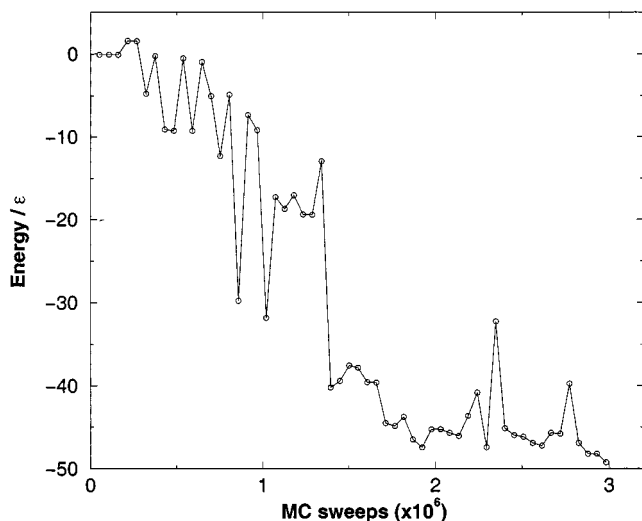


Figure 3. Potential energy of the BLN 46-mer versus MC sweeps. Each circle represents an intermediate classical configuration obtained at the end of a QTAR cycle.

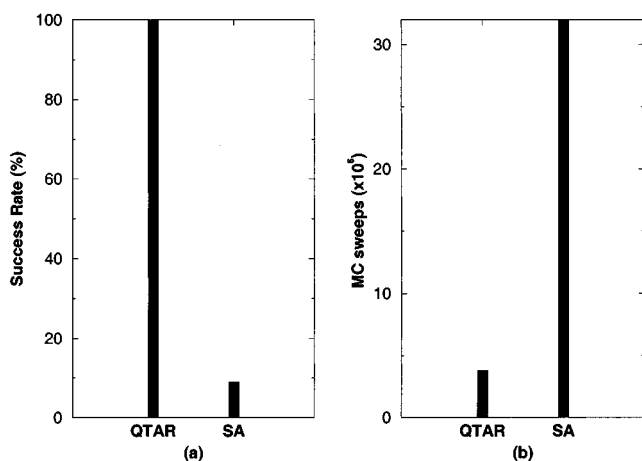


Figure 4. The results of global optimization of the 46-residue BLN protein with the QTAR algorithm, in comparison to those obtained with SA (simulated annealing). Results for QTAR are averaged over 20 trials, those for SA are from ref 6. (a) The success rates for locating the global minimum of the BLN 46-mer with QTAR and SA. (b) The average number of MC sweeps used to find the global minimum over 20 QTAR trials. For SA, the total number of MC sweeps used in each trial is shown.

approximately 13% more CPU time than that for SA, due to the extra overhead associated with using PIMC over regular MC mentioned earlier. Consequently, the amount of computational effort used for QTAR is 7.5 times less than that of SA, yet its success rate is 11.1 times higher than the latter.

As a further test, we conduct sets of 20 SA runs using 4.3×10^6 MC sweeps per run. If one takes into account the 13% CPU overhead associated with utilizing PIMC over regular MC, each set of 20 SA runs takes the same amount of CPU time as a set of 20 QTAR runs which uses an average of 3.8×10^6 MC sweeps per run. We perform 3 different sets of such 20 SA trials, each with a different linear annealing schedule. Here, we denote the number of temperature annealing steps as n_{dT} and the number of MC sweeps used at each temperature as n_T , such that the total number of MC sweeps used in each SA run is $n_{dT} \times n_T$. The 3 SA schedules have $n_{dT} \times n_T = 215 \times 20000$, 2150×2000 , and 21500×200 . They are labeled as SA1, SA2 and SA3, respectively, in Figure 5a, and SA(avg) represents the averaged results of SA1 to SA3. From the figure, it is clear

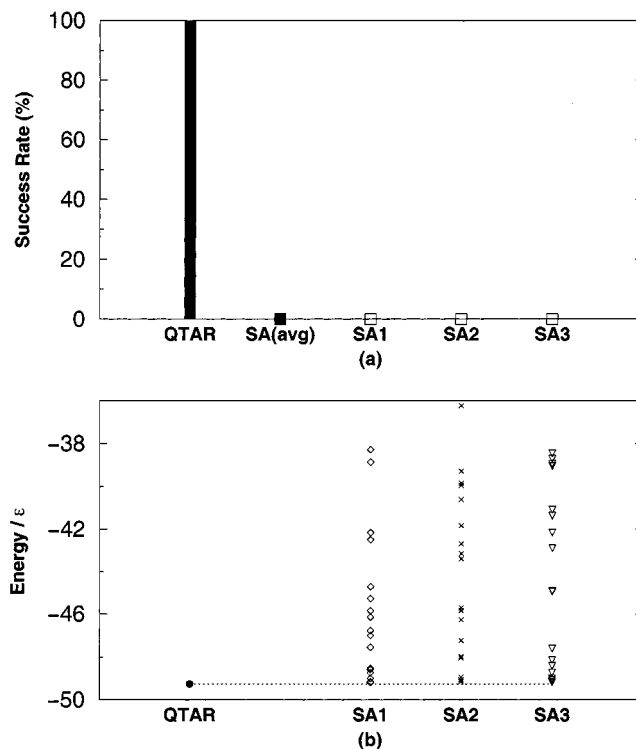


Figure 5. (a) The success rates of locating the global minimum of the 46-residue BLN protein with QTAR and SA. Each bar represents 20 independent simulation runs done using the same total amount of CPU time. SA(avg) is the average of SA1 to SA3. None of the SA schedules is able to locate the global minimum even once, while QTAR is able to achieve 100% success rate. (b) The corresponding minimum energies attained with the fore-mentioned QTAR and SA schedules. The dotted line represents the global minimum energy of the BLN 46-mer.

that none of these SA schedules is able to locate the global minimum of the 46-mer even once. In fact, SA is prone to getting stuck at metastable states with higher energies. This is evident from the large spread in the spectrum of lowest energy attained, as shown in Figure 5b. Another significant aspect of the aforementioned figure is the collection of local minima which have energies very close to that of the global minimum. This is a clear indication of the frustrated nature of the 46-mer. On the other hand, with the same computational effort, QTAR is able to achieve 100% success rate in locating the global minimum of this highly frustrated protein.

The present QTAR algorithm also represents an improvement over the previous QTA scheme.⁶ For example, in QTA, the global minimum was found 60% of the time out of a series of 20 simulation runs each utilizing 32×10^6 MC sweeps, whereas in QTAR, the global minimum is found 100% of the time using an average of 3.8×10^6 MC sweeps per trial over 20 trials. Thus, QTAR is at least 8 times more efficient than QTA. While the present results are very encouraging, it is not inconceivable that other QTAR schedules, together with systematic parameter optimization, might improve the results even further.

V. Discussion

For an unbiased global optimization algorithm to be successful, it has to take into account at least the following: (a) the size of the search space, (b) the ability to overcome potential barriers, (c) the ability to explore the search space in an effective and efficient manner. Although the configurational space for QTAR is of a higher dimension than the underlying classical system, simple arithmetic and simulation results show that the

use of renormalization overcomes the computational cost of the increased dimensionality effectively. One could also view the renormalization process as a systematic way for performing dimensional annealing. In fact, it has been shown that dimensional strategies do very well in optimization problems.^{20–24} QTAR finds its power in its ability to utilize both quantum and thermal effects in overcoming energy barriers efficiently on the PES. With the use of suitable parameters, there would be no barriers on the PES that the system could not overcome thermally and/or tunnel through quantum mechanically. In addition, QTAR is capable of generating higher dimensional *collective* moves (i.e., global and staging moves that displace whole and partial segments of quantum chains, respectively). This serves the need for collective rearrangements of protein configurations mentioned at the beginning of this article, albeit at higher dimensions. Hence, the QTAR algorithm is able to address all three criteria (a)–(c) mentioned above. Furthermore, it is possible to incorporate future improvements in any of the three categories. The current method is completely unbiased. However, for a general optimization problem, one would also want to make use of any physical insights into the problem, if available, by including system-specific techniques and heuristics. This would also help further increase the effectiveness of the algorithm.

VI. Conclusion

The present study demonstrates that QTAR consistently and efficiently locates the global minimum of a highly frustrated system, a BLN model protein with 46 residues, in completely unbiased fashion. A perfect success rate of 100% is attained with this new global optimization algorithm. This is achieved using significantly less computational effort than previous methods, including thermal annealing which is very ineffective for this system. In the future, we plan to test the QTAR algorithm on more realistic protein models, as well as other physical, chemical, biological or even abstract computational systems.

Acknowledgment. This work has been supported by the National Institutes of Health under Grants GM43340 and RR-06892.

References and Notes

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

- (2) Amara, P.; Hsu, D.; Straub, J. E. *J. Phys. Chem.* **1993**, *97*, 6715.
 (3) Doll, J. D.; Freeman, D. L. *IEEE Comput. Sci. Eng.* **1994**, *1*, 22.
 (4) Finnila, A. B.; Gomez, M. A.; Sebenik, C.; Stenson, C.; Doll, J. D. *Chem. Phys. Lett.* **1994**, *219*, 343.
 (5) Kadowaki, T.; Nishimori, H. *Phys. Rev. E* **1998**, *58*, 5355.
 (6) Lee, Y. H.; Berne, B. J. *J. Phys. Chem. A* **2000**, *104*, 86.
 (7) Lee, Y. H.; Berne, B. J. *Ann. Phys. (Leipzig)* **2000**, *9–10*, 668.
 (8) Stillinger, F. H. *Phys. Rev. B* **1985**, *32*, 3134.
 (9) Stillinger, F. H.; Weber, T. A. *J. Stat. Phys.* **1988**, *52*, 1429.
 (10) Li, Z.; Scheraga, H. A. *Proc. Natl. Acad. Sci. U.S.A.* **1987**, *84*, 15.
 (11) Piela, L.; Kostrowicki, J.; Scheraga, H. A. *J. Phys. Chem.* **1989**, *93*, 3339.
 (12) Ma, J.; Straub, J. E. *J. Chem. Phys.* **1994**, *101*, 533.
 (13) Amara, P.; Straub, J. E. *J. Phys. Chem.* **1995**, *99*, 14840.
 (14) Shalloway, D. *Global Optimization* **1992**, *2*, 281.
 (15) Li, Z.; Scheraga, H. A. *Proc. Natl. Acad. Sci. U.S.A.* **1987**, *84*, 6611.
 (16) Wales, D. J.; Doye, J. P. K. *J. Phys. Chem. A* **1997**, *101*, 5111.
 (17) Berg, B. A.; Neuhaus, T. *Phys. Lett. B* **1991**, *267*, 249.
 (18) Hansmann, U. H. E.; Okamoto, Y. *Phys. A* **1994**, *212*, 415.
 (19) Xu, H.; Berne, B. J. *J. Chem. Phys.* **2000**, *112*, 2701.
 (20) Purisima, E. O.; Scheraga, H. A. *Proc. Natl. Acad. Sci. U.S.A.* **1986**, *83*, 2782.
 (21) Purisima, E. O.; Scheraga, H. A. *J. Mol. Biol.* **1987**, *196*, 697.
 (22) Beutler, T. C.; van Gunsteren, W. F. *J. Chem. Phys.* **1994**, *101*, 1417.
 (23) Faken, D. B.; Voter, A. F.; Freeman, D. L.; Doll, J. D. *J. Phys. Chem. A* **1999**, *103*, 9521.
 (24) Stolovitzky, G.; Berne, B. J. *Proc. Natl. Acad. Sci. U.S.A.* **2000**, *97*, 11164.
 (25) Brooke, J.; Bitko, D.; Rosenbaum, T. F.; Aeppli, G. *Science* **1999**, *284*, 779.
 (26) Honeycutt, J. D.; Thirumalai, D. *Proc. Natl. Acad. Sci. U.S.A.* **1990**, *87*, 3526.
 (27) Honeycutt, J. D.; Thirumalai, D. *Biopolymers* **1992**, *32*, 695.
 (28) Nymeyer, H.; Garcia, A. E.; Onuchic, J. N. *Proc. Natl. Acad. Sci. U.S.A.* **1998**, *95*, 5921.
 (29) Miller, M. A.; Wales, D. J. *J. Chem. Phys.* **1999**, *111*, 6610.
 (30) Feynman, R. P.; Hibbs, A. R. *Quantum Mechanics and Path Integrals*; McGraw-Hill: New York, 1965.
 (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) Pollock, E. L.; Ceperley, D. M. *Phys. Rev. B* **1984**, *30*, 2555.
 (36) Tuckerman, M. E.; Berne, B. J.; Martyna, G. J.; Klein, M. L. *J. Chem. Phys.* **1993**, *99*, 2796.
 (37) Migdal, A. A. *Sov. Phys. JETP* **1975**, *42*, 413.
 (38) Migdal, A. A. *Sov. Phys. JETP* **1975**, *42*, 743.
 (39) Kadanoff, L. P. *Ann. Phys.* **1976**, *100*, 359.