PHYSICAL REVIEW
LETTERS

VOLUME 82 | 12 APRIL 1999 | NUMBER 15

Stochastic Tunneling Approach for Global Minimization of Complex Potential Energy Landscapes

W. Wenzel and K. Hamacher
Institut für Physik, Universität Dortmund, D-44221 Dortmund, Germany
(Received 2 November 1998)

We investigate a novel stochastic technique for the global optimization of complex potential energy surfaces that avoids the freezing problem of simulated annealing by allowing the dynamical process to tunnel energetically inaccessible regions of the potential energy surfaces by way of a dynamically adjusted nonlinear transformation of the original potential energy surfaces. We demonstrate the success of this approach, which is characterized by a single adjustable parameter, for three generic hard minimization problems. [S0031-9007(99)08916-4]

PACS numbers: 02.70.Lq, 02.50.Ey

The development of methods that efficiently determine the global minima of complex and rugged energy landscapes remains a challenging problem with applications in many scientific and technological areas. In particular, for *NP*-hard [1,2] problems, stochastic methods offer an acceptable compromise between the reliability of the method and its computational cost, which scales only as a power law with the number of variables [3] (for a fixed probability to locate the true minimum). In such techniques the global minimization is performed through the simulation of a dynamical process for a "particle" on the multidimensional potential energy surface. Widely used is the simulated annealing (SA) technique [4], where the potential energy surface (PES) is explored in a series of Monte Carlo (MC) simulations at successively decreasing temperatures. Its success often depends strongly on the choice of the cooling schedule, yet even the simplest geometric cooling schedule is characterized by three parameters (starting temperature, cooling rate and number of cooling steps) which must be optimized to obtain adequate results. For many difficult problems with rugged energy landscapes, SA suffers from the notorious "freezing" problem, because the escape rate from local minima diverges with decreasing temperature. To ameliorate this problem many variants of the original algorithm [5–7] have been proposed, often at the expense of introducing additional parameters.

In this Letter we investigate the stochastic tunneling method, a generic physically motivated generalization of SA. This approach circumvents the freezing problem, while reducing the number of problem-dependent parameters to one. We demonstrate the success of this approach for three hard minimization problems: the Coulomb spin glass (CSG), the traveling salesman problem (TSP), and the determination of low-autocorrelation binary sequences (LABS) in comparison with other techniques.

Method.—The freezing problem in stochastic minimization methods arises when the energy difference between "adjacent" local minima on the PES is much smaller than the energy of intervening transition states separating them. As an example consider the dynamics on the model potential in Fig. 1(a). At high temperatures a particle can still cross the barriers, but not differentiate between the wells. As the temperature drops, the particle will eventually become trapped with almost equal probability in any of the wells, failing to resolve the energy difference between them. The physical idea behind the stochastic tunneling method (STUN) is to allow the particle to "tunnel" [8] forbidden regions of the PES, once it has been determined that they are irrelevant for the low-energy properties of the problem. This can be accomplished by applying a nonlinear transformation to the PES:

$$f_{\mathrm{STUN}}(x) = 1 - \exp[-\gamma(f(x) - f_0)], \qquad (1)$$

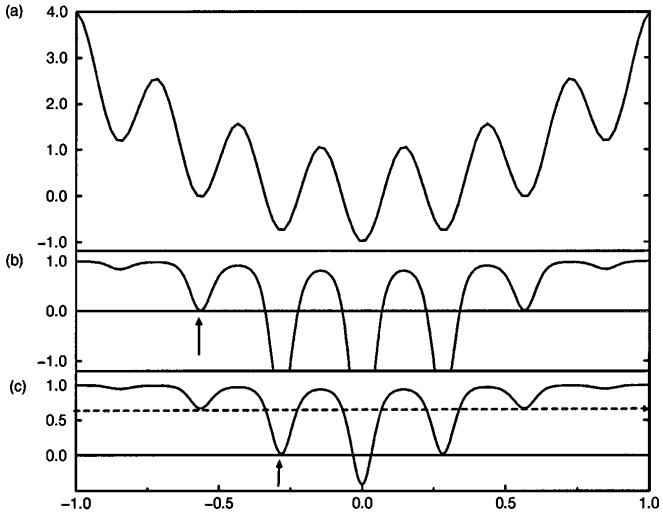

FIG. 1. (a) Schematic one-dimensional PES and (b) STUN effective potential, where the minimum indicated by the arrow is the best minimum found so far. All wells that lie above the best minimum found are suppressed. If the dynamical process can escape the well around the current ground-state estimate, it will not be trapped by local minima that are higher in energy. Wells with deeper minima are preserved and enhanced. (c) After the next minimum has been located, wells that were still pronounced in (b) are also suppressed. Once the true ground state has been found (not shown), all other wells have been suppressed and will no longer trap the dynamical process. The dotted line in (c) illustrates an energy threshold $0 < f_t < 1$ to classify the nature of the dynamics. Adjusting the temperature to maintain a particular average effective energy balances the tunneling and the local-search phases of the algorithm.

where f_0 is the lowest minimum encountered by the dynamical process thus far [see Figs. 1(b) and 1(c)] [9]. The effective potential preserves the locations of all minima, but maps the entire energy space from f_0 to the maximum of the potential onto the interval $[0, 1]$. At a given finite temperature of O(1), the dynamical process can therefore pass through energy barriers of arbitrary height, while the low-energy region is still well resolved. The degree of steepness of the cutoff of the high-energy regions is controlled by the tunneling parameter $\gamma > 0$. Continuously adjusting the reference energy f_0 to the best energy found so far successively eliminates irrelevant features of the PES which would trap the dynamical process.

To illustrate the physical content of the transformation we consider a MC process at some fixed inverse temperature β on the STUN-PES. An MC step from x_1 to x_2 with $\Delta = f(x_2) - f(x_1)$ is accepted with probability $\tilde{w}_{1\to 2} = \exp[-\beta(f_{\mathrm{STUN}}(x_2) - f_{\mathrm{STUN}}(x_1))]$. In the limit $\gamma\Delta \ll 1$ this reduces to $\tilde{w}_{1\to 2} \approx \exp(-\tilde{\beta}\Delta)$ with an effective energy-dependent temperature $\tilde{\beta} = \beta\gamma e^{\gamma(f_0 - f(x_1))} \leq \beta\gamma$. The dynamical process on the STUN potential energy surface with fixed temperature can thus be interpreted as an MC process with an energy-dependent temperature on the original PES. In the latter process the temperature rises rapidly when the local energy is larger than f_0 and the particle diffuses (or tunnels) freely through potential barriers of arbitrary height. As better and better minima are found, ever larger portions of the high-energy part of the PES are flattened out. In analogy to the SA approach this behavior can be interpreted as a self-adjusting cooling schedule that is optimized as the simulation proceeds.

Since the transformation in Eq. (1) is bounded, it is possible to further simplify the method: On the fixed energy scale of the effective potential, one can distinguish between phases corresponding to a local search and "tunneling" phases by comparing f_{STUN} with some fixed problem-independent predefined threshold f_t [see Fig. 1(c)]. For the success of the method it is essential that the minimization process spends some time tunneling and some time searching at any stage of the minimization process. We therefore adjust the parameter β accordingly during the simulation: If a short-time moving average of f_{STUN} exceeds the threshold $f_{\mathrm{thresh}} \approx 0.03$, β is reduced by some fixed factor, otherwise it is increased. Following this prescription the method is characterized by the single problem-dependent parameter (γ).

Applications.—In order to test the performance of this algorithm we have investigated three families of complicated *NP*-hard minimization problems. For each problem we have determined either the exact ground-state energy or a good estimate thereof. We computed the average error of the various optimization methods as a function of the computational effort to determine the computational effort required to reach a prescribed accuracy.

(i) (CSG) The determination of low-energy configurations of glassy PES is a notoriously difficult problem. We have verified by direct comparison that the method converges quickly to the exact ground states [10] for two-dimensional short-range Ising spin glasses of linear dimension 10 to 30 with either discrete or Gaussian distributions of the coupling parameters. Next we turned to the more demanding problem of the Coulomb spin glass, where classical charges $\{s_i\}$ with $s_i = \pm 1$ are placed on fixed randomly chosen locations within the unit cube. The energy of the system,

$$E(\{s_i\}) = \sum_{ij}^{N} \frac{s_i s_j}{|\vec{r}_i - \vec{r}_j|}, \qquad (2)$$

is minimized as a function of the distribution of the $\{s_i\}$.

The results of grand-canonical simulations for ten replicas of $N = 100$ and $N = 500$ charges are shown in Fig. 2. We first conducted twenty very long STUN runs for each replica to determine upper bounds for the true ground-state energy. For the same charge distributions we then averaged the error of STUN, SA, simulated tempering (ST) [6], and parallel tempering (PT) [7] for twenty runs per replica as a function of the numerical effort. We found that the average STUN energy converged in

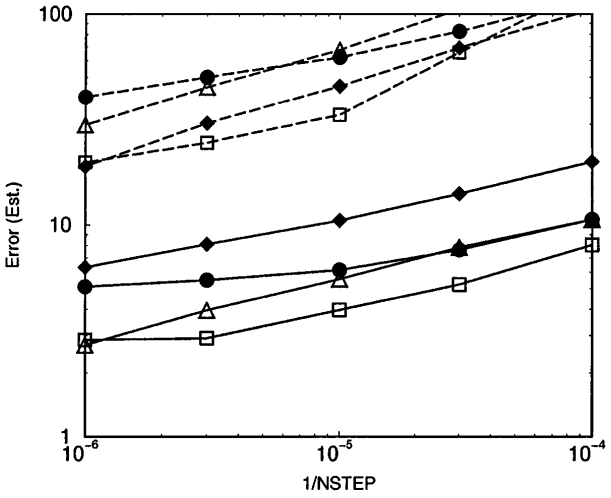

FIG. 2. Average estimated error for the ground-state estimates of the Coulomb glass using SA (circles), STUN (squares), ST (triangles), and PT (diamonds) for $N = 100$ (full lines in lower part) and $N = 500$ (dashed lines in upper part).

10^6 MC steps to within 1% of the estimated true ground-state energy. Fitting the curves in the figure with a power-law dependence we estimate that STUN is about 2 orders of magnitude more efficient than SA.

We found no consistent ranking of ST and PT relative to SA for the two system sizes considered. Both methods offer alternative routes to overcome the freezing problem in SA. In PT the configurations of concurrent simulations at a variety of temperatures are occasionally exchanged. In ST only a single simulation is undertaken, but its temperature is considered to be a dynamical variable. In both methods, a configuration can escape a local minimum when the instantaneous temperature is increased. The choice of the temperature set is system dependent and must be optimized much like the annealing schedule in SA. In accordance with other studies our results indicate that ST performs significantly better than SA for long simulation times, while PT was successful only for the larger system ($N = 500$). STUN converged faster than any of the competing methods, but showed a tendency to level off at high accuracy. In the limit of large computation, its accuracy was matched by ST for $N = 100$ and PT for $N = 500$.

(ii) (TSP) The "traveling salesman" problem is another ubiquitous *NP*-hard minimization problem [11,12]. We have investigated the problem in its simplest incarnation, i.e., as a minimization of the Euclidian distance along a closed path of N cities. Using long-range updates, i.e., the reversal and exchange of paths of arbitrary length, we found that both SA and STUN perform about equally well and reach the global optimum for $N = 20$, 50, and 100 very quickly (see right side of Table I).

Nevertheless, it is instructive to analyze this model somewhat further as it provides insight into the interplay of move construction and the complexity of the minimization problem. The unconstrained TSP is a rare instance among *NP*-hard minimization problems, where it is possible to construct efficient "long-range" hops on the PES.

TABLE I. Estimates for the optimal path length for the traveling salesman problem with $N = 20$, 50, and 100 sites using either only local (left side) or global (right side) moves as described in the text. For global moves, both SA and STUN are equally efficient to obtain low-energy paths. Using only local moves, the existence of barriers hampers the progress of SA. As a result, SA becomes less efficient than STUN. By virtue of its temperature exchange mechanism, PT also allows the random walk to cross the barriers, but is less efficient than STUN. The effort is given in thousands of steps; note that the evaluation of a local move is much less costly than that of a global move. The path length indicates the average optimal energy for 20 runs and the best energy found.

		Local Moves			Global Moves	
N	Effort	SA	PT	STUN	SA	STUN
20	50	4.85/3.55	4.35/3.55	3.60/3.55	3.94/3.61	3.55/3.55
20	100	4.52/3.58	4.02/3.55	3.62/3.55	3.93/3.55	
20	500	4.08/3.55	3.57/3.55	3.55/3.55	3.82/3.56	
20	1000	4.08/3.55	3.55/3.55			
20	5000	3.75/3.55				
50	100	12.5/10.61	13.72/12.58	11.06/9.39	5.74/5.65	5.72/5.65
50	500	11.0/8.68	11.55/10.65	8.32/5.83	5.70/5.65	5.67/5.65
50	1000	11.0/8.84	10.70/9.82	7.75/5.78	5.68/5.65	5.67/5.65
50	5000	9.84/8.10	8.99/7.89	7.16/5.78	5.66/5.65	5.65/5.65
50	10000	9.87/8.31		6.70/5.72	5.66/5.65	
100	200				8.42/8.11	8.40/8.01
100	500				8.18/8.01	8.18/7.97
100	1000				8.08/7.94	8.03/7.95
100	5000				8.01/7.94	8.01/7.96

In most practical applications of minimization problems related to the TSP, the construction of global moves is severely complicated by the existence of "hard constraints" on the routes taken. For such problems, as well as the other examples reported here, the alteration of just a few variables of the configurations leads to unacceptably high energies in almost all cases. As a result, the construction of global moves is not an efficient way to facilitate the escape from local minima. When only local moves, i.e., transpositions of two adjacent cities, are considered, high barriers, circumvented in the presence of global moves, hamper the progress of SA. The results on the left side of Table I demonstrate that in this scenario SA performs significantly worse than STUN.

(iii) (LABS) Finally, we turn to the construction of low-autocorrelation binary sequences [12,13]. The model can be cast as a ground-state problem for a one-dimensional classical spin-1/2 chain with long-range four-point interactions,

$$E = \frac{1}{N} \sum_{k=1}^{N-1} \left[\sum_{j=1}^{N-k} s_j s_{j+k} \right]^2, \tag{3}$$

and is one of the hardest discrete minimization problems known [14]. Even highly sophisticated and specialized optimization algorithms [12] have failed to find configurations anywhere near (within 20%) the ground-state energy that can be extrapolated from exact enumeration studies for small systems ($N < 50$) [15,16]. The reason for this difficulty has been attributed to the "golf-course" character of the energy landscape, and there is convincing evidence that SA will fail to converge to the ground-state energy even in the limit of adiabatic cooling [13]. The situation is significantly improved if the original potential energy surface is replaced by a piecewise constant energy surface that is obtained by a local minimization of the original PES at each point [17]. The latter surface preserves all ground-state configurations and energies of the original PES, but eliminates many "plateaus" of the golf-course landscape. By using the modified energy surface we are able to compare SA to STUN, since SA can now determine the ground-state energy of medium size systems ($N = 49$) with a large, but finite, computational effort. Table II summarizes the results for the average error of 20 SA and STUN runs for system sizes $N = 49$ and $N = 101$ as a function of the computational effort. In direct comparison we find that STUN is 2 orders of magnitude more efficient than SA. Both methods are at least a dozen orders of magnitude more efficient than SA on the original PES.

By using three *NP*-hard minimization problems with high barriers separating local minima we have demonstrated that the stochastic tunneling approach offers a reliable, generic, and efficient route for the determination of low-energy configurations. One chief advantage of the method lies in the fact that only a single parameter must be adjusted to adapt it to a specific problem. Fixing the functional form of the transformation and the "cooling schedule" demonstrates the adequacy of these choices but does not guarantee that they are optimal even for the problems considered here. One of the drawbacks of STUN is that, in contrast to, e.g., PT, no thermodynamic expectation values for the system can be obtained from the simulation. Also, because the nonlinear transformation will map any unbounded PES onto an interval bounded from above, the dynamical process in STUN will experience tunneling phases at any finite temperature. For PES that do not contain high barriers, or in the presence of efficient global moves, STUN may therefore be less efficient than competing methods. In many realistic optimization problems, where the construction of global moves is exceedingly difficult or very expensive, the tunneling approach can ameliorate the difficulties associated with the existence of high-energy barriers that separate local minima of the PES.

TABLE II. Average and best ground-state estimates for LABS for the $N = 49$ and $N = 101$ using SA and STUN on the locally minimized PES described in the text. The effort is given in thousands of steps, each step consists of a multispin flip followed by a local minimization.

Effort	SA	SC
	$N = 49$	
10	212.48/176	185.12/136
50	196.64/164	168.72/136
100	191.68/144	161.60/136
500	177.68/136	151.76/136
1000	175.52/136	139.44/136
	$N = 101$	
10	987.44/914	918.08/810
50	946.44/854	880.08/790
100	927.84/846	865.76/766
500	894.32/822	
1000	891.68/818	

We gratefully acknowledge stimulating discussions with C. Gros and U. Hansmann.

[1] C. H. Papadimitriou, *Computational Complexity* (Addison-Wesley, Reading, MA, 1994).
[2] M. R. Garey and D. S. Johnson, *Computers and Intractability—A Guide to the Theory of NP-Completeness* (Freeman, New York, 1979).
[3] K. Hamacher and W. Wenzel, Phys. Rev. E **58**, 938 (1999).
[4] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, Science **220**, 671 (1983).
[5] B. A. Berg and T. Neuhaus, Phys. Lett. B **267**, 249 (1991); J. Lee, Phys. Rev. Lett. **71**, 211 (1993); B. Hesselbo and R. B. Stinchcombe, Phys. Rev. Lett. **74**, 2151 (1995).
[6] A. P. Lyubartsev, A. A. Martinovski, S. V. Shevkunov, and

P. N. Vorontsov-Velyaminov, J. Chem. Phys. **96**, 1776 (1992); E. Marinari and G. Parisi, Europhys. Lett. **19**, 451 (1992).
[7] U. H. E. Hansmann, Chem. Phys. Lett. **281**, 140 (1997).
[8] J. Barhen, V. Protoposecu, and D. Reister, Science **276**, 1094 (1997).
[9] While the transformation in Eq. (1) is only one possible choice, a number of features constrain its construction: (i) It must be strongly nonlinear in the high-energy regime, to yield a nearly constant effective PES for high energies and true tunneling. (ii) There must be a parameter to modulate the compression (γ), as the height of the energy barriers is problem dependent. (iii) Requiring an essentially flat PES at high energy (for typically unbounded PES) requires a transformation that maps the interval $[f_0, \infty]$ onto some finite interval.
[10] C. Simone, M. Diehl, M. Jünger, P. Mutzel, and G. Reinelt, J. Stat. Phys. **80**, 487 (1995).
[11] W. H. Press *et al.*, *Numerical Recipes in C* (Cambridge University Press, Cambridge, England, 1995).
[12] F.-M. Dittes, Phys. Rev. Lett. **76**, 4651 (1996).
[13] J. Bernasconi, J. Phys. **48**, 559 (1987).
[14] W. Krauth and M. Mezard, Z. Phys. B **97**, 127 (1995).
[15] M. J. E. Golay, IEEE Trans. Inf. Theory **28**, 543 (1982).
[16] G. F. M. Beenker, T. Claasen, and P. W. C. Hermes, Philips J. Res. **40**, 289 (1985).
[17] A. Nayeem, J. Vila, and H. A. Scheraga, J. Comput. Chem. **12**, 594 (1991).