Proc. Natl. Acad. Sci. USA
Vol. 84, pp. 6611–6615, October 1987
Chemistry

Monte Carlo-minimization approach to the multiple-minima problem in protein folding*

(Metropolis criterion/Markov process/absorbing state/global optimization)

ZHENQIN LI AND HAROLD A. SCHERAGA

Baker Laboratory of Chemistry, Cornell University, Ithaca, NY 14853-1301

Contributed by Harold A. Scheraga, June 5, 1987

ABSTRACT **A Monte Carlo-minimization method has been developed to overcome the multiple-minima problem. The Metropolis Monte Carlo sampling, assisted by energy minimization, surmounts intervening barriers in moving through successive discrete local minima in the multidimensional energy surface. The method has located the lowest-energy minimum thus far reported for the brain pentapeptide [Met^5]enkephalin in the absence of water. Presumably it is the global minimum-energy structure. This supports the concept that protein folding may be a Markov process. In the presence of water, the molecules appear to exist as an ensemble of different conformations.**

Optimization procedures are required for an ultimate understanding as to how interatomic interactions lead to the folded, most-stable conformation of a protein from a linear polypeptide chain. A major problem in locating the global minimum of the empirical potential function that describes the conformations of a protein arises from the existence of many local minima in the multidimensional energy surface: the multiple-minima problem (1). This problem exists even for a system as small as a terminally blocked amino acid and becomes aggravated as the size of the system increases. Whereas algorithms are available for minimizing a function of many variables, none exist for passing from one local minimum, over an intervening barrier, to the next local minimum—and ultimately to the global minimum—in a many-dimensional space (1, 2). Several procedures have been developed to overcome this problem (1); these include the "buildup" method (3), optimization of electrostatics (4), relaxation of dimensionality (5, 6), adaptive importance sampling Monte Carlo (7–10), pattern recognition based on factor analysis of protein data (11, 12), use of distance constraints (13), and use of short-, medium-, and long-range interactions (14). Most of these procedures have been tested so far on short oligopeptides (up to 20 residues, in some cases), and their possible extension to proteins containing of the order of 100 residues would be of great interest. In our continual search for procedures to overcome this problem, we have developed an approach that appears to work very efficiently on the pentapeptide [Met^5]enkephalin (H-Tyr-Gly-Gly-Phe-Met-OH) and hopefully can be extended to larger structures. The application of this procedure to enkephalin is reported here.

The multiple-minima problem is not unique to protein folding but arises in many other fields of biology, chemistry, and physics whenever complexity appears (e.g., for intrinsically heterogeneous systems with a large number of strongly coupled degrees of freedom). A protein, composed of chemically distinct amino acids in a unique sequence, is a heterogeneous system that is fundamentally different from a homopolymer, and its many degrees of freedom contribute to the formidable difficulty of the multiple-minima problem. From a computational point of view, the multiple-minima problem is reminiscent of the NP (nondeterministic polynomial time) problem (15), in that the total number of possible conformations is an exponential function of the total number of degrees of freedom.

The approach taken here combines the power of conventional energy minimization (16) to find local minima and that of the Metropolis Monte Carlo method (17) in global combinatorial optimization (18). When implemented, it generates a Markov walk on the hyperlattice of all (discrete) energy minima, with Boltzmann transition probabilities. The working hypothesis ("Markovian hypothesis") underlying this method is (*i*) protein folding is a Markov process with Boltzmann transition probabilities and (*ii*) for a natural biologically active protein, such a Markov process leads to a unique absorbing state (19) (one in which equilibrium is reached after a sufficiently long time and in which the stationary probability of occurrence approaches unity), corresponding to the native structure of a protein. The method has been tested extensively on [Met^5]enkephalin, with interaction energies computed by the ECEPP/2 (empirical conformational energy program for peptides) algorithm (20–22). In the absence of water, the Monte Carlo-minimization procedure converges consistently to the same global minimum (a type II′ β-bend structure, the central two residues of which are Gly-Phe) for as many as 12 random starting conformations (and an additional one selected to have a different β-bend structure). In the presence of water, the molecule undergoes considerable structural fluctuations, with no unique stable structure, suggesting that a large ensemble of distinct conformations coexist at equilibrium.

THE MONTE CARLO-MINIMIZATION METHOD

Motivation. Experimental studies (23) have demonstrated that a protein is not a static structure but instead undergoes fluctuations. Based on photodissociation studies of carbon monoxide bound to myoglobin, it has been suggested that a protein can exist in a large number of conformational substates separated by barriers, with transitions among substates constituting equilibrium fluctuations (24). A recent molecular dynamics study of myoglobin (25) reported the existence of many minima in the vicinity of the native protein; these corresponded to relative reorientations of the α-helices coupled with rearrangements of the side chains, as a consequence of the internal dynamics of the protein. It follows, as a necessary condition that a structure be stable, that the native conformation of a protein must be stable not only against small disturbances but also against larger-scale thermal fluctuations; i.e., the native structure must be able to recover from any thermal impulse, even though the latter may (temporarily) lead to a different local minimum-energy

The publication costs of this article were defrayed in part by page charge payment. This article must therefore be hereby marked "*advertisement*" in accordance with 18 U.S.C. §1734 solely to indicate this fact.

*A portion of this work was presented at the Biophysical Society meeting, New Orleans, LA, February, 1987.

structure. A structure determined by energy minimization alone, which is stable only against small distortions, is very likely to be thermally unstable and hence cannot be admitted as a candidate for the native structure. These considerations suggest that thermal fluctuations play an indispensable role in selecting the native structure.

The fact that a protein in a thermodynamic environment can fold into its native structure within a time scale of milliseconds to seconds implies that, if we can reproduce the natural processes theoretically or at least simulate their essential and most relevant features (mainly thermal fluctuations and energetic processes), we may be able to devise a sufficiently efficient method to fold a protein.

Since a Metropolis Monte Carlo (26) method simulates natural thermal processes, by taking into account both random fluctuations and energetic considerations, it might be applicable to protein folding. The successful application of the simulated annealing method (18), which is essentially a Metropolis Monte Carlo simulation technique with an artificial "temperature," to the computationally difficult "traveling-salesman problem" [which is in the class of NP problems (15)] is very similar to the multiple-minima problem, in that the total number of possible solutions is a nonpolynomial function of the number of cities.

A straightforward application of the Metropolis Monte Carlo method to polypeptides, however, has proven to be very inefficient (7, 27, 28), or even impossible, because we have to search a high-dimensional conformational space rather than discrete states. Conventional Metropolis Monte Carlo samples the whole space by making small increments in each step. The large energy barriers in the conformational space of a protein make such a method impractical because, for most of the time, the sampling is confined to a very restricted region of the whole conformational space. A different type of Monte Carlo algorithm (7–9) is an alternative approach to this problem. To overcome these difficulties, we have developed the Monte Carlo-minimization method, which randomly samples only the discrete set of energy minima instead of the whole conformational space.

Implementation. The method consists of three components.

Step i. The first is a Monte Carlo sampling strategy, which satisfies the ergodicity requirement; i.e., any local minimum is accessible from any other one after a finite number of random sampling steps. In addition to this thermodynamic condition, we find it necessary to restrict the random sampling to only a few variables at each time in order to maintain sampling efficiency and to simulate the kinetic nature (26) (i.e., the short time correlations) of natural thermal processes in which physical fluctuations are localized. Even with these two requirements, there is still considerable freedom to choose the form of the sampling technique. We have chosen to make a completely random change ($-180° \leq \theta \leq 180°$) in one randomly selected dihedral angle θ (where θ stands for ϕ, ψ, ω, or the χ values of the given residue) among all the variable dihedral angles of the molecule. The choice of random changes in *two* randomly picked dihedral angles was also tested and led to the same global minimum for [Met^5]enkephalin, but the acceptance ratio (about 20%) was considerably lower than when one variable at a time was chosen (about 45%); this inefficiency arises because this modified procedure samples the energetically unfavorable regions more frequently. In this step, a possible stochastic event is generated, leading to a conformation that generally is not at a local minimum.

Step ii. The randomly chosen conformation of step *i* is then subjected to conventional minimization with the SUMSL routine (16) using the ECEPP/2 energy function (20–22) to reach the nearest local minimum (a state on the hyperlattice of all energy minima).

Step iii. This local minimum is examined by the Metropolis criterion (17) to compare it with the previously accepted local minimum to update the current conformation. As a consequence, the transition probabilities of the series of local minima generated in the Markov process satisfy the Boltzmann distribution (17, 26). Then step *i* is repeated to continue the iteration process, which generates a Markov sequence with Boltzmann probabilities.

The energy minimization (step *ii*) is the most time-consuming part of the whole procedure and consumes about 95% of the overall computational time. For a sample run on [Met^5]enkephalin on an IBM-3090 computer, one Monte Carlo-minimization iteration (steps *i* to *iii*), involving about 100 energy evaluations, took about 4 sec. An arbitrary cut-off of 10,000 iterations was needed for frequent convergence to the global minimum, for a total time of about 10 hr; in some runs, convergence occurred in as little as 1 hr.

The procedure described above pertains to calculations carried out in the absence of water. In the presence of water, the method is basically the same, except that the "energy" in the above procedure is replaced by the "total energy" [i.e., the sum of the ECEPP energy and an empirical hydration free energy (29)]. Since energy minimization involving such hydration terms is extremely time-consuming, we decided not to include the hydration term in the energy minimization (step *ii*), to speed up the procedure; instead, after minimization of the ECEPP energy, the hydration free energy was included in the total energy at the local minimum, for the selection of the current conformation by the Metropolis criterion (step *iii*).

RESULTS

[Met^5]Enkephalin in the Absence of Water. Twenty-four variable backbone and side-chain dihedral angles are required to specify the conformation of [Met^5]enkephalin. The total number of possible local minima in its conformational energy surface is estimated to be more than 3^{24}, or 10^{11}. Nevertheless, the apparent global minimum of this peptide (Fig. 1; a type II′ β-bend involving Gly-Gly-Phe-Met) was first located by the Monte Carlo-minimization method in 4 hr and was subsequently also achieved by another independent method developed in this laboratory (6). To show that the structure in Fig. 1 has a high probability of being the global minimum, we repeated the procedure by starting from 16 additional independently generated random conformations (chosen from the whole conformational space) and from one conformation that was selected arbitrarily to have a β-bend in a location different from that of the global minimum (i.e., involving Tyr-Gly-Gly-Phe). Within 10,000 iterations, 12 out of 17 randomly generated initial conformations, plus the arbitrarily selected β-bend structure, led to the same global minimum-energy structure that had identical values (within 0.5°) of all backbone and side-chain dihedral angles ϕ, ψ, ω, and χ's and identical energies (within 0.01 kcal/mol) (Table 1). The other five randomly initiated conformations had not yet converged (within 10,000 iterations, or about 10 hr on the IBM-3090) and all had energies at least 2.0 kcal/mol higher than that at the global minimum and shapes that differed considerably from the global-minimum structure at the times that these five runs were terminated. These results clearly demonstrate the power of the Monte Carlo-minimization procedure to reach the apparent global minimum from random starting conformations that are distinctively different.

Based on the fact that 12 out of 17 random starting conformations converged to the same lowest-energy structure, the probability of finding another structure with energy lower than that at the apparent global minimum is estimated to be less than 0.01%. These results demonstrate that the

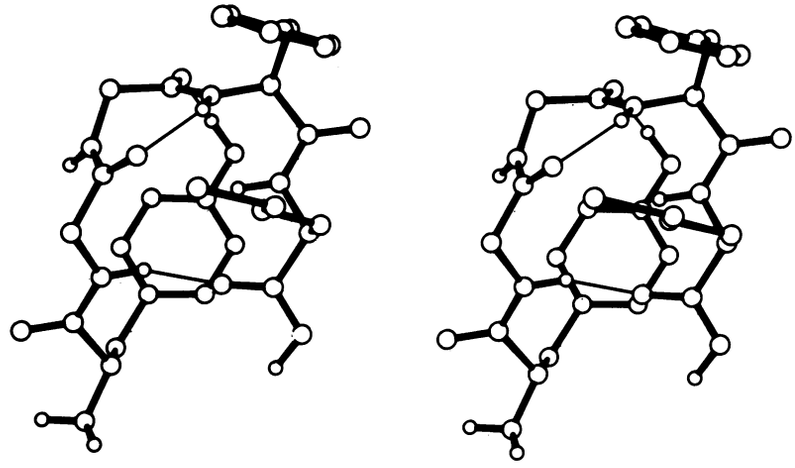

FIG. 1. Stereoview of the global minimum-energy structure of [Met^5]enkephalin in the absence of water, located by the Monte Carlo-minimization method (at −20°C and 40°C) using the ECEPP/2 parameters. It is a type II′ β-bend involving Gly-Gly-Phe-Met.

Monte Carlo-minimization procedure is free of bias in the choice of starting conformations. This situation is probably encountered because our method generates a Markov process with presumably good ergodic properties (which, by definition, is independent of the initial state after a long enough time), as a result of the random change in a randomly chosen variable, followed by energy minimization to bypass large energy barriers. In principle, the global minimum can be reached from any starting conformation after a long enough time needed to reach equilibrium; these results demonstrate that the required time is a practical one for this pentapeptide with presently available computers.

The global minimum-energy structure has been reached at more than one temperature, within a range of −20°C to 40°C. Four of the 17 runs from random starting conformations were carried out at 40°C, and one of them reached the global minimum. The remaining 13 were carried out at −20°C, and 11 of them reached the global minimum. The one run from a nonrandom start was carried out at 40°C and reached the global minimum. At these temperatures, fluctuations are significantly large; structures with energies up to 2 kcal/mol above the global minimum occurred quite frequently in these simulations. In fact, fluctuating structures surrounding the global minimum [with minor variations (usually within 10°) in a few variables] generally occur more frequently than the global minimum itself (even after the first appearance of the global minimum), a natural statistical mechanical phenomenon. Nevertheless, the global minimum appears sufficiently frequently to indicate convergence of the algorithm and establishment of thermal equilibrium. It is worth noting that fluctuating changes in some dihedral angles can be as large as 60° without a major increase in energy. Our experience with this algorithm thus far shows that it is most efficient (in reaching the global minimum, in the absence of water) at the lower end of the temperature range explored (i.e., at −20°C).

Table 1. Global minimum-energy structure of [Met^5]enkephalin in the absence of water

Residue	Dihedral angles, degrees						
	ϕ	ψ	ω	χ^1	χ^2	χ^3	χ^4
Tyr-1	−86	156	−177	−173	79	−166	
Gly-2	−154	83	169				
Gly-3	84	−74	−170				
Phe-4	−13	19	−174	59	95		
Met-5	−164	160	−180	53	175	−180	−59

The conformational energy is −12.90 kcal/mol, based on the latest ECEPP/2 parameters (21, 22).

[Met^5]Enkephalin in the Presence of Water. As many as five runs of the Monte Carlo-minimization have been performed to simulate [Met^5]enkephalin in water at 20°C; each run involved more than 10,000 iterations. Four out of the five runs were carried out from random starting conformations; the other one started from the global-minimum structure in the absence of water. In contrast to the results in the absence of water, where the Monte Carlo-minimization converged consistently to the same global minimum for [Met^5]enkephalin, all of the five runs led to different conformations with comparable total "energies" (ECEPP energy plus hydration free energy). Among the sequence of accepted conformations for each run of 10,000 iterations, we were not able to identify any stable structure that occurred with significant frequency. These results indicate that [Met^5]enkephalin in water at 20°C is likely to be in an "unfolded" state, in which a large ensemble of distinct conformations coexist at equilibrium. The validity of this suggestion can be tested experimentally—e.g., by comparing the circular dichroism or nuclear magnetic resonance spectrum of [Met^5]enkephalin in solution with that in ordered structures—as in various crystalline environments.

DISCUSSION

The apparent success of the Monte Carlo-minimization method in locating the global minimum-energy structure of [Met^5]enkephalin in the absence of water (the fluctuating neighborhood, which also happens to be the most populated in the temperature range investigated) may have implications for the dynamic processes involved in protein folding. Although this method involves an artificial sampling strategy, some key features of the procedure may reflect the process by which the multiple-minima problem is overcome in nature.

The Monte Carlo-minimization method is an implementation of the underlying working hypothesis stated earlier (we propose that it be called the Markovian hypothesis): protein folding is a Markov process whose transition probabilities satisfy the detailed balance (26) (equilibrium between two states) required by equilibrium statistical mechanics; the folded structure of a biologically active protein corresponds to an absorbing state, or a tight set of neighboring absorbing states, of the Markov process in a given thermodynamic environment. This hypothesis implies both that proteins fold by a Markov process with Boltzmann transition probabilities and that the native conformations are stochastically stable with respect to random fluctuations. As an operational hypothesis, it may describe the way proteins fold in nature and may be used as the basis of a method to predict the

structures of native proteins. Under such a hypothesis, the unfolding of a protein is exactly the same physical process as protein folding, except that the absorbing state(s) diffuses out as a result of different thermodynamic conditions. In the following discussion, we shall concentrate only on protein folding, although the physical picture is the same for unfolding.

The Monte Carlo-minimization is one of many possible procedures used to implement the Markovian hypothesis by generating a Markov process with Boltzmann transition probabilities, to simulate natural protein folding. The artificiality and arbitrariness of the (ergodic) random sampling strategy in this method becomes irrelevant as the Markov process approaches equilibrium in the long-time limit: the equilibrium probabilities are independent of the particular choice of sampling strategy. The existence of an absorbing state of the Markov process would guarantee convergence of such a procedure to the thermally most stable structure from *any* starting conformation.

The formalism of a Markov process is a simple description of a nonequilibrium stochastic process whose long-time limit leads to the equilibrium behavior (19, 26). At the beginning of real protein folding, the transition probabilities may not satisfy detailed balance; i.e., they may not obey the Boltzmann distribution. But, since protein folding occurs on a time scale much longer than that needed for (spatial and temporal) local equilibrium, under the assumption that folding is a Markov process the initial nonequilibrium character of the transition probabilities is forgotten after a long enough time; thus, the Monte Carlo-minimization procedure with Boltzmann transition probabilities may be an accurate simulation of a real folding process on a long time scale.

If the Markovian hypothesis of protein folding is correct, then protein folding has both thermodynamic and kinetic features. The transition probabilities are thermodynamic in origin, due to local equilibrium among individual variables within the time scale of thermal transitions. On the longer time scales needed for equilibration among more degrees of freedom (characteristic of real protein folding), however, the thermally accessible conformations at any given moment are always limited, depending on the conformation at the immediate past, under the assumption that physical fluctuations are local with coherence at any time extending to only a few among the many variables. (At equilibrium, an instantaneous large-scale change involving *all* degrees of freedom is considered to be unphysical and is not allowed to occur in our Monte Carlo-sampling procedure.) These features, especially the rapid equilibrium among individual degrees of freedom and the relatively long-time structural correlations for many degrees of freedom, are characteristic of large heterogeneous systems.

Furthermore, it can be demonstrated that the thermodynamic hypothesis (30) and the kinetic hypothesis (31) of protein folding can be incorporated into the Markovian hypothesis as two specific cases. For a given protein in water, the relaxation time of the Markov (folding) process is well defined, corresponding to the time needed for full equilibration of the total system. If this relaxation time is shorter than the experimental time scale, then the thermodynamic hypothesis holds (i.e., protein folding is thermodynamically controlled); if otherwise, then the kinetic hypothesis is true (i.e., protein folding is a kinetically controlled process).

As an implementation of the Markovian hypothesis, the Monte Carlo-minimization method owes its efficiency to the interplay between random fluctuations (through Monte Carlo sampling) and directional energetic processes (by energy minimization and the Metropolis criterion). One aspect of this interplay manifests itself in the balance of thermodynamic and kinetic features. Fluctuations must be present, yet they must be localized (i.e., incremental or gradual); otherwise, an energetically favorable structure would be too difficult to find or too easy to lose. To obtain an efficient procedure, we have limited the Monte Carlo sampling along the trajectory to a few variables each time, allowing large changes in these few variables even though real protein folding may simultaneously sample all degrees of freedom with small fluctuational amplitudes. Another more obvious aspect of this interplay is through the balance of entropic and energetic processes (i.e., random sampling and energy minimization-plus-Metropolis criterion) at a finite temperature. To illustrate this, let us consider two extreme cases in the Monte Carlo-minimization method. Suppose the temperature is close to infinity, so that the energetic processes are not important in selecting the energy minima. Then the Monte Carlo-minimization simulates purely random fluctuations. It is just an exhaustive random walk on the whole hyperlattice of energy minima; therefore, it is not efficient. Now, suppose the temperature is close to absolute zero. Under this condition, only structures with energy lower than the present one are accepted by the Metropolis criterion. The acceptance ratio is so low that fluctuations do not play a significant role: the protein could be trapped in some local minima for a long time. This is the problem encountered in many other global minimization methods (32–34), which correspond to the Monte Carlo-minimization procedure at absolute zero. At an optimal temperature (which is system-dependent), both random fluctuations and energetic processes are important in locating the global minimum. The computational time needed to reach the global minimum depends critically on the physical nature of the problem (i.e., on the interaction energies and the temperature), along with a dependence on the sampling strategy, starting conformations, and intermediate stochastic sequences produced by the random-number generator.

The possibility that protein folding is a Markov process was investigated previously (35) on the basis of Monte Carlo studies of a two-dimensional lattice model of proteins, but prior to the development of our Monte Carlo-minimization method, it has not been possible to implement this hypothesis directly on a realistic system (e.g., [Met^5]enkephalin) using realistic energy parameters (e.g., ECEPP) to test out the implied consequences of this hypothesis for a real protein.

The thermalization procedure of Levitt and Warshel (36) is similar to our Monte Carlo-minimization method. There is, however, one major difference: the thermal fluctuations in their procedure are distributed uniformly; these do not satisfy the Boltzmann distribution, so energetically unfavorable structures (local minima) are likely to arise. In contrast, the Metropolis criterion used in our method ensures that the fluctuations obey the Boltzmann distribution, which is crucial in stochastically driving a protein toward a folded structure under favorable thermodynamic conditions, by eliminating thermally unstable conformations. In a changing thermodynamic environment, it is again this Metropolis criterion, which is critically important in mediating a folding–unfolding transition in a protein. In fact, the Metropolis criterion is the distinguishing factor between the Monte Carlo-minimization method discussed in this paper and other stochastic minimization methods (32–34, 36, 37); in the latter case, the dynamic process of equilibrium thermal transitions during structure formation is not adequately addressed.

The Monte Carlo-minimization method formulated in this article is based on classical mechanics and classical statistical mechanics, although the potential energy function used may be of quantum mechanical origin. Entropic contributions to the statistical weight of a structure, which are measured by the relative frequency of occurrence of the structure in a long enough Monte Carlo-minimization simulation, are automatically taken into account through Monte Carlo random sampling. This method is best adapted to study long-time and global properties of proteins undergoing large-scale structur-

al changes. In contrast, molecular dynamics gives us better pictures of the detailed dynamics of proteins, although it is usually limited to short time scales (of the order of picoseconds) with a given starting conformation. Because the Monte Carlo-minimization method can overcome energy barriers by random changes and reminimization, it can be viewed as operating on a "renormalized" or "coarse-grained" time scale characteristic of thermally activated processes; therefore, it is more relevant than molecular dynamics (as currently practiced) in studies of structural changes on a time scale longer than that of elementary thermal transitions, in particular in studies of protein folding.

PERSPECTIVES

Protein folding is a prototypical example of the self-organization of complex systems in nature. The application of the Monte Carlo-minimization method to oligopeptide folding not only provides us with insights into the physical nature of protein folding but also may be a powerful predictive scheme for elucidating native protein structures, once we know precisely all relevant energy parameters: the interactions among protein constituents, that between protein constituents and water, and that between water molecules. Furthermore, as a method capable of studying large-scale structural changes, of which protein folding is only one example, the Monte Carlo-minimization method may be of value in the investigation of properties and structure–function relationships of proteins. Since the method is a simple and general approach to global optimization of any continuous function (2), it should be applicable to many other important problems where the multiple-minima problem appears [e.g., determination of the structures of other biomolecules and those of organic molecules, of complex surfaces (38), and of chemical reaction pathways (39)]. These applications require only a knowledge of realistic parameters characterizing the interactions and proper identification of the degrees of freedom of the system, which are just the same information needed in conventional energy minimization. It remains to be seen whether the Monte Carlo-minimization procedure can be extended efficiently to larger polypeptide (protein) structures, with the incorporation of realistic interaction terms to take hydration into account.

We thank K. D. Gibson, G. Némethy, E. O. Purisima, and M. Vásquez for helpful comments on this paper. This work was supported by research grants from the National Science Foundation (DMB84-01811) and from the National Institute of General Medical Sciences of the National Institutes of Health (GM-14312). The computations were carried out on the Cornell Production Supercomputer Facility of the Center for Theory and Simulation in Science and Engineering, which is funded, in part, by the National Science Foundation, New York State, and IBM Corporation.

1. Gibson, K. D. & Scheraga, H. A. (1987) in *Structure and Expression: Nucleic Acids and Proteins*, eds. Sarma, R. H. & Sarma, M. H. (Adenine, Guilderland, NY), in press.
2. Dixon, L. C. W. & Szegö, G. P. (1978) in *Towards Global Optimization*, eds. Dixon, L. C. W. & Szegö, G. P. (North Holland, New York), Vol. 2, pp. 1–15.
3. Vásquez, M. & Scheraga, H. A. (1985) *Biopolymers* **24,** 1437–1447.
4. Piela, L. & Scheraga, H. A. (1987) *Biopolymers* **26,** S33–S58.
5. Purisima, E. O. & Scheraga, H. A. (1986) *Proc. Natl. Acad. Sci. USA* **83,** 2782–2786.
6. Purisima, E. O. & Scheraga, H. A. (1987) *J. Mol. Biol.* **196,** 697–709.
7. Paine, G. H. & Scheraga, H. A. (1985) *Biopolymers* **24,** 1391–1436.
8. Paine, G. H. & Scheraga, H. A. (1986) *Biopolymers* **25,** 1547–1560.
9. Paine, G. H. & Scheraga, H. A. (1987) *Biopolymers* **26,** 1125–1162.
10. Scheraga, H. A. & Paine, G. H. (1986) *Ann. N.Y. Acad. Sci.* **482,** 60–68.
11. Kidera, A., Konishi, Y., Oka, M., Ooi, T. & Scheraga, H. A. (1985) *J. Protein Chem.* **4,** 23–35.
12. Kidera, A., Konishi, Y., Ooi, T. & Scheraga, H. A. (1985) *J. Protein Chem.* **4,** 265–297.
13. Wako, H. & Scheraga, H. A. (1981) *Macromolecules* **14,** 961–969.
14. Wako, H. & Scheraga, H. A. (1982) *J. Protein Chem.* **1,** 5–45, 85–117.
15. Garey, M. R. & Johnson, D. S. (1979) *Computers and Intractability: A Guide to the Theory of NP-Completeness* (Freeman, San Francisco).
16. Gay, D. M. (1983) *ACM Trans. Math. Software* **9,** 503–524.
17. Metropolis, N. A., Rosenbluth, A. W., Teller, A. H. & Teller, E. (1953) *J. Chem. Phys.* **21,** 1087–1092.
18. Kirkpatrick, S., Gelatt, C. D., Jr., & Vecchi, M. P. (1983) *Science* **220,** 671–680.
19. Bharucha-Reid, A. T. (1960) *Elements of the Theory of Markov Processes and Their Applications* (McGraw-Hill, New York).
20. Momany, F. A., McGuire, R. F., Burgess, A. W. & Scheraga, H. A. (1975) *J. Phys. Chem.* **79,** 2361–2381.
21. Némethy, G., Pottle, M. S. & Scheraga, H. A. (1983) *J. Phys. Chem.* **87,** 1883–1887.
22. Sippl, M. J., Némethy, G. & Scheraga, H. A. (1984) *J. Phys. Chem.* **88,** 6231–6233.
23. Gurd, F. R. N. & Rothgeb, T. M. (1979) *Adv. Protein Chem.* **33,** 73–165.
24. Ansari, A., Berendzen, J., Bowne, S. F., Frauenfelder, H., Iben, I. E. T., Sauke, T. B., Shyamsunder, E. & Young, R. D. (1985) *Proc. Natl. Acad. Sci. USA* **82,** 5000–5004.
25. Elber, R. & Karplus, M. (1987) *Science* **235,** 318–321.
26. Binder, K. (1986) in *Monte Carlo Methods*, ed. Binder, K. (Springer, New York), 2nd Ed., pp. 1–46.
27. Hagler, A. T., Stern, P. S., Sharon, R., Becker, J. M. & Naider, F. (1979) *J. Am. Chem. Soc.* **101,** 6842–6852.
28. Rapaport, D. C. & Scheraga, H. A. (1981) *Macromolecules* **14,** 1238–1246.
29. Kang, Y. K., Némethy, G. & Scheraga, H. A. (1987) *J. Phys. Chem.* **91,** 4105–4120.
30. Anfinsen, C. B. (1973) *Science* **181,** 223–230.
31. Wetlaufer, D. B. & Ristow, S. (1973) *Ann. Rev. Biochem.* **42,** 135–158.
32. Robson, B. & Osguthorpe, D. J. (1979) *J. Mol. Biol.* **132,** 19–51.
33. Finn, P. W., Robson, B. & Griffiths, E. C. (1984) *Int. J. Peptide Protein Res.* **24,** 407–413.
34. Rao, G. S., Tyog, R. S. & Mishra, R. K. (1981) *Int. J. Quantum Chem.* **20,** 273–279.
35. Gō, N. (1983) *Ann. Rev. Biophys. Bioeng.* **12,** 183–210.
36. Levitt, M. & Warshel, A. (1975) *Nature (London)* **253,** 694–698.
37. Saunders, M. (1987) *J. Am. Chem. Soc.* **109,** 3150–3152.
38. Chadi, D. J. (1978) *Phys. Rev. Lett.* **41,** 1062–1065.
39. Morokuma, K. (1982) in *Frontiers of Chemistry*, ed. Laidler, K. J. (Pergamon, New York), pp. 143–158.