

Fig. 7 A schematic north-south section through the Skaergaard intrusion showing structural and volumetric relationships consistent with known data.

As a further test of the experimentally determined liquid compositions, one can calculate the compositions and proportions of mineral assemblages that must be removed from each liquid to proceed to its successor. These fractionated assemblages should match corresponding units of the Layered Series and Border Groups for appropriate levels. When mass balance calculations of this kind are carried out for the new Skaergaard data, it is found that consistent solutions are obtained, but there is some uncertainty in the composition and relative volumes of the Border Groups, for which few data are available. For this reason, a reliable solution of the volume problem must await a thorough study of additional samples collected in 1974; using data obtained so far, it has been possible to make preliminary estimates of the relative volumes of the principal units of the intrusion. By relating these proportional volumes to measured stratigraphic thicknesses one can estimate the lateral extent of each layer and from this the form of the intrusion. Figure 7 illustrates a schematic section that is consistent with geophysical data and average compositions of the rocks and liquids as they are known at this time.

Apart from the great reduction of the volume of the Hidden Zone, the principal differences between these and previous estimates are in the greater lateral extent of the middle portions of the intrusion from Lower Zone B through the Middle Zone.

A further complication arises from the recent discovery that the phase layering of the Layered Series transgresses structural and stratigraphic horizons. The sequence of appearance of mafic phases is displaced upward in the eastern part of the intrusion relative to a section on the west side, so that cumulus clinopyroxene and magnetite appear and olivine disappears at levels that rise toward the east relative to distinctive horizons, such as the conspicuous Triple Group. The fact that the composition of plagioclase does not seem to be affected by this phenomenon suggests that the liquidus temperatures of iron-bearing minerals reflected a difference in oxygen fugacity from one side of the intrusion to the other. This feature was only recognised very recently and requires more study.

This work was supported by grants from the NSF, the National Geographic Society and the Geological Society of America. Experimental work was carried out at the University of Edinburgh and the Carnegie Institution of Washington.

Received December 16, 1974.

- 1 Wager, R. L., *J. Petrol.*, **1**, 364-398 (1960).
- 2 Taylor, H. P., and Forester, R. W., *Trans. Am. Geoph. Un.*, **54**, 500 (1973).
- 3 Leeman, W. P., Goles, G. G., McBirney, A. R., and Dasch, E. J., *A. Meet. Geol. Soc. Am.*, **5**, 710-711 (1973).
- 4 Chayes, F., *J. Petrol.*, **11**, 1-14 (1970).
- 5 Blank, H. R., and Gettings, M. E., *Trans. Am. Geoph. Un.*, **54**, 507 (1973).
- 6 Williams, R. J., *Am. J. Sci.*, **271**, 132-146 (1971).
- 7 Tilley, C. E., Yoder, H. S. Jr, and Schairer, J. F., *Yb. Carnegie Instn. Wash.*, **62**, 77-84 (1963).
- 8 Henderson, P., *J. Petrol.*, **11**, 463-473 (1970).
- 9 Holgate, N., *J. Geol.*, **62**, 439-480 (1954).
- 10 Philpotts, A. R., *Econ. Geol.*, **62**, 303-330 (1967).
- 11 Biggar, G. M., and O'Hara, M. J., *Min. Mag.*, **37**, 198-205 (1969).
- 12 Roedder, E., and Weiblen, P. W., *Science*, **167**, 641-644 (1970).
- 13 McBirney, A. R., and Nakamura, Y., *Trans. Am. Geoph. Un.*, **55**, 450 (1974).
- 14 Roedder, E., *Am. Min.*, **36**, 282-286 (1951).
- 15 Greig, J. W., *Am. J. Sci.*, ser. 5, **13**, 1-44, 133-154 (1927).
- 16 Weiblen, P. W., and Roedder, E., *Geochim. cosmochim. Acta.*, Suppl. 4, **1**, 681-703 (1973).
- 17 Coles, G. G., and McBirney, A. R., *Trans. Am. Geoph. Un.*, **54**, 1225 (1973).
- 18 De Aniruddha, *Geol. Soc. Am. Bull.*, **85**, 471-474 (1974).

Computer simulation of protein folding

Michael Levitt* & Arieh Warshel*

Department of Chemical Physics, Weizmann Institute of Science, Rehovoth, Israel

A new and very simple representation of protein conformations has been used together with energy minimisation and thermalisation to simulate protein folding. Under certain conditions, the method succeeds in 'renaturing' bovine pancreatic trypsin inhibitor from an open-chain conformation into a folded conformation close to that of the native molecule.

PROTEIN molecules owe their enormous functional versatility to the fact that they spontaneously fold into complicated and unique conformations determined by the particular amino-acid sequence¹. Discovering the relationship between protein sequence and conformation is a fascinating theoretical problem of fundamental importance. Most previous theoretical work has used the concept of 'local structure', in which the conformation of a short segment of polypeptide chain is supposed to depend almost entirely on the sequence of that segment. Although this approach has helped understand local secondary structure^{2,3}, it has not shown how residues distant along the chain can come together to form the overall conformation. The only promising attempt to study the tertiary folding of a

protein, in this case myoglobin, was based on the packing of cylinders supposed to represent α helices⁴. The method was not implemented on a computer and cannot be applied more generally to other proteins not built entirely from helices.

Here we tackle the problem differently. First, we simplify the representation of a protein by averaging over the finer details. This is done both to make the calculations much more efficient and also to avoid having to distinguish between many conformations that differ only in these finer details. Second, we simulate the folding of this simple structure by the combination of use of convergent energy minimisation and normal mode thermalisation, which accelerate the process by avoiding many non-productive random fluctuations that occur in nature. Tests of the procedure on bovine pancreatic trypsin inhibitor (PTI), show that under certain conditions it can rapidly reproduce the correct overall folding of this small protein molecule.

Simple representation of protein structure

Even the smallest protein (say 50 residues) is extremely complicated, with about 750 atoms and 200 degrees of freedom (single bond torsion angles). Calculating its free energy presents severe computational difficulties, in particular when considering interactions with the rapidly moving solvent molecules and

*Present address; MRC Laboratory of Molecular Biology, Hills Road, Cambridge, UK.




Fig. 1 Relationship between structure into proteins. The simplified model and the C^α. E is the torsion angle (i-1, i, i+1), same simplified torsion angle molecule weight protein conformation.

thermal motion designed to overcome assumptions: to be eliminated folding can be variables (those changes in conformation).

Averaging over simplified structures, the C^α interactions are while the C^α amino-acid residues angle about the helix. Although a simple method has been used before, it has been applied to reduce the degrees of freedom of the protein. The effect of these variables is in the conformation used. The equivalent conformations obtained by averaging a dipeptide are done on the gly, ala-pro. The effective protein conformation is determined by the minimum at a




Fig. 1 Relationship between the simplified model of protein structure introduced here and the real all-atom structure of proteins. The two reference points for each residue in the simplified model correspond to the centroid of the side chain and the C^α . Each residue is only allowed one degree of freedom: the torsion angle α between the 4 successive C^α s of residues ($i-1, i, i+1, i+2$). All the side chains of a given type have the same simplified geometry. The bond lengths, bond angles, and torsion angles used to define the geometry of the simplified molecule were taken as the average values found in eight protein conformations, though they could just as well have been taken from amino-acid model compounds.

thermal motion of parts of the protein itself. Our method is designed to overcome these problems and is based on two assumptions: (1) that much of the protein's fine structure can be eliminated by averaging, and (2) that the overall chain folding can be obtained by considering only the most effective variables (those that vary most slowly yet cause the greatest changes in conformation).

Averaging over groups of atoms in the full structure gives a simplified structure with each residue represented by only two centres, the C^α atom, and the centroid of the side chain. Interactions are assumed to occur only between side chains, while the C^α positions define the chain path (Fig. 1). Each amino-acid residue only has one degree of freedom, the torsion angle about the line joining two adjacent C^α s (known here as α). Although a simple representation based on virtual bonds has been used before to study polypeptide random coils⁵, it has never been applied to ordered globular proteins. This simplification reduces the degrees of freedom by a factor of four and the number of interaction centres by a factor of fifteen. One might also hope that the reduced space used here to describe different conformations would have many fewer energy minima. The space is of lower dimension and the side chains are smooth curves without all the minor bumps of the all-atom structures. The effect of the fine details and more rapidly changing variables is included in the effective time-averaged potential functions used. (By the ergodic theorem⁶, this time averaging is equivalent to Boltzmann weighted spatial averaging over conformations generated by changing the fast variables.) For calculations about the torsion angle α , the effective potential is obtained by averaging the energy over all those conformations of a dipeptide that have a particular value of α . As it was not possible to study all 400 different dipeptides, calculations were done on six considered most representative: ala-ala, ala-gly, ala-pro, gly-gly, gly-ala and pro-ala. This showed that the effective potential only depended on the nature of the preceding amino acid, giving different potentials for the α preceding ala, gly, and pro. The alanine potential had a deep minimum at $\alpha = 210^\circ$ (twisted β chain) and a more shallow minimum at $\alpha = 45^\circ$ (α helix); the glycine potential had a

broad minimum at $\alpha = 0^\circ$ (reverse turn); and the proline potential had two sharp minima at $\alpha = 60^\circ$ and $\alpha = 210^\circ$. Because aspartic acid and asparagine were found to occur as frequently in the reverse turns of known protein conformations as glycine, the same potential was used for all three. The alanine-type potential was used for all other amino acids except proline. All the atoms were included in these dipeptide calculations, which used energy parameters derived from crystals of amino acids, amides, and hydrocarbons.

The interaction potential between a pair of identical amino-acid side chains was also calculated by spatial averaging. Each side chain was assumed to be spherically symmetrical with a radius equal to the average radius of gyration of that side chain. The effective potential was calculated at various distances apart as a sum of the interaction energy of all atoms anywhere in one sphere with all atoms anywhere in the other




Fig. 2 Stereo ribbon drawings of PTI in: *a*, the idealised native molecule, *b*, the minimum energy conformation generated starting at the idealised native conformation and *c*, the best conformation generated by folding from an extended chain with a terminal helix (the final conformation in Fig. 4). (The programs used to rotate the molecules into the same orientation and then draw the ribbon between C^α s were provided by Dr A. D. McLachlan.)

sphere. This effective potential between identical side chains was approximated by a Lennard-Jones type function, and potentials between pairs of different side chains were obtained by a geometric mean combining law. Because proteins fold in water not a vacuum, interactions with the solvent are included by assigning to each side chain a hydrophobic energy taken from the solubilities of amino acids in water and in ethanol⁷. In the calculations, the energy of transfer between these two solvents was taken as the difference in energy of the side chain when isolated in water and when completely surrounded by other residues. When surrounded by an intermediate number of neighbours, the hydrophobic energy was varied according to a sigmoid function. More complicated models that include hydrogen bonds and S-S bridges will be described elsewhere, together with full details of the standard geometry and all energy parameters used.

The folding of this idealised protein can be simulated by solving the equations of molecular dynamics at sufficiently small time intervals. In a viscous medium like water, these equations of motion can be approximated by Langevin equations, where the change in the variables is directed down the energy gradient with a random deflection due to Brownian motion⁸. For greater computational efficiency we neglect these thermal fluctuations while the chain folds, and the end point of the trajectory is the potential energy minimum accessible from the starting conformation. We minimise the energy of the idealised protein chain with respect to all the α angles using a powerful quadratically convergent method (VA09A, by R. Fletcher and taken from the Harwell Subroutine Library). After reaching a minimum, thermal fluctuations are reintroduced and the conformation is considered to be vibrating about the minimum so that each normal mode has average kinetic energy $kT/2$ (where k is the Boltzmann constant and T the absolute temperature). A new starting conformation for the next pass of energy minimisation is chosen by suddenly stopping the thermal vibration. At this time each normal-mode coordinate will be displaced randomly from the minimum by $(R(n)kT/\lambda)^{1/2}$, so that the associated energy becomes $R(n)kT/2$. Here λ is the eigenvalue of the energy second derivative matrix corresponding to the particular normal mode, and $R(n)$ is a random number uniformly distributed between 0 and 1. (An exponential distribution of random numbers between 0 and ∞ would be more realistic.) Normal-mode thermalisation avoids non-productive changes in conformation for it knows which combinations of angle changes should cause the greatest change in conformation for a given energy increase.

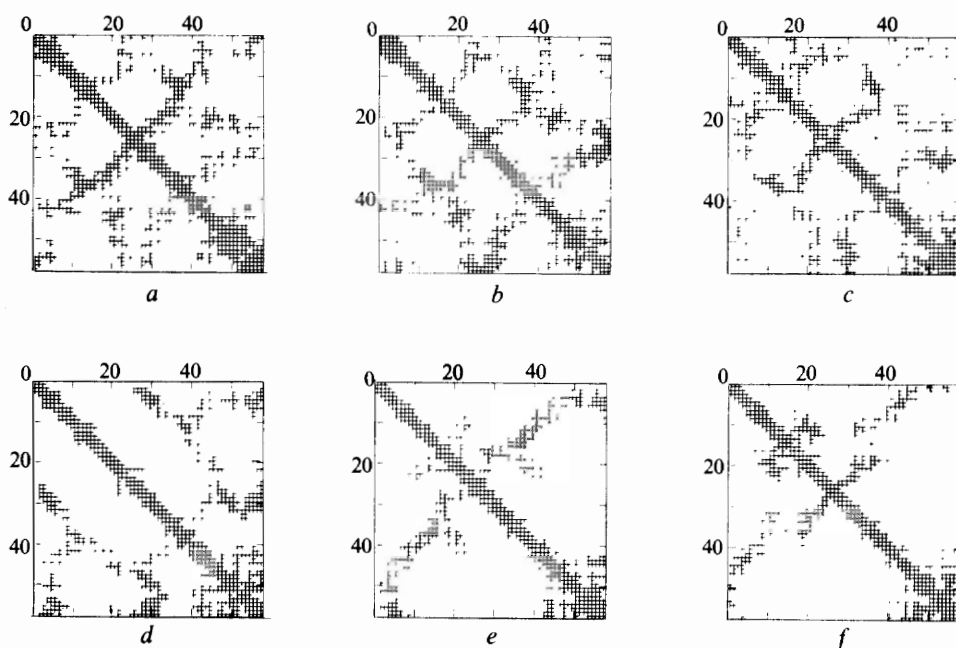


Fig. 3 Contact maps¹¹ of the following conformations: *a*, idealised native; *b*, folded from idealised native; *c*, folded from extended chain with terminal helix; *d*, *e*, *f*, folded from the same starting conformation as *c*, but using different sequences of random numbers for the thermalisation. The six structures shown here have energies of -332.0, -52.0, -48.9, -32.5, -28.7 kcalorie mol⁻¹, respectively and r.m.s. deviations from the idealised native conformation of 1.1, 3.4, 5.3, 6.3, and 12.4 Å respectively. The energy of the idealised native conformation is so high because it has not been minimised. A contact at the intersection of row i and column j indicates that residue i is within 10 Å of residue j . In these maps, helices feature a broadening of the diagonal (as from top left to bottom right), antiparallel β sheet as a band perpendicular to the diagonal, and parallel β sheet as a band running parallel to the diagonal.

Testing the simplified representation

The drastic simplifications used in the present representation of a protein conformation were tested by minimisation from near the native folded conformation. Bovine pancreatic trypsin inhibitor was chosen for this test as it is the only small protein (less than 100 residues) of known conformation that has a single polypeptide chain and no additional prosthetic groups. As a first step, a simplified native PTI conformation was obtained by taking the C α positions and side-chain centroids from the X-ray coordinates⁹ (kindly supplied by Drs Huber and Steigemann). Next an idealised chain, based on the PTI sequence and having the same geometry for all side chains of the same type, was made to fit the simplified native coordinates by adjusting the α torsion angles. This conformation, known as the idealised native structure, deviates by 1.1 Å r.m.s. from the simplified native structure. The r.m.s. deviation is

$$\left\{ \frac{1}{N} \sum_{i,j} (\Delta r_{ij})^2 \right\}^{1/2}$$

where Δr_{ij} is the difference, in the two structures, of the distance between side-chain centroids (i and j). Energy minimisation from this starting conformation was then carried out to reproduce the stability of native PTI.

After 558 cycles a perfect minimum is reached at an energy of -52.0 kcalorie mol⁻¹ and a r.m.s. deviation of only 3.3 Å from the simplified native conformation. Thermal randomisation about this minimum does not lead to further movement from the native molecule on subsequent minimisation. Randomly disturbing the initial best-fit angles (with a disturbance between -15° and +15°) has little effect on the conformation obtained by subsequent minimisation. Figure 2 compares minimum energy and native chain folding in stereo. Figure 3 compares the contact maps. Because of a general twisting of the molecule, the comparison of the two structures should be done in stereo. Because main-chain hydrogen bonds have been omitted, the terminal helix becomes distorted and consequently packs too tightly against the β hairpin centre at residue 27.

Simulation of PTI folding

Having shown that so simple a model can represent the stable conformation of a folded protein, we tried to simulate the actual process of folding. Most tests were done with two ex-




Fig. 4 Simulation of PTI folding (with $\alpha = 45^\circ$). Note that the simulation was restarted in the direction of randomisation in the plus and minus directions, which show the energy of the native structure. The last five cycles have energies of -46.0, -46.0, -46.0, -46.0, -46.0 kcalorie mol⁻¹ at the end of a minimum.

starting conformations: one extended and one folded, except for the terminal helix. The prediction of the native conformation. In the latter case, the prediction was used to guess the native conformation. Figure 4 shows the energy of the native conformation, raised to bend back on itself. At this point, the prediction then reaches a minimum that could be reached by randomisation of the prediction. The terminal helix at residue 27. More minimisation of the two top loops of the N-terminal helix of the folded conformation (Figure 3) bends back on itself. In the native conformation, the terminal helix is known to be within 10 Å. It is

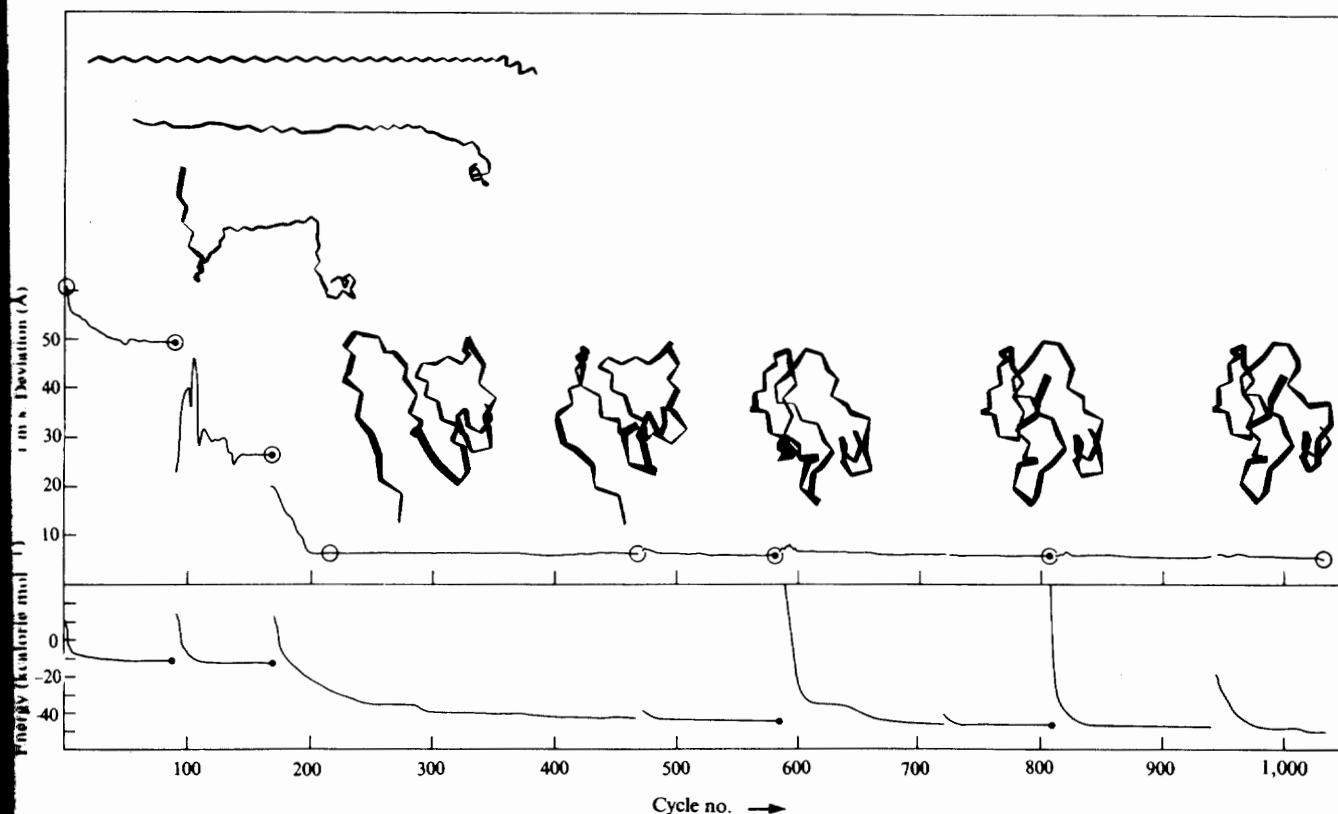


Fig. 4 Simulation of PTI folding from an extended starting conformation with the terminal helix ($\alpha = 180^\circ$ for all except 48 to 58 where $\alpha = 45^\circ$). No knowledge whatsoever about native PTI is used during this simulation (apart from setting the terminal helix). The conformation was thermalised at the end of each minimisation except near cycles 490 and 730 when the energy rises slightly because the minimisation was restarted after rounding the torsion angles to one degree. In the first two thermalisations, each normal mode was perturbed in the plus direction to raise the associated energy to $R(n)kT/2$ with $T = 1,000$ K. In the other three thermalisations, the perturbations were randomly in the plus and minus directions but always such as to raise the energy by $kT/2$ with $T = 300$ K. (Because the random numbers are distributed uniformly rather than exponentially, these temperatures do not correspond to the macroscopic temperature.) The 8 ribbon diagrams, which show the C^α chain path, refer from left to right to the 8 conformations at the circled points on the r.m.s. deviation curve, respectively. The last five conformations have progressively lower energies and are each a little closer to the native structure ($E = -43.3, -45.7, -46.0, -46.9$ and -48.9 kcalorie mol^{-1} , respectively; r.m.s. deviation = 6.08, 5.7, 5.6, 5.4 and 5.3 Å, respectively). The solid dots at the end of a minimisation indicate that a perfect minimum was reached (r.m.s. gradient less than 10^{-6} kcalorie $\text{mol}^{-1}\text{-rad}^{-1}$). One cycle takes about 0.6 s on an IBM 370/165 computer.

starting conformations, one fully extended (all $\alpha = 180^\circ$, and one extended apart from the C-terminal helix ($\alpha = 180^\circ$, except for residues 48 to 58 where $\alpha = 45^\circ$). Retaining the terminal helix from the native structure is justified here as we are more concerned with the process of folding than with prediction of the native conformation of an unknown protein. In the latter case a statistical rule (see ref. 2) could be used to guess the position of the α helices in the starting conformation. Figure 4 shows the iteration history of minimisation from the second of these starting points, which was the most successful run of those obtained to date. Thermal randomisation about the first minimum, an irregular but extended conformation, raises the energy and in this case causes the chain to bend back on to itself decreasing the r.m.s. deviation. From this point, minimisation first opens the molecule again, but then reaches a new minimum where the chain now has kinks that could become the bends of β hairpins. After a second randomisation, minimisation rapidly folds the molecule bringing the terminal helix close to the β hairpin centred on residue C. More minimisation and thermalisation first brings together the two top loops (near residues 15 and 40), and then brings the N-terminal tail on to the rest of the molecule. The final folded conformation of Fig. 4 is remarkably like the native molecule (Figs 2 and 3). In both conformations the chain bends back on itself near residues 14, 27, and 40. In both conformations, the pairs of half-cystine residues that are experimentally known to form S-S bridges, are close together (<10 Å). It is interesting that the C-terminal helix is the part

of the native molecule reproduced least well even though these residues had been set to a perfect helix in the starting conformation; this is due to the omission of peptide-peptide hydrogen bonds which stabilise the helix and could now be introduced.

Repeating the folding simulation from the fully extended starting conformation also lead to a compact structure with many of the features of native PTI, although after the same number of cycles the r.m.s. fit was a little worse (7.7 Å instead of 6.5 Å) and the energy was higher (-27.7 kcalorie mol^{-1} instead of -44.9 kcalorie mol^{-1}). Almost all the differences in conformation of these two folded structures involved the last 10 residues which remained extended if not pre-set to a helix and consequently failed to pack against the rest of the molecule.

Variation of folding conditions

Changing either set of starting torsion angles by a random value between -15° and 15° had little effect on the final conformation. Folding at a lower initial temperature ($T = 300$ K, rather than $T = 1,000$ K) failed to reach a compact conformation, as the thermal disturbances were too small to get out of the local minima corresponding to an extended chain.

Four additional runs of 600 cycles, under the same conditions as those used in Fig. 4 but based on different sequences of random numbers, gave rise to different folded shapes. In one of these, the folded molecule was close to the native structure (r.m.s. deviation of 6.3 Å), although the β sheet was formed between parallel rather than antiparallel chains (Fig. 3f). In

two others, the antiparallel β sheet centred near residue 27 was formed, but this hairpin did not subsequently fold on to itself to give a compact shape (Fig. 3*d* and *e*). Conformations that deviated more from the native structure always had higher energies, which gives an independent criterion for choosing the best conformation and suggests that more passes of thermalisation and minimisation lead to a conformation closer to the native one. Of the five runs using different random numbers for the thermalisation step, two succeeded in getting to within 6.5 Å of the simplified native structure in less than 600 cycles. That certain folding pathways are less successful is consistent with the experimental results of Creighton¹⁰ who has analysed the predominant kinetic intermediates present at different times after starting PTI renaturation and found several with the wrong tertiary fold.

General model for protein folding?

It seems remarkable that so simple a model based on time-averaged forces can account for the stability and folding of a molecule as complicated as a protein. Looking at known protein conformations closely, one is struck by the precise geometry of the interatomic contacts that stabilise the molecule: all possible interior hydrogen bonds are well formed, and many of the nonpolar side chains interlock to form a close-packed interior. As the forces responsible for this precise geometry fall off rapidly with distance and improper orientation, it would seem that folding must depend on a very rare random fluctuation that happened to bring the right residues close together with sufficient precision for the short-range forces to take effect. It therefore seems unlikely that these short-range forces could 'direct' the folding from an open disordered structure. In view of the present results, however, the time average of these short-range forces may play an important role in directing protein folding. These effective forces, which are weak, fairly long range, and not too dependent on orientation, restrict the number of low energy conformations severely; they cause the chain to fold into the approximate shape rapidly and without having to pass through many local minima. Because this approximately folded molecule corresponds to a large region in the space of possible protein conformations, folding would not be so rare an event.

As a general model for protein folding we propose that initially, when the chain has a flexible open structure, the effective time-averaged forces between the residues play a

central role, folding the chain into a compact shape with most groups close to their final positions (say within 5 Å). Once the chain becomes compact with less freedom of movement, the specific short-range interatomic forces become important; they form a precise conformation provided that the resulting gain in enthalpy overcomes the loss in entropy. The process would be rather like crystallisation, with the atoms simply falling into place from their nearby positions in the approximately folded conformation.

To simulate this second step one switches over to progressively more detailed models gradually incorporating more atoms and ending with the all-atom structures considered in earlier work^{12,13}. Although calculating the energy of the all-atom molecule would be time consuming, one would have the great advantage of starting close to the right conformation and could minimise successive overlapping zones of a few residues at a time without having to search through many local minima.

The general concept of using a simple model based on effective time-averaged forces when the detailed forces are too complicated has many potential applications; for example, the formation of protein quaternary structure and many enzyme complexes, virus assembly and so on. At each level of complexity, forces would be time averaged over those substructures that are relatively fixed or seem to play an important role in the assembly. Such a hierarchical approach might eventually lead to an understanding and simulation of very complicated biological assembly processes.

We thank the Weizmann Institute Computer Center for facilities and the European Molecular Biology Organisation for supporting one of us (M.L.).

Received October 29, 1974; revised January 15, 1975.

- 1 Anfinsen, C. B., *Science*, **181**, 223-230 (1973).
- 2 Schulz, G. E., Barry, C. D., Friedman, J., Chou, P. Y., Fasman, G. D., Frenkel, A. V., Lim, V. I., Ptitsyn, O. B., Kabat, E. A., Wu, T. T., Leach, R., Robson, B., and Nagano, K., *Nature*, **250**, 140-142 (1974).
- 3 Scheraga, H. A., in *Current topics in biochemistry* (edit. by Anfinsen, C. B., Schechter, A. N.), 1-42 (Academic, New York, 1974).
- 4 Ptitsyn, O. B., *Vestnik Akad. Nauk S.S.S.R.*, **5**, 57-68 (1973).
- 5 Flory, P. J., in *Statistical Mechanics of Chain Molecules*, 248-306 (Wiley, New York, 1969).
- 6 Hill, T. L., in *Statistical Mechanics*, 15-17 (McGraw-Hill, New York, 1956).
- 7 Nozaki, Y., and Tanford, C., *J. biol. Chem.*, **246**, 2211-2217 (1971).
- 8 Simon, E. M., *Biopolymers*, **10**, 973-989 (1971).
- 9 Huber, R., Kukla, D., Ruhlmann, A., and Steigemann, W., *Cold Spring Harbor Symp. quant. Biol.*, **36**, 141-148 (1971).
- 10 Creighton, T. E., *J. molec. Biol.*, **87**, 603-624 (1974).
- 11 Phillips, D. C., in *British Biochemistry, Past and Present* (edit. by G. D. F. Williams), 11-28 (Academic, London, 1970).
- 12 Levitt, M., and Lifson, S., *J. molec. Biol.*, **46**, 269-279 (1969).
- 13 Warne, P. K., and Scheraga, H. A., *Biochemistry*, **13**, 757-767 (1974).

letters to nature

Two kinds of stellar collapse

HERE I shall show that the events which produce compact objects may be divided into at least two classes on the basis of the impulse given to a binary system containing the collapsing object. This impulse may be related to the amount and symmetry of mass loss resulting from the collapse event.

If one member of a binary system in circular orbit loses mass in a spherically symmetric manner and on a time scale much less than the orbital period, then the barycentre of the remaining system acquires the velocity¹

$$S_g = \frac{M_1 M_2 (1-q)}{(M_1 + M_2)(qM_1 + M_2)} S_0 \quad (1)$$

where S_0 is the initial relative orbital velocity of the two masses,

M_1 and M_2 are the initial masses and qM_1 and M_2 are final masses. If more than half of the total initial mass of binary is lost the system becomes unbound; if not, the orbit becomes eccentric¹. If the initial orbit is eccentric, the explosion is asymmetric, or the effects^{2,3} of the collision of the debris with the companion star are included, then several more parameters are needed to describe the event and to determine the new orbit. Conservation of momentum arguments indicate that explosions producing loss of a few tenths of the binary mass, or giving the compact remnant a recoil velocity of several tenths of S_0 will generally (except for fortuitous combinations of parameters) lead to orbits of substantial eccentricity. If $M_1 \sim M_2$ then barycentre recoil velocities of a few tenths of S_0 will also result. Slightly more violent events produce hyperbolic orbits.

I first summarise the evidence for the expulsion of large