Global Optimization Methods for Protein Folding Problems

Richard H. Byrd, Elizabeth Eskow, André van der Hoek, Robert B. Schnabel, Chung-Shang Shao, and Zhihong Zou

ABSTRACT. The problem of finding the naturally occurring structure of a protein is believed to correspond to minimizing the free, or potential, energy of the protein. This is generally a very difficult global optimization problem, with a large number of parameters and a huge number of local minimizers including many with function values near that of the global minimizer. This paper presents a new global optimization method for such problems. The method consists of an initial phase that locates some reasonably low local minimizers of the energy function, followed by the main phase that progresses from the best current local minimizers to even lower local minimizers. The method combines portions that work on small subsets of the parameters, including small-scale global optimizations using stochastic methods, with local minimizations involving all the parameters. In computational tests on the protein polyalanine with up to 58 amino acids (116 internal parameters), the method appears to be very successful in finding the lowest energy structures. The largest case is particularly significant because the lowest energy structures that are found include ones that exhibit interesting tertiary as opposed to just secondary structure.

1. Introduction

The naturally occurring three dimensional structure of a protein, called its "tertiary structure", is believed to be uniquely determined from its "primary structure", the sequence of amino acids of which the protein is composed. The problem of finding the tertiary structure given the primary structure is known as the protein folding problem. Scientists believe that the tertiary structure of a protein is the structure that minimizes its free, or potential energy. Thus the protein folding problem can be posed as an optimization problem. This optimization problem is an extremely challenging global optimization problem for even moderately sized proteins, because the number of optimization parameters for such problems is 100 or more, and the number of local minimizers is believed to be exponential in the number of parameters. Moreover, many of the local minimizers have function values near that of the global minimizer.

This paper presents a new global optimization method for protein folding problems, and computational results for problems with up to 116 parameters. This problem of minimizing potential energy is an instance of the general global optimization problem of finding the lowest minimizer of a nonlinear function $f(x)$ that has multiple local minima, for x in a domain D defined by lower and upper bounds on each parameter x_i. The functions $f(x)$ used in this research are well known empirical potential energy functions which are described in more detail in later sections.

A major challenge in solving large-scale global optimization problems is that there is no clean mathematical basis for efficiently reaching a global minimizer, such as an analog to steepest descent or Newton's method for local minimization. Many methods have been developed for such global optimization problems (see the survey [17]), and these therefore tend to be heuristic and often require large amounts of computation time. This is true of the method presented here as well. We do aim in this research, however, to develop a global optimization methodology that can be applied to a wide range of large-scale problems, albeit with some application-dependent modifications. We also aim to assure that our methods can effectively utilize powerful parallel computers, although this aspect of our research is not addressed in this paper.

The new method described in this paper has the same top level structure as the global optimization methods for molecular cluster problems that we introduced in [4, 2]. These methods have two main phases. The first, initialization phase generates a set of configurations that are local minimizers of the energy function and have reasonably low energy values. The second phase, which accounts for the overwhelming majority of the computational effort, is an improvement phase that successively finds new and hopefully lower local minimizers from the current lowest local minimizers. The key feature of each step of this phase is the solution of a small-scale global optimization problem (typically with 3-9 parameters) by a stochastic global optimization method. That is, the relatively well developed methodology for small-scale global optimization problems is utilized within the large-scale global optimization algorithm. At the level described above, the methods for molecular cluster problems and the new methods for protein folding are similar; beneath this level they differ significantly.

The remainder of the paper motivates and describes our new global optimization method. Section 2 briefly summarizes our global optimization algorithm for molecular cluster problems, which is the starting point for our new method, and also very briefly summarizes our computational experience with this algorithm to justify that it appears to be a viable approach. Section 3 present our global optimization method for protein folding problems and computational results on the protein polyalanine. Section 4 gives some conclusions and directions for future research in this area.

2. Global Optimization Methods for Molecular Cluster Problems

The basic strategy of our algorithm was developed first to solve global optimization problems for finding the configurations of macro-molecular clusters. The approach was applied to potential energy functions for two different cluster problems. The first problem was that of identical atoms using the Lennard-Jones potential energy function, and the second problem involved clusters of water molecules using

the Coker/Watts energy function [6]. These molecular cluster problems have some of the same challenges as the protein folding problem, including many parameters, a huge number of local minimizers, and many low local minimizers. In addition, the energy functions for both cluster and protein problems include $O(N^2)$ 2-body energy terms, composed of similar functional forms, for an N atom problem. The characteristic that makes molecular cluster problems much different from protein folding problems is the lack of the chain structure that is inherent in proteins. This enables the movement of a single atom or molecule in the structure to a completely different location without compromising the integrity of the structure as would be the case in a chain of amino acids.

Lennard-Jones problems assume that the potential energy of the cluster is given by the sum of the pairwise interactions between atoms, with these interactions being Van der Waals forces given by the Lennard-Jones 6-12 potential. That is, if we define the position of the molecular cluster by

$$x = (x_1, x_2, ..., x_N)$$

where x_i is a three-dimensional vector denoting the coordinates of the i^{th} atom, then the potential energy function is

$$(2.1) \qquad f(x) = \sum_{i=1}^{N} \sum_{j=1}^{i-1} \left[\frac{1}{d_{ij}^{12}} - 2 * \frac{1}{d_{ij}^{6}} \right]$$

where d_{ij} is the Euclidean distance between x_i and x_j. In this formulation, the pairwise equilibrium distance is scaled to one and the pairwise minimum energy is scaled to -1. Lennard-Jones problems characterize the interaction of inert gasses such as argon; more importantly, Lennard-Jones terms are a key part of many empirical energy models including all commonly used energy functions for proteins.

The potential energy of water clusters that we use was developed by Coker and Watts [6], and has also been used in global optimization research by Long [14]. It has the form

$$(2.2) \qquad f(x) = \sum_{i \neq j} u(x_i, x_j) + \sum_i v(x_i)$$

where each vector x_i gives the coordinates of three atoms in the i-th molecule, and the functions u and v are composed of a number of terms that give the interaction energy between pairs of molecules and the internal energy of the molecule, respectively. Note that modeling water will be important in future research on protein folding, which will need to consider the behavior of proteins in a water solvent.

The basic framework of our global optimization methods for molecular cluster problems is outlined in Algorithm 2.1 below. As mentioned in Section 1, the method combines an initial phase that locates some low local minimizers with a second phase for progressing from low to even lower local minimizers that accounts for most of the computational effort, and the success, of the method. Both phases make use of techniques that vary only a small subset of the variables (an atom for Lennard-Jones problems, a molecule for water) at once. In the initial phase this approach is used to improve the sample points one atom or molecule at a time in step 1b. In the second phase, this approach is used to move one atom or molecule in an existing configuration to new positions via the one atom/molecule global optimization in step 2c. This global optimization finds the best possible position for the selected atom or molecule in the current configuration with the remainder

of the configuration temporarily fixed. In both phases, the small subset of the variables that is selected is the one that contributes the most to the energy function (a quantity that is readily calculated due the partial separability of the energy function), an indicator that moving this atom or molecule may be a fruitful way to reduce the overall energy of the cluster. Both of the one atom/molecule steps are relatively inexpensive due to the small number of variables involved and the partial separability of the energy function, which allows the energy to be recalculated at $O(1/N)$ of the cost of a full energy evaluation if only one atom or molecule is moved. In both phases, the one atom/molecule steps are followed by local minimizations in all the variables; these steps allow the entire configuration to change. This overall approach has been found to lead to improvements in the sample points and local minimizers far more effectively and efficiently than strategies that always work on all the variables at once.

The overall progress of the improvement phase is governed by the heuristic in step 2a. This heuristic has been developed to explore the enormous search space of possible configurations in a manner that combines breadth and depth. Configurations that were passed to Phase 2 from Phase 1 are considered to be roots of trees that first are each grown to a specified depth, regardless of the energy function values produced at each consecutive level. After this initial balancing portion in which all of the trees have been explored to an equal (specified) depth, selections of configurations are based solely on energy criteria, which tends to encourage only a few trees (with the lowest energy values) to continue growing. At both stages, the energy values do not necessarily decrease monotonically as one descends a tree of configurations, since the best new configuration resulting from a given configuration may not necessarily have a lower energy value. The balancing of work over all of the initial configurations before working exclusively on the best few has proven to be an extremely beneficial strategy in the overall success of the algorithm.

A final important feature of Algorithm 2.1 is the expansion step 2b. It was found, particularly for water clusters, that by expanding the cluster prior to the one-molecule global optimization step, the small global optimization could locate better new positions for the molecule and the overall algorithm could make better progress. For Lennard-Jones problems this step is sometimes helpful as well. There is no analog to this step in our current protein methods.

2.1. Results for Cluster Problems. There has been extensive development of and experimentation with very special purpose methods for Lennard-Jones problems [12, 16, 20], and many researchers have applied more general approaches to these problems as well [7, 15, 13, 18]. As a result, the optimal configurations are believed to be known for most clusters with up to 100 or so atoms. Algorithm 2.1 has been applied to all Lennard-Jones clusters with 5 to 76 atoms (15 to 228 parameters), and finds the best known solutions in all cases. A new best minimizer was found for the 75 atom problem, with an energy value of -396.282, whereas the best previously reported value for this case was -396.239 [20]. Recently discovered lowest energy values for 66 and 72 cases [7] were also matched by Algorithm 2.1. These results appear to be the best currently reported by any global optimization method on Lennard-Jones problems.

For the water problem, we have mainly run our algorithm on clusters of 20, 21 and 32 water molecules (180, 189, and 288 parameters), because results of minimizing these same clusters and energy function, using a dynamic simulated annealing

Algorithm 2.1 – Framework of the Large-Scale Global Optimization Algorithm for Molecular Cluster Problems

1. **Initial Generation of Configurations Phase :**
   (a) **Sampling in Full Domain :** Randomly generate the coordinates of the sample points in the sampling domain, and evaluate $f(x)$ at each new sample point. Discard all sample points whose function value is below a global "cutoff level".
   (b) **One-atom/molecule Sampling Improvement :** For each remaining sample point : While the energy of the sample point is above the threshold value, Repeat:
       • Select the atom/molecule that contributes most to the energy function value
       • Randomly sample on the location of the selected atom/molecule
       • Replace this atom/molecule in the sample point with the new sample coordinates that give the lowest energy value.
   (c) **Start Point Selection :** Select a subset of the improved sample points from step 1b to be start points for local minimizations.
   (d) **Full-Dimensional Local Minimizations :** Perform a local minimization from each start point selected in step 1c. Collect some number of the best of these minimizers for improvement in Phase 2.

2. **Improvement of Local Minimizers Phase:** For some number of iterations:
   (a) **Select a Configuration :** From the list of full-dimensional local minimizers, select the local minimizer and atom/molecule to be optimized.
   (b) **Expansion :** Transform the configuration by multiplying the position of each atom relative to the center of mass of the configuration by a constant factor of between 1.0 (no expansion) and 1.75.
   (c) **One Atom/Molecule Global Optimization :** Apply a global optimization algorithm to the expanded configuration with only the atom/molecule chosen in step 2a as a variable.
   (d) **Full-Dimensional Local Minimization :** Apply a local minimization procedure, using all the atoms/molecules as variables, to the lowest configurations that resulted from the one-atom/molecule global optimization.
   (e) **Merge the New Local Minimizers :** Merge the new lowest configurations into the existing list of local minimizers.

procedure, have been obtained by Long [14]. We have obtained many configurations with significantly lower energies than those obtained in [14]. The best solutions obtained by running our algorithm have energies of -0.348183, -0.369011 and -0.585062 atomic units (a.u.) for 20, 21 and 32 molecules respectively. These energy values are approximately 0.005, 0.01 and 0.02 a.u. lower than the best structures found in [14], whereas at room temperature, only vibrational states with energies about 0.001 a.u. above the ground state are possible. The structures for 20 and 21 molecules have collapsed dodecahedral (for 20) and dodecahedral (for 21) shapes, which are in agreement with the shapes of the lowest energy clusters found in [14].

These results appear to indicate that the approach of Algorithm 2.1 is a useful one for the global optimization of energy functions for molecular clusters.

3. Methods for Protein Conformation Problems

The problem of finding the native three dimensional structure of a protein given its primary sequence of amino acids differs from the molecular cluster problems discussed in the previous section, due mainly to the inherent chain-like structure of proteins. Before expounding on the ramifications of this structural difference in our methods, we address some fundamental issues regarding the protein problem.

The basic building block of proteins is an amino acid. The amino acids are bonded together by peptide bonds to form a polypeptide chain, and the sequence of amino acids in the chain is referred to as the primary sequence of the protein. There are 20 amino acids that exist in nature, and all proteins are composed from these. The secondary structure of a protein refers to regular structure in portions of the polypeptide chain, such as alpha helices, beta sheets and the turns that link them, whereas tertiary structure describes the overall shape adopted by the polypeptide chain within a domain [8]. The tertiary structure is believed to be the structure that minimizes the free, or potential, energy of of the protein.

Chemists have developed a number of empirical potential energy functions to describe the interaction between atoms in proteins. In this study we use a version of the CHARMM energy function [1] to compute the potential energy. This function depends on the coordinates of the atoms in the protein and has the form

$$(3.1) \qquad f(x) = \sum_{i,j} E_{bond} + \sum_{i,j,k} E_{angle} + \sum_{i,j,k,l} E_{dihedral} + \sum_{m,n} (E_{VdW} + E_{elec})$$

. The first term, the bond length potential, is a sum over all pairs of of bonded atoms i and j, and depends on the deviation of the distance $\|x_i - x_j\|$ from its equilibrium value. For each three atoms such that atoms i and k are both bonded to atom j, the second sum has a term depending on the bond angle formed by the three atoms, with vertex at x_j. The last bonded summation includes terms depending on certain dihedral angles formed by sets of four atoms connected by bonds. The dihedral angle is formed by two planes each containing a specified set of three of the four atoms, and is a measure of the torsion of the configuration. The final summation in (3.1) runs over all pairs of atoms not involved in bonds. It includes the Van der Waals (Lennard-Jones 6-12 as discussed above)) interactions and the electrostatic potential between pairs of atoms.

The CHARMM function computes potential energy as a function of the cartesian coordinates of the atoms, but it is convenient to express the protein configuration in terms of an equivalent set of internal parameters given in terms of bond lengths, bond angles and dihedral angles. This parameterization is useful for global optimization because the dihedral angles are the crucial parameters to be varied in the optimization. Since the energy to perturb bond lengths and bond angles from their equilibrium values is relatively large, it is conventional in the bio-chemistry community to consider these parameters as fixed, and to address the global optimization problem with just the dihedral angles as the variables, and this is how we have formulated the problem. The main advantage of this parameterization is that the number of variables is greatly reduced, generally by a factor of ten or more. A disadvantage is that the energy function is no longer partially separable in this parameterization, meaning that it is no longer much less expensive to re-evaluate

Algorithm 3.1 – Framework of the Large-Scale Global Optimization Algorithm for Protein Problems

1. **Initial Generation of Configurations Phase :**
   (a) **Protein Sample Point Buildup:** Build up sample configurations from one end of the protein to the other by sequentially generating each dihedral angle in the protein: randomly sample the current dihedral angle a fixed number of times and select the dihedral angle that gives the lowest energy function value for the partial protein generated so far.
   (b) **Start Point Selection :** Select a subset of the best sample points from step 1a to be start points for local minimizations.
   (c) **Full-Dimensional Local Minimizations :** Perform a local minimization from each start point selected in step 1b. Collect some number of the best of these minimizers for improvement in Phase 2.

2. **Improvement of Local Minimizers Phase :** For some number of iterations:
   (a) **Select a protein:** From the list of full-dimensional local minimizers, select a local minimizer and a small subset of dihedral angles from that minimizer to be optimized.
   (b) **Global Optimization on a small subset of variables :** Apply a fairly exhaustive small-scale global optimization algorithm to the energy of the selected configuration using the selected small subset of the dihedral angles as variables.
   (c) **Full-Dimensional Local Minimization :** Apply a local minimization procedure, with all dihedral angles as variables, to the lowest configurations that resulted from the global optimization of the step 2b.
   (d) **Merge the New Local Minimizers :** Merge the new lowest configurations into the existing list of local minimizers.

the energy if only a few parameters change than if they all change. This loss of partial separability has some ramifications for our algorithmic strategy that are mentioned below.

The framework of our global optimation method for protein problems is outlined in Algorithm 3.1 below. The highest level structure of the algorithm is the same as in the molecular cluster algorithm. It consists of an initialization phase that locates some low minimizers, followed by an improvement phase that finds better local minimizers from the current best local minimizers. As in the molecular cluster algorithm, the improvement phase utilizes a small-scale global optimization on a selected subset of the parameters followed by local minimizations on all the parameters. Beyond this level there are many differences between the two algorithms due to the chain structure of proteins.

One main difference of Algorithm 3.1 from the molecular cluster algorithm is in the sampling improvement strategy of the first phase, where the cluster methods move one particle at a time to random locations in order to improve the sample configurations. Because it is not possible to move an arbitrary atom of a chain without affecting the bond lengths and angles attached to neighboring atoms, the sample generation for proteins is done entirely differently. The chain is built up one

dihedral angle at a time, by sampling dihedral angles in a sequential manner along the chain. Each dihedral angle is sampled a fixed number of times, and the value that gives the configuration under construction the best partial energy is selected. This initialization step seems to have worked well, and overcomes the lack of partial separability since the change in energy due to adding on to the end of a chain can be calculated efficiently. It is not the crucial part of the method, however, since it just supplies some starting configurations for the main, improvement phase. In fact, if a scientist knows some good candidate configurations, these could be used as input to the improvement phase, instead of or in addition to the ones supplied by the initialization phase.

The other key difference from the molecular cluster algorithm is that the small-scale global optimization in the improvement phase of the method can no longer select and move a single atom or molecule, for the reasons mentioned above, but instead optimizes over a small number of dihedral angles with the remaining dihedral angles temporarily fixed. The selection of the angles to optimize over is very important to the success of the method, and involves interesting new tradeoffs. Selecting a small number of dihedral angles that are dispersed throughout the protein chain offers the greatest possibility for change in the tertiary structure of the protein during the small-scale global optimization. However it means that the small-scale global optimization is relatively expensive because the energy evaluations are nearly as expensive as full energy evaluations. Conversely, selecting a small contiguous set of dihedral angles offers less possibility for change in the overall structure of the protein during the small-scale global optimization, but it allows the small-scale global optimization to be relatively inexpensive because the energy evaluations can performed efficiently, as in the small subproblems for molecular clusters. So far, our algorithms have been based on the first option, selecting a small, dispersed set of dihedral angles, but we intend to experiment with the second possibility as well.

Given this overall strategy of choosing a small set of possibly dispersed dihedral angles as the parameters for the small-scale global optimization, our algorithm must determine which angles to select. The method for doing this is different and more difficult than the analogous selection for cluster problems. Once again, we wish to use the energy function to indicate which parameters, if varied, offer the greatest potential for reduction in the energy function. But we can no longer calculate this directly by utilizing partial separability to compute the partial energy associated with each parameter. Instead, Algorithm 3.1 uses two different methods for selecting the parameters to optimize, both based upon the interaction energies between all the atoms to the left of a given dihedral angle and all those to its right. The methods differ in their normalization of the interaction energies. The first method computes the interaction energy just described for each possible dihedral angle, and normalizes it by the product of the number of atoms to the left times the number to the right. Some specified number (generally five) of dihedral angles with the highest normalized interaction energies are then selected. The second method computes the interaction energy in the same manner as the first, but normalizes it by the maximum of the number of atoms to the left and right. The first method exhibits a bias toward selecting dihedral angles in the middle part of the conformation, whereas the second method tends to pick dihedral angles closer to the ends. The first method appears to be more effective in the earlier stages of the computation when the conformation is still far from the global minimum, while

the second method appears appropriate for polishing the ends of the conformation after the middle part has been optimized.

In our experiments with the 58 amino-acid polyalanine mentioned in Section 3.1, we have also used a third method for selecting the dihedral angles for the small-scale global optimization. This method, used only in the last stages of the calculation, is to chose those dihedral angles that are close to where a hydrogen bond is expected to be but is not currently present in the conformation. So far, this choice has been made manually by visualizing the protein, although it should be possible to make it algorithmically. Of course, including this type of heuristic in the code makes it a less general-purpose global optimization method, but may be important in its ability to solve difficult protein problems. We discuss this issue futher in the section on future research.

Finally, the heuristics used in step 2a of the cluster method to balance the breadth and depth of the search have proven to be valuable in searching the space of possible protein conformations as well, and are used in essentially the same manner.

3.1. Results for Protein Problems. We have run Algorithm 3.1 on the protein polyalanine with 20, 30, 40 and 58 residues (amino acids). Polyalanine is a simple polypeptide in that each amino acid is the same and the side chain in the amino acid is simple and contributes no additional dihedral angles to the parameterization. It has been used commonly as a test problem in the bio-chemistry community (see e.g. [**19, 10**]). In internal parameters, there are two parameters per amino acid, meaning that the problems tested have up to 116 optimization parameters. Our calculations have mainly been performed on a KSR1 multiprocessor using up to 64 processors, using a parallelization strategy similar to that described in [**3**].

In the problems with 20, 30, and 40 amino acids, the best structure found by our algorithm in each case is a very regular alpha-helix. These results are expected, although we have no comparative energy values from other methods using the same input data and the CHARMM energy function to check against. Interestingly, in the 20 amino acid problem the initialization phase already finds the best configuration, although the improvement phase can also reach this configuration from poor starting configurations. In the 30 and 40 amino acid problems, the initialization phase does not produce the best configuration but it is found by the improvement phase.

The 58 residue problem is considerably more interesting. There are conflicting results in the literature regarding what shape to expect for this problem. Wilson and Cui [**19**] found alpha-helical structures for polyalanine with up to 80 residues using a simulated annealing method; however, Head-Gordon and Stillinger [**10**] show that for 58 residue polyalanine, a lower energy minimizer exists whose structure is two parallel alpha-helices connected by a turn. Our work on this problem is still in progress, but interestingly, so far our algorithm has found two very low energy minimizers with almost identical energies, and with the two different shapes just mentioned. These configurations were located using the normalized interaction energy strategies mentioned in the discussion of Algorithm 3.1. Following this, we have been able to further improve the energies of both of these conformations using the third technique for selecting parameters for the small-scale global optimization mentioned above, that is visually choosing dihedral angles near missing hydrogen

bonds. At this point, the best straight conformation appears to be a totally regular alpha helix with all of the expected hydrogen bonds present. The best bent conformation has a somewhat lower energy value than the best straight conformation, and its shape has changed such that the two alpha helices are no longer parallel, but are almost crossing over each other. This conformation has a potential energy about 4 kcal/mole lower than that of the best straight configuration. Which of these two ultimately is "best" is not the key issue, and in fact it appears that this may vary if one uses different potentials or cartesian versus dihedral angle parameterizations. The most exciting aspect of these results is that they indicate that our algorithm has the potential to locate complex tertiary structures without knowing of them apriori.

4. Conclusions and Future Results

We have presented a new global optimization algorithm for locating the minimum energy configurations of proteins. The algorithm utilizes small-scale global optimization calculations on selected subsets of the parameters, performed by a stochastic global optimization method, as a key part of its approach. Its structure is related to our previous global optimization methods for molecular clusters but there are many important differences due to the chain structure of proteins. In computational tests on the protein polyalanine with up to 58 amino acids (116 internal parameters), the method appears to be very successful in finding the lowest energy structures. The largest case is particularly significant because the lowest energy structures that are found include ones that exhibit interesting tertiary as opposed to just secondary structure.

The research reported here is part of an ongoing project that is continuing in many directions. One of these is the extension of the methodology presented here to more complex proteins composed of different amino acids and more complex side chains. This raises a number of new algorithmic issues. A second is the exploration of new strategies for choosing the parameters to optimize over in the small-scale global optimizations. This includes the option discussed in Section 3 of selecting a set of consecutive dihedral angles, as well as automatic strategies that correspond to the use of visual information discussed at the end of Section 3. Another direction is exploring ways to incorporate additional partial information that scientists have about the structure of proteins. For example, scientists appear able to predict the secondary structure of portions of proteins with high but not perfect accurary [11, 5], and it would seem useful to be able to utilize these predictions in the global optimization algorithm in some manner.

References

[1] B. R. Brooks, R. E. Bruccoleri, B. D. Olafson, D. J. States, S. Swaminathan and M. Karplus, *A program for Macromolecular Energy, Minimization, and Dynamics Calculations*, J. Comp. Chem. **4** (1983), pp. 187–217.

[2] R. H. Byrd, T. Derby, E. Eskow, K. Oldenkamp and R. B. Schnabel, *A New Stochastic/Perturbation Method for Large-Scale Global Optimization and its Application to Water Cluster Problems*, in Large-Scale Optimization: State of the Art, W. Hager, D. Hearn, and P. Pardalos, eds., Kluwer Academic Publishers, Dordrecht, The Netherlands, 1994, pp. 71–84.

[3] R. H. Byrd, E. Eskow, A. van der Hoek, R. B. Schnabel, K. P. B. Oldenkamp, *A Parallel Global Optimization Method for Solving Molecular Cluster and Polymer Conformation Problems*, Proceedings of the Seventh Siam Conference on Parallel Processing for Scientific Computing,

D. H. Bailey, P. E. Bjørstad, J. R. Gilbert, M. V. Mascagni, R. S. Shreiber, H. D. Simon, V. J. Torczon, L. T. Watson, eds., SIAM, Philadelphia, 1995, pp. 72–77.

[4] R. H. Byrd, E. Eskow and R. B. Schnabel, *A New Large-Scale Global Optimization Method and its Application to Lennard-Jones Problems*, University of Colorado Technical Report CU-CS-630-92.

[5] P. Prevelige Jr. and G. D. Fasman, *Chou-Fasman Prediction of the Secondary Structure of Proteins*, in Prediction of Protein Structure and the Principles of Protein Conformation, G. D. Fasman, ed., Plenum Press, New York, 1989, pp. 391–416.

[6] D. F. Coker and R. O. Watts, *Structure and Vibrational Spectroscopy of the Water Dimer Using Quantum Simulation*, J. Phys. Chem.,**91**, (1987), pp.2513–2518.

[7] T. Coleman, D. Shalloway, and Z. Wu, *A Parallel Build-up Algorithm for Global Energy Minimizations of Molecular Clusters Using Effective Energy Simulated Annealing* Technical Report, Center for Theory and Simulation in Science and Engineering, Cornell University, 1993.

[8] N. J. Darby, and T. E. Creighton, *Protein Structure*, Oxford University Press, New York, 1993.

[9] C. J. Epstein, R. F. Goldberger and C. B. Anfinsen, *The genetic control of tertiary protein structure: Studies with model systems.* Cold Spring Sympos. Quant. Biol.**28** (1963), pp. 439–449.

[10] T. Head-Gordon and F. H. Stillinger, *Enthalpy of Knotted Polypeptides*, J. Phys. Chem., **96**, (1992), pp. 7792–7796.

[11] T. Head-Gordon and F. H. Stillinger, *Optimal neural networks for protein-structure prediction*, Phys. Rev. E., **48**, (1993), pp. 1502-1516.

[12] M. R. Hoare and P. Pal, *Physical cluster mechanics: statics and energy surfaces for monatomic systems*, Adv. Phys., **20**, (1971) pp. 1196–196.

[13] J. Kostrowicki, L. Piela, B. J. Cherayil and H. A. Scheraga, *Performance of the diffusion equation method in searches of optimum structures of clusters of Lennard-Jones atoms*, J. Phys. Chem., **95**, (1991) pp. 4113–4119.

[14] X.Long, *Molecular Dynamics Simulations of Clusters – Impure van der Waals and e^- – $(H_2O)_n$ Systems*, Ph.D. Thesis, Department of Chemistry, University of California, San Diego, 1992.

[15] C. D. Maranas and C. A. Floudas, *A global optimization approach for Lennard-Jones microclusters*, J. Chem. Phys., **97** (1992) p. 7667.

[16] J. A. Northby, *Structure and binding of Lennard-Jones clusters:$13 \leq N \leq 147$* J. Chem. Phys., **87**, (1987) pp. 6166–6177.

[17] P.M. Pardalos, D. Shalloway and G.L. Xue, *Optimization Methods for Computing Global Minima of Nonconvex Potential Energy Functions*, Journal of Global Optimization, **4** (1994), pp. 117-136.

[18] D. Shalloway, *Packet Annealing: A Deterministic Method for Global Minimization: Application to Molecular Conformation*, in Recent Advances in Global Optimization, C. Floudas and P. Pardalos (eds.), Princeton University Press, Princeton, N.J., (1992) pp. 433-477.

[19] S. R. Wilson and W. Cui, *Applications of Simulated Annealing to Peptides*, Biopolymers, **29**, (1990), pp. 225–235.

[20] G. L. Xue, *Improvement of the Northby Algorithm for Molecular Conformation: Better Solutions*, J. Global Optimization, **4**, (1994), p. 425.

DEPARTMENT OF COMPUTER SCIENCE, UNIVERSITY OF COLORADO AT BOULDER, CAMPUS BOX 430, BOULDER, CO 80309-0430