Chapter 2

Local elevation: A method for improving the searching properties of molecular dynamics simulation

2.1 Summary

The concept of memory has been introduced into a molecular dynamics algorithm. This was done so as to persuade a molecular system to visit new areas of conformational space rather than be confined to a small number of low energy regions. The method was demonstrated on a simple model system and the 11 residue cyclic peptide Cyclosporin A. For comparison, calculations were also performed using simulated temperature annealing and a potential energy annealing scheme. Although the method can only be applied to systems with a small number of degrees of freedom, it offers the chance to generate a multitude of different low energy structures where other methods only give a single or few. This is clearly important in problems such as drug design where one is interested in the conformational spread of a system.

2.2 Introduction

Dynamic simulation methods are very powerful tools for refinement of molecular structures with respect to the potential energy. In general, one has a penalty function, which may represent the potential energy of a molecule as a function of conformation or perhaps is a measure of how well a conformation fits an experimentally measured property. If one can calculate the derivative of this function with respect to coordinates, one should be able to simply run a classical dynamics scheme until all low energy conformations are visited. In practice this idea breaks down for several reasons. First, the conformational space is vast. Even a 10 residue cyclic peptide has at least 10 important degrees of freedom, so one has a very high-dimensional function surface to search on. Second, the penalty function is usually very complex with many barriers. All but the smallest are surmounted very infrequently. This in turn leads to the tendency to repeatedly visit a small set of minima.

There have been many approaches to the multiple minimum problem [1]. The most commonly applied is probably simulated annealing [2], which relies on temperature evolution in time. Potential energy annealing (PEACS) [3] attempts to take advantage of the shape of the potential hypersurface. Several other methods exist, which rely on temporarily increasing the dimensionality by adding artificial degrees of freedom [4–8] Finally, there are systematic search methods [9] which soon become inapplicable as the system size grows, because of the astronomical computational effort.

Our approach here – in some way related with the tabu search technique [10,11] – is rather more radical and attempts to specifically address the problem of repeatedly visiting certain areas of conformational space. Somewhat as a side effect, this also has an influence on the crossing of energy barriers. To avoid resampling conformations, a kind of learning process is employed so as to recognise conformations sampled before. When a previously visited conformation is recognised, it is easy to penalise the system and make this region of conformational space less energetically attractive. The advantage of this idea is that the system is kept in low energy regions and will only be driven to high energy areas if no alternative is available. Thus, little time is spent searching in high entropy high energy areas and after being driven across a barrier, the system will soon return to low energy areas.

2.3 Theory

For penalising visited conformations, we used a Gaussian function

$$V_{mem}(\chi) \;=\; k_{mem}\; n_{\chi^0}\; \exp(-\frac{(\chi-\chi^0)^2}{2w^2}) \tag{2.1}$$

with χ the current conformation, $\chi^{\mathbf{0}}$ a conformation visited before, n_{χ^0} the number of times this conformation has been sampled before (including the dimension) k_{mem} (> 0) defining the magnitude and w the width of the memory penalty function. If χ is expanded to a set of variables x_i, one causes the function to decrease very quickly with small conformational changes, because the penalty function has the form of a product over a set of Gaussians for independent variables. This can be seen by rewriting equation (2.1),

$$V_{mem}(\chi) \;=\; k_{mem}\; n_{\chi^0}\; \exp(-\frac{(\chi-\chi^0)^2}{2w^2}) \tag{2.2}$$

$$\begin{aligned} V_{mem}(\{x_k\}) &= k_{mem}\; n_{\chi^0}\; \exp(-\frac{\sum_i (x_i - x_i^0)^2}{2w^2}) \\ &= n_{\chi^0} \prod_i k'_{mem}\; \exp(-\frac{(x_i - x_i^0)^2}{2w^2}) \end{aligned} \tag{2.3}$$

Next, to put equation (2.1) into a useful form, we chose dihedral angles to define the conformation. With ϕ denoting a dihedral angle, the penalty function (2.3) can be rewritten

$$V_{mem}(\{\phi_k\}) \;=\; n_{\chi^0} \prod_i k'_{mem}\; \exp(-\frac{(\phi_i - \phi_i^0)^2}{2\; w^2}) \tag{2.4}$$

and it is straightforward to calculate the partial force on particle j.

$$\mathbf{F}_{\mathbf{r_j}} \;=\; \sum_i V_{mem}(\{\phi_k\})\; \frac{(\phi_i - \phi_i^0)}{w^2} \cdot \frac{\partial}{\partial \mathbf{r_j}} \phi_i \tag{2.5}$$

At the heart of the method is the memory of the system's history. For this purpose, a storing/comparing routine has been developed for conformations keeping two points in mind. First, it must be able to handle a large set of stored conformations and second, the loss of speed compared to a free simulation (without memory) should be acceptable. Both criteria can be met by using data reduction and packing. Figure 2.1 shows the flow diagram of the so called local elevation (LE) method. First the complete data describing a conformation are reduced to the relevant set of dihedral angles. After further reduction of the precision, the dihedral angle values are packed into integers. Bearing in mind that in peptides only a few well separated minima of a dihedral angle are usual, a maximum number of 16 states per dihedral angle is sufficient to distinguish real conformational changes. Given only 16 states per dihedral angle, one needs only 4 bits of information per angle. If we consider only the peptide backbone angles then the conformation of an N residue peptide is crudely stored in N bytes. If we consider the special case of small cyclic peptides, we can take advantage of the fact that changes in ϕ and ψ are highly correlated and store only one of the angles. This means that the conformation of an N residue cyclic peptide can be approximately represented in N/2 bytes. This rough reduction of the data and the use of fast integer arithmetic makes the method practicable. Aside from rapid comparison of conformations, there is another advantage to our approach. The range of the penalty function is a consequence of the degree of reduction of the data. A high compression factor of the information is equivalent to a rough high dimensional grid with a big range of the penalty function, whereas a smaller compression factor leads to a finer grid and smaller range of the penalty term. No further calculations to determine a cutoff radius for the penalty term are necessary.

2.4 Methods

The method was first tested on a simple model, pentane in a united atom representation. This has only two interesting degrees of freedom, the dihedral angles involving only carbon atoms. Stochastic dynamics (SD) simulations [12] of 100 ps were performed using a integration step size of 2 fs, employing the SHAKE algorithm [13] to maintain bond lengths, and using a tight coupling to a temperature bath [14] with a relaxation time of $\tau_T = \Delta t = 2fs$. All simulations were performed using GROMOS, the Groningen molecular simulation package and the GROMOS interaction function [15]. Two simulations were performed, one free simulation, and one using the local elevation method. Because of the small size of the system, no cutoff was applied in the evaluation of the nonbonded forces. The force constant

Flow diagram of the local elevation routine

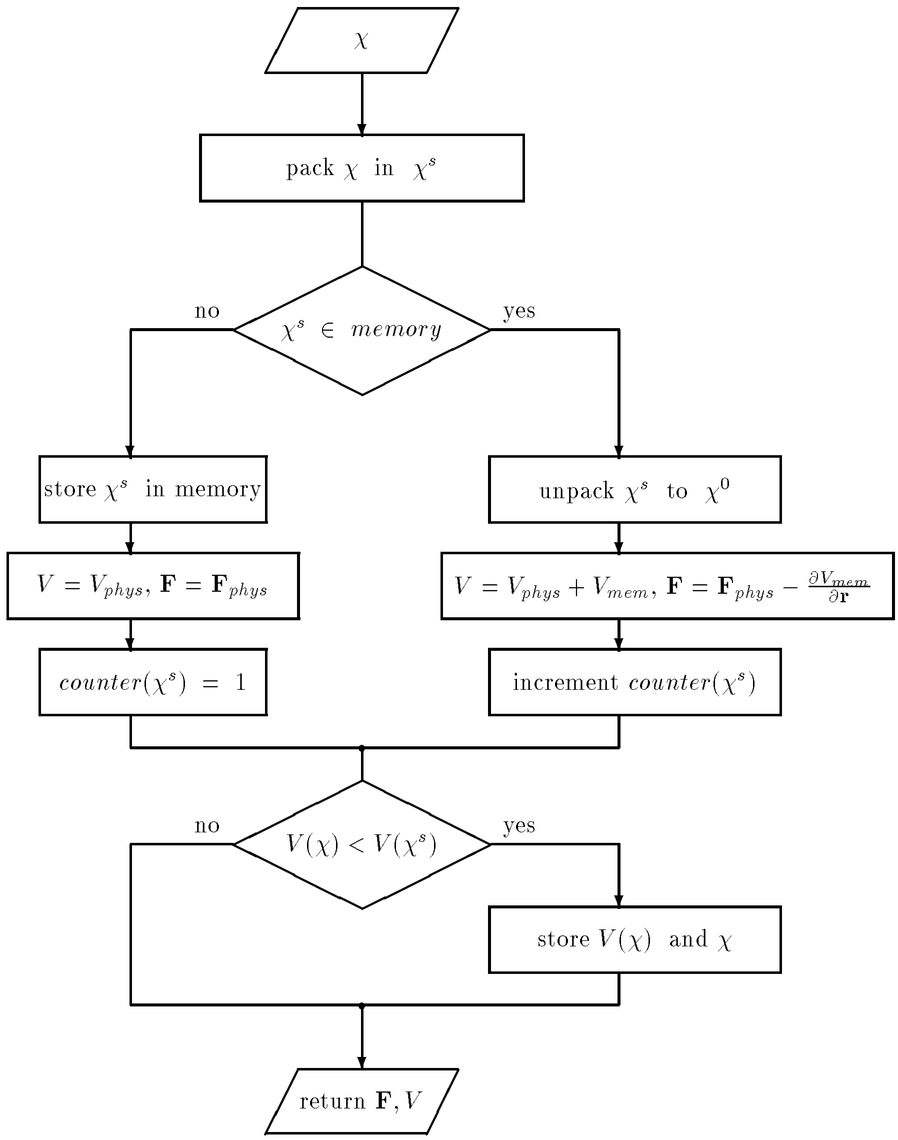

Figure 2.1: Flow diagram of the local elevation routine. the following abbreviations are used: χ actual conformation, χ^s actual conformation in packed form with loss of information by encoding, χ^0 actual conformation, but reduced information through encoding decoding procedure, V_{phys} standard physical potential energy, V_{mem} penalty potential energy term involving the conformational memory, $\mathbf{F}$ vector of force.

of the Gaussian penalty function in (2.4) was set to $k'_{mem} = 5$, $n_{\chi^0} = 1kJ\ mol^{-1}$ per resampling of a conformation, the width to $w = 22.5°$ and the grid length to define the cut off of the penalty function, i.e. the difference between two sequential ϕ^0 values, also to 22.5°. This is equivalent to recognising 16 states per dihedral angle. The starting structure was one of the possible gauche–gauche conformations (dihedral angle $\phi_1 = \phi_2 = -60°$).

For a real test case we chose Cyclosporin A, a well studied cyclic undecapeptide [16]. With 90 atoms in the united atom representation, it has a size typically of interest in rational drug design. For this molecule the following calculations were performed:

1. Free molecular dynamics (MD) at different temperatures
2. Local elevation (LE) molecular dynamics
3. Simulated annealing (SA) molecular dynamics
4. Potential energy annealing (PEACS) molecular dynamics

The first two calculations were the same as for the pentane example, but with Newtonian molecular dynamics (MD) rather than Langevin dynamics (SD), a longer simulation time of 1 ns, the penalty function force constant reduced to $k'_{mem} = 2.5$ and the width w of the penalty function and the ϕ-grid spacing doubled to 45°.

For the LE calculation, a set of approximations allowed each conformation of Cyclosporin A to be stored in a pair of 4 byte integers. First, only backbone dihedral angles were considered. Next, ω angles were assumed to be in the trans conformation, except for ω_{MeLeu^9} in the cis conformation, and so not stored. Finally, as discussed above, only the ϕ angle of each residue was stored, because of the correlation of changes in ϕ and ψ angles. This last assumption is only valid in small cyclic peptides, but is not unreasonable in these cases [17]. All simulations were started from the same well minimised but arbitrary conformation. Because of residual strain in this conformation a long free simulation of 1ns at 300 K was performed to guarantee total relaxation. The resulting conformation was used as starting conformation for both free and LE simulations.

For further comparison of different methods, we ran simulated annealing and PEACS simulations. To test the performance of the annealing techniques, method specific parameters were varied and standard simulation parameters were kept fixed. The dependence of specific variables gives an impression of the stability and general utility of a method and of whether a variation of the setup of the simulation is necessary to obtain different low energy structures. For simulated temperature annealing (SA) simulations, the velocities were taken from a Maxwell distribution at 1200 K and, after a variable equilibration time (see Table 1), the system was cooled down exponentially to 50 K over different time periods.

Performing PEACS simulations, an equilibration time of 2 ps at 900 K was used and the speed of cooling to 300 K and the characteristic average times to calculate the reference level of potential energy were changed with respect to ref. [3]. A summary of the annealing protocols can be found in Table 1.

All the annealing simulations we started from the same, well minimised structure which was also used in the free and LE simulations. No equilibration period at 300 K was used, because of the fast relaxation at high temperature. The number of annealing calculations was chosen so as to use CPU time comparable with LE simulation, with about 30% more effort on simulated annealing.

Before comparing conformations, each was minimised with the conjugate gradients method until the gradient was lower than $0.001 kJmol^{-1}nm^{-1}$. Structures were compared using both dihedral angles and internal distances. This was done so as to avoid misleading results when only one measure is used. For example, large but compensating changes in adjacent dihedral angles will result in little overall structural change. With respect to dihedral angles, structures were compared using a similarity grid as was used to store conformations during the LE simulation. A pair of conformations was defined to be similar, if the difference of each pair of dihedral angles fell within a predefined width. Another similarity criterion with respect to dihedral angle coordinates is the root mean square difference of corresponding dihedral angles of two different conformations $\{\phi_i\}$ and $\{\phi_i'\}$ each consisting of N dihedral angles called DHAD [18].

$$DHAD = \sqrt{\frac{1}{N}\sum_{i}^{N}(\phi_i - \phi_i')^2} \tag{2.6}$$

Using distance information instead of angle information, the so called distance matrix error (DME) [18] of the C_α–C_α distances of two different conformations $\{\mathbf{r}_i\}$

and $\{\mathbf{r}'_\mathrm{i}\}$ containing N C_α atoms becomes

$$DME = \sqrt{\frac{2}{N(N-1)} \sum_{i<j}^{N} (d_{ij} - d'_{ij})^2} \tag{2.7}$$

where d_{ij} is the distance between each pair of atoms i and j.

Table 2.1: Protocols of temperature annealing and potential energy annealing conformational search simulations with an exponential decay of temperature

	SA		PEACS	
nr	t_{equil} [1]	t_{cool} [2]	τ_V [3]	t_{cool} [2]
1	0	100	2.5	200
2	10	90	3.5	200
3	20	80	1.5	100
4	30	70	2.5	100
5	40	60	3.5	100
6	50	50	5.0	100
7	60	40	1.5	20
8	70	30	2.5	20
9	80	20	3.5	20
10	5	50	5.0	20
11	10	50		
12	15	50		
13	20	50		
14	25	50		
15	30	50		
16	35	50		
17	40	50		
18	45	50		
19	50	50		

[1] equilibration time at 1200 K [ps]

[2] time of exponential cooling [ps]

[3] relaxation time of potential energy [ps], see ref. [3]

2.5 Results and Discussion

2.5.1 pentane

The virtue of pentane in a united atom representation is that, with two dihedral angles, its conformations can be represented completely in a two dimensional figure. Figure 2.2 shows the Ramachandran-like plot of these two dihedral angles for a free simulation and one using the local elevation method. In the free simulation, the system repeatedly visits the energetically favoured conformations and gives little information about other parts of conformational space. In practice, it is unlikely, that one could ever simulate long enough to locate the other, higher energetic regions. In contrast, with the local elevation method, the conformational space is searched in the direction from low to high energy regions. Nearly the whole available conformational space is sampled after 20 ps, and only regions around the high energy cis conformations are not visited. After 100 ps, the total potential energy of every conformation is roughly equal and the conformational space is nearly uniformly sampled. This small test system is useful as a demonstration of principle, but says little about real utility.

2.5.2 Cyclosporin A

Calculations with Cyclosporin A yield comparable behaviour, which is more difficult to monitor, because of the high dimensionality of the configurational space spanned by the degrees of freedom of interest. In contrast to other searching methods in which only the final structure and its energy are of interest, in the local elevation method, the trajectory of the system maps out the low energy regions of conformational space. The fact that the system prefers low energy regions, but also crosses energy barriers is shown in Figure 2.3c. The stepwise increasing potential energy is a consequence of the applied potential energy penalty term. After lifting the potential energy of the system up to a value to overcome the lowest barrier surrounding a minimum, the system crosses it, and is suddenly freed from the local memory penalty term. From the number of steps in Figure 2.3c, one can estimate the number of distinct minima which were visited. The slow overall increase of the potential energy during the simulation reflects the overall lifting of a region in conformational space and can be used to get an idea of the local shape of the potential

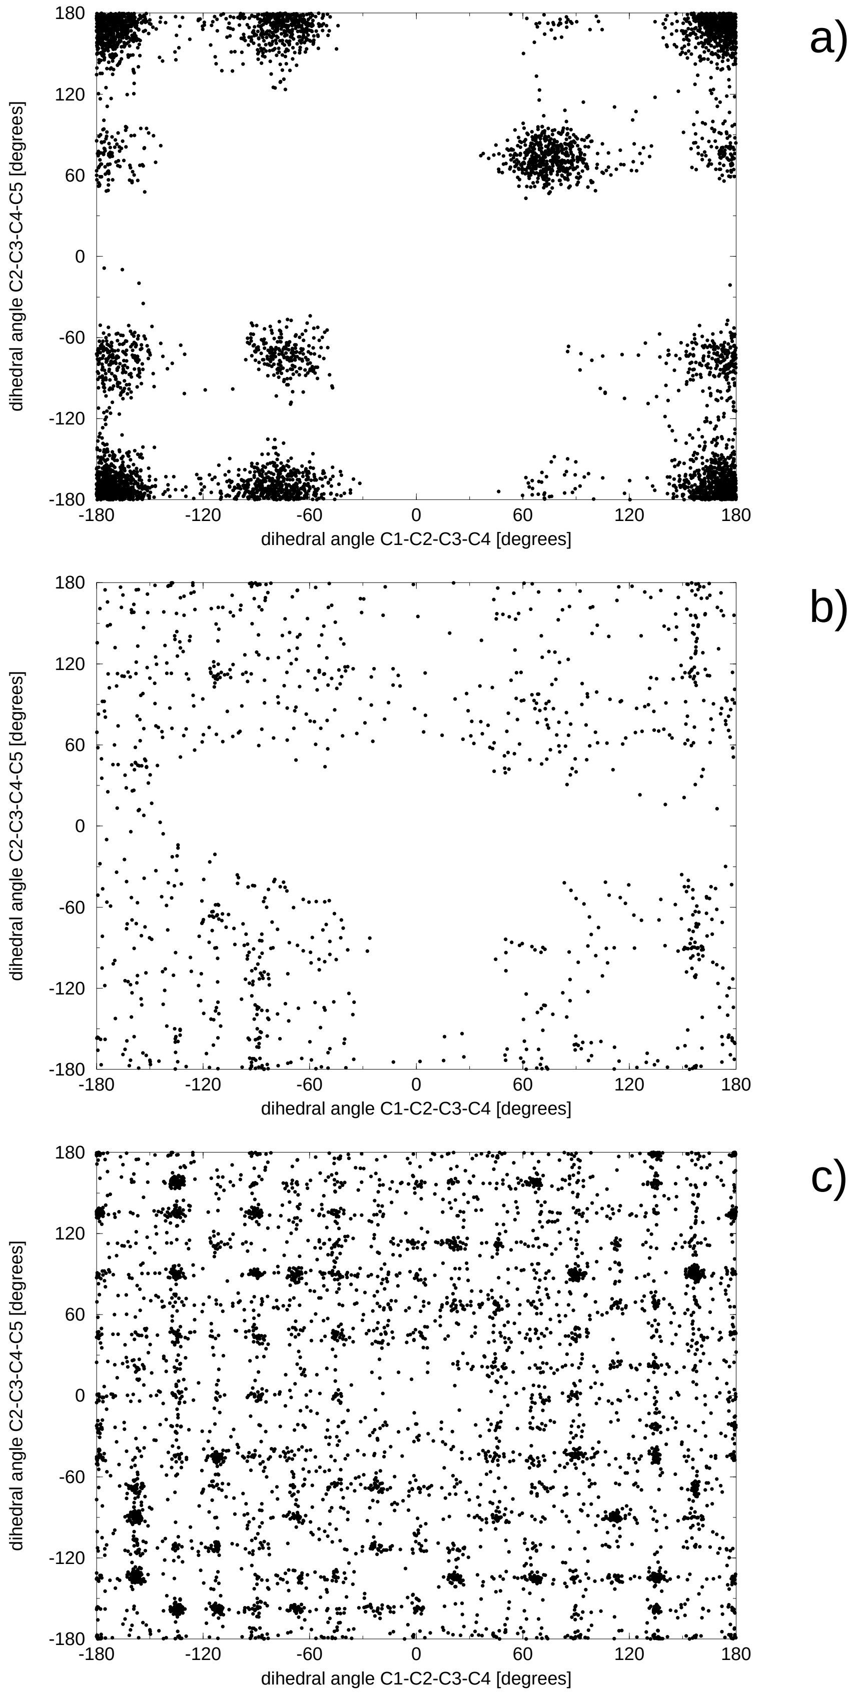

Figure 2.2: Visited conformations of pentane in a (a) 100 ps free stochastic dynamics simulation, (b) 20 ps local elevation simulation, (c) 100 ps local elevation simulation.

energy surface.

Improved searching performance can be shown by comparing similarity matrices corresponding to the LE and the free MD simulations (Figure 2.4). Of course there is no unique description for the similarity of conformations in a two dimensional plot. Therefore, we will use two other common similarity measures, DHAD and DME, apart from the method based on grid similarity and show that depending on the similarity criterion used, different interpretations of the results are possible.

The similarity matrices were calculated from trajectories of 1 ns with 2500 structures (every 400 fs one structure). In the grid similarity matrix, a dot marks whenever two conformations on the time axes fall within our criterion of similarity. We take advantage of the symmetry of the plots to show different similarity criteria on either side of the diagonal in Figure 2.4.

Figure 2.4a shows that in the free simulation at 300 K, only similar conformations are visited during the whole trajectory. In the free simulation at higher temperature (Figure 2.4b) the grid similarity decreases, but there is still the tendency to visit conformations repeatedly. In contrast, the local elevation method generates a trajectory of highly dissimilar structures in terms of grid similarity. To better quantify the number of states visited, one can consider the conformational space of the system as being described by an N-dimensional grid. Each of the $N = 11$ ϕ dihedral angles represents one axis of the grid and is divided into 8 domains of 45°. One can then count the number of different N-dimensional cubes (maximum number 8^{11}), which are visited at least once during the simulation. In the free simulation at 300 K, 3047 different grid cells were counted, in the simulation at 600 K the value increases to 19423 and in the LE simulation at 300 K 307395 different cells were counted from a total of 500000 simulation steps. This reflects an average number of sampling one of these visited grid cells of 164.1, 25.7 and 1.63 respectively. These numbers demonstrate that repeated visits of certain conformations are quite a problem for exploring conformational space, but are effectively prevented with the LE technique. The similarity plot of the local elevation simulation also shows that revisiting of areas of conformational space is possible, even if a penalty function is built up during the simulation and drives the system through conformational space. That is, the artificial peaks of V_{mem} do not necessarily form unsurmountable barriers which might prevent further searching.

Another commonly used similarity measure is the root mean square difference

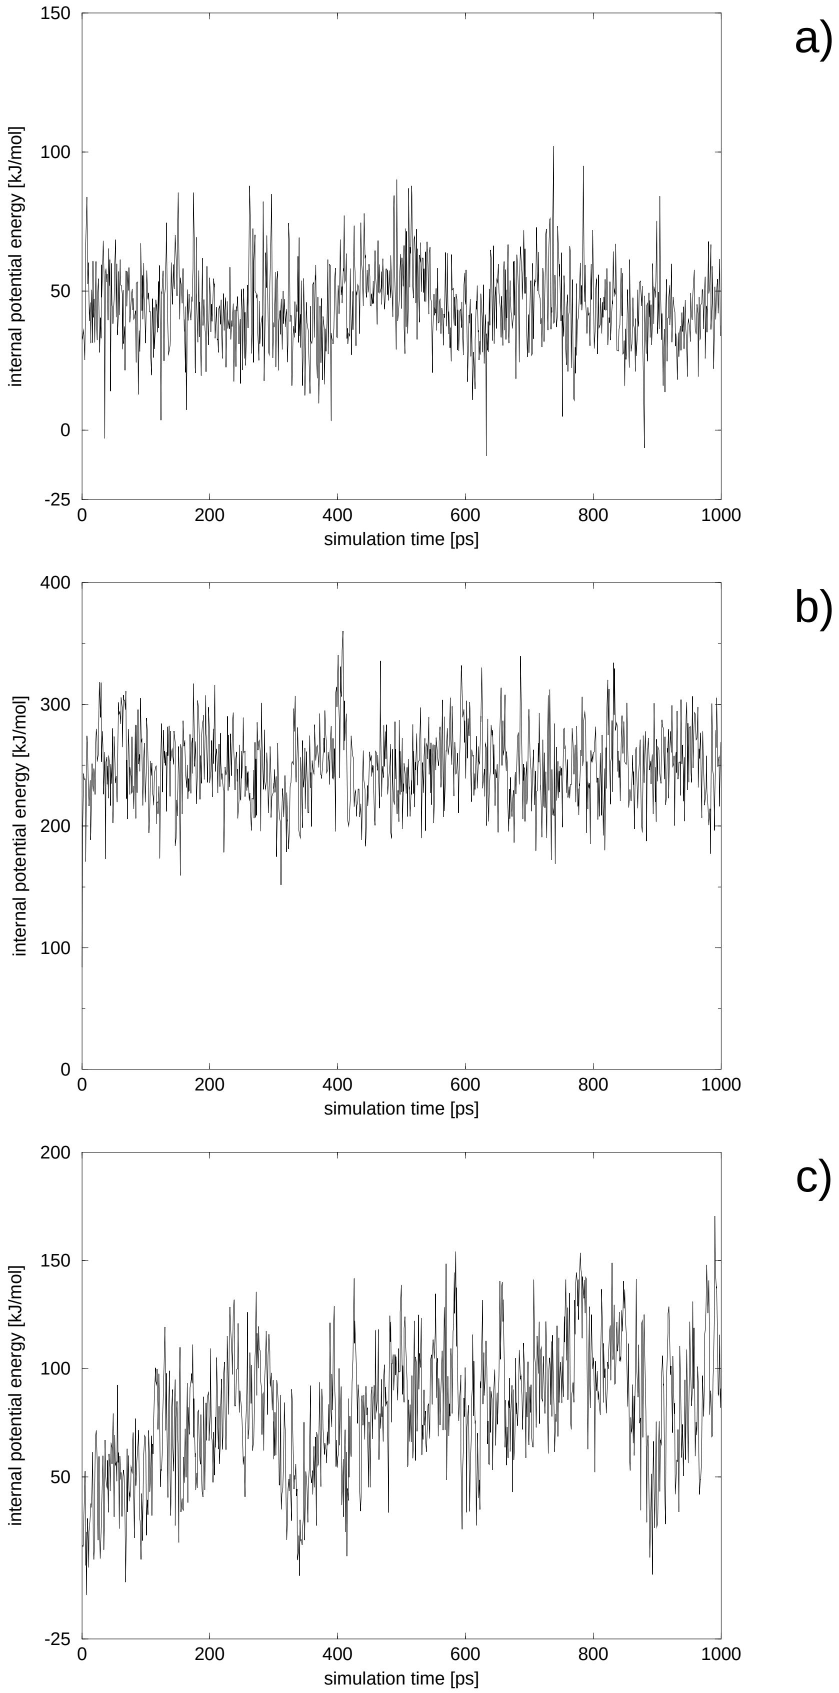

Figure 2.3: Internal potential energy as a function of time in simulations of Cyclosporin A. (a) free simulation at 300 K, (b) free simulation at 600 K, (c) local elevation simulation at 300 K.

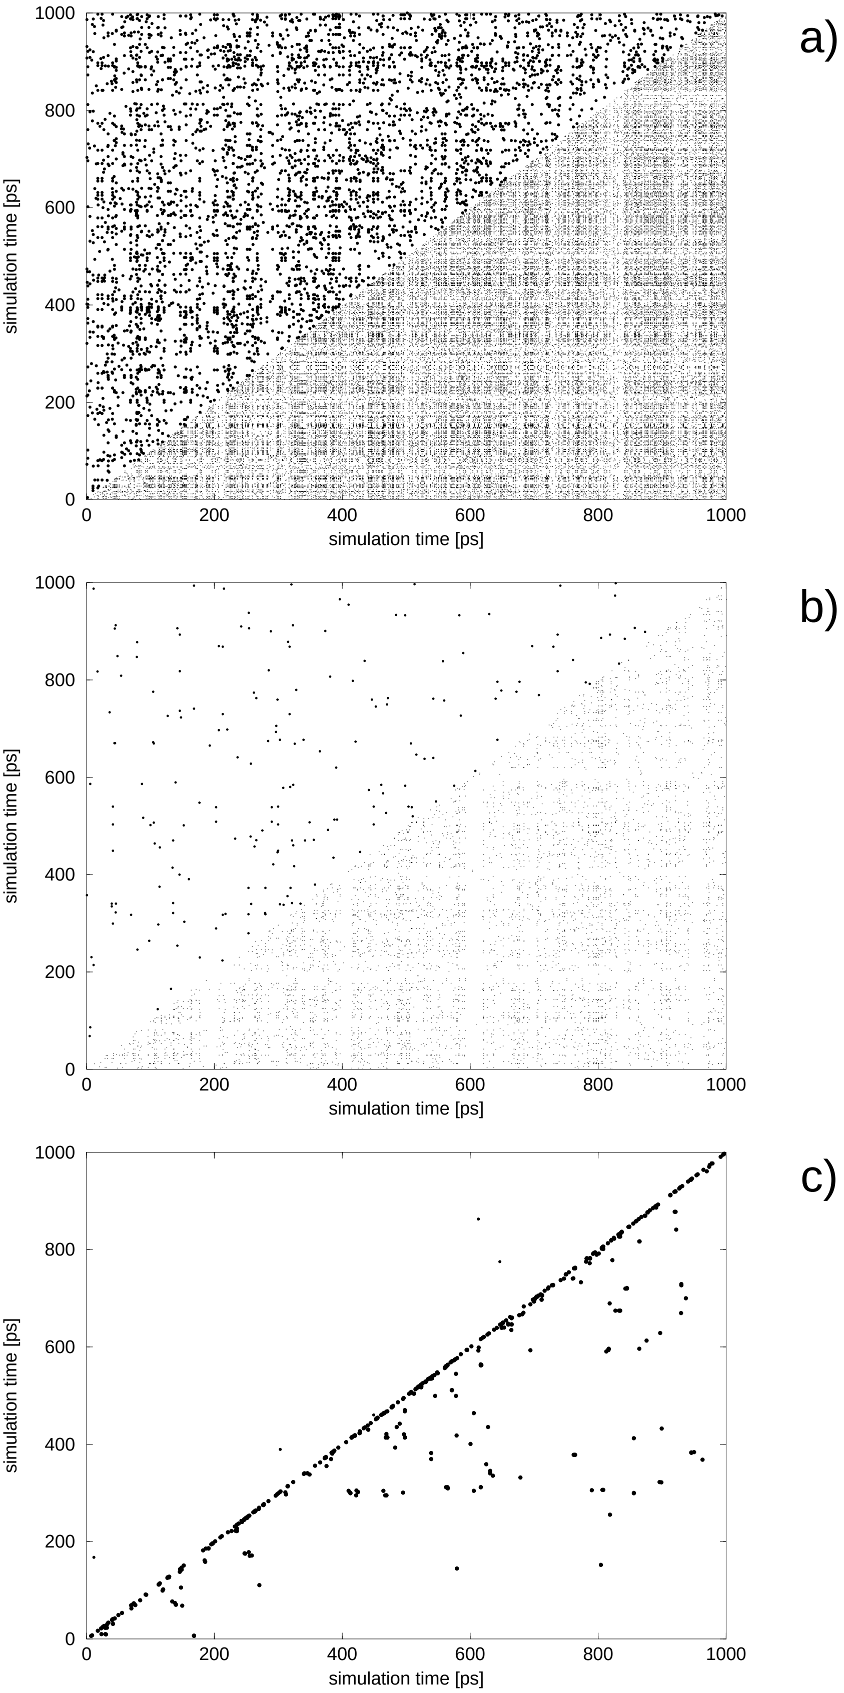

Figure 2.4: Grid similarity matrix of 11 ϕ dihedral angles in Cyclosporin A. Similarity criteria are $\Delta\phi \leq 30°$ for each of the 11 dihedral angles (upper triangle) and $\Delta\phi \leq 45°$ (lower triangle) for the free simulations, and $\Delta\phi \leq 45°$ (upper triangle) and $\Delta\phi \leq 60°$ (lower triangle) for the LE simulation. (a) free simulation at 300 K, (b) free simulation at 600 K, (c) local elevation simulation at 300 K.

between corresponding dihedral angles (DHAD). Whereas the grid similarity reflects the single largest difference, the DHAD reflects an average dihedral angle difference over all dihedral angles constituting a conformation. Averages bear the problem that local conformational changes may be averaged out, but in general the DHAD is a reliable measure to describe similarity in dihedral space. A comparison (data not shown) shows that grid similarity and DHAD similarity, both measures in dihedral space, lead to comparable results.

Our third measure to describe conformational resemblance is the DME of C_α–C_α distances. In our simulations, many correlated changes of ϕ_i and ψ_{i-1} dihedral angles were seen, but most of these cause only small changes in the C_α–C_α distances. While both kinds of variations are conformational changes, this kind of projection onto a single value may mask the information. Yet the DME and DHAD measures yield to a quit similar picture.

Molecules in solution are usually flexible and adopt conformations which constitute the state of lowest free energy. Therefore, we are not interested in a single molecular conformation, but in an ensemble of conformations with low internal energy. With the LE method, we generate a set of low energy conformations, different at least in terms of DHAD. To compare the results of LE simulations with other optimisation techniques, 20 conformations with lowest energy were picked from the trajectory. After minimisation of the potential energy, the structures were compared with respect to potential energy and the DHAD and DME with 19 structures obtained from simulated annealing calculations and 10 structures generated by potential energy annealing. All 20 structures from the LE simulation, using a memory penalty grid with width of 45° have potential energies in a range similar to those obtained by annealing techniques (Figure 2.5). Within the group, they also differ in both DHAD as well as in DME.

A quick assessment of the influence of grid size was made by repeating the calculations with a spacing of 22.5°. This results in a group of structures of distinct low energy as shown in Figure 2.5. Unfortunately the finer grid results in poorer searching. From these calculations it would seem that the choice of grid size introduces a somewhat arbitrary trade off. A finer grid results in lower energy structures, and a courser grid results in better searching. Figure 2.5 shows that the temperature annealed simulations produce the 4 of lowest energy structures, apart the LE ones generated using a finer grid. There is a group of 11 SA structures in the range of

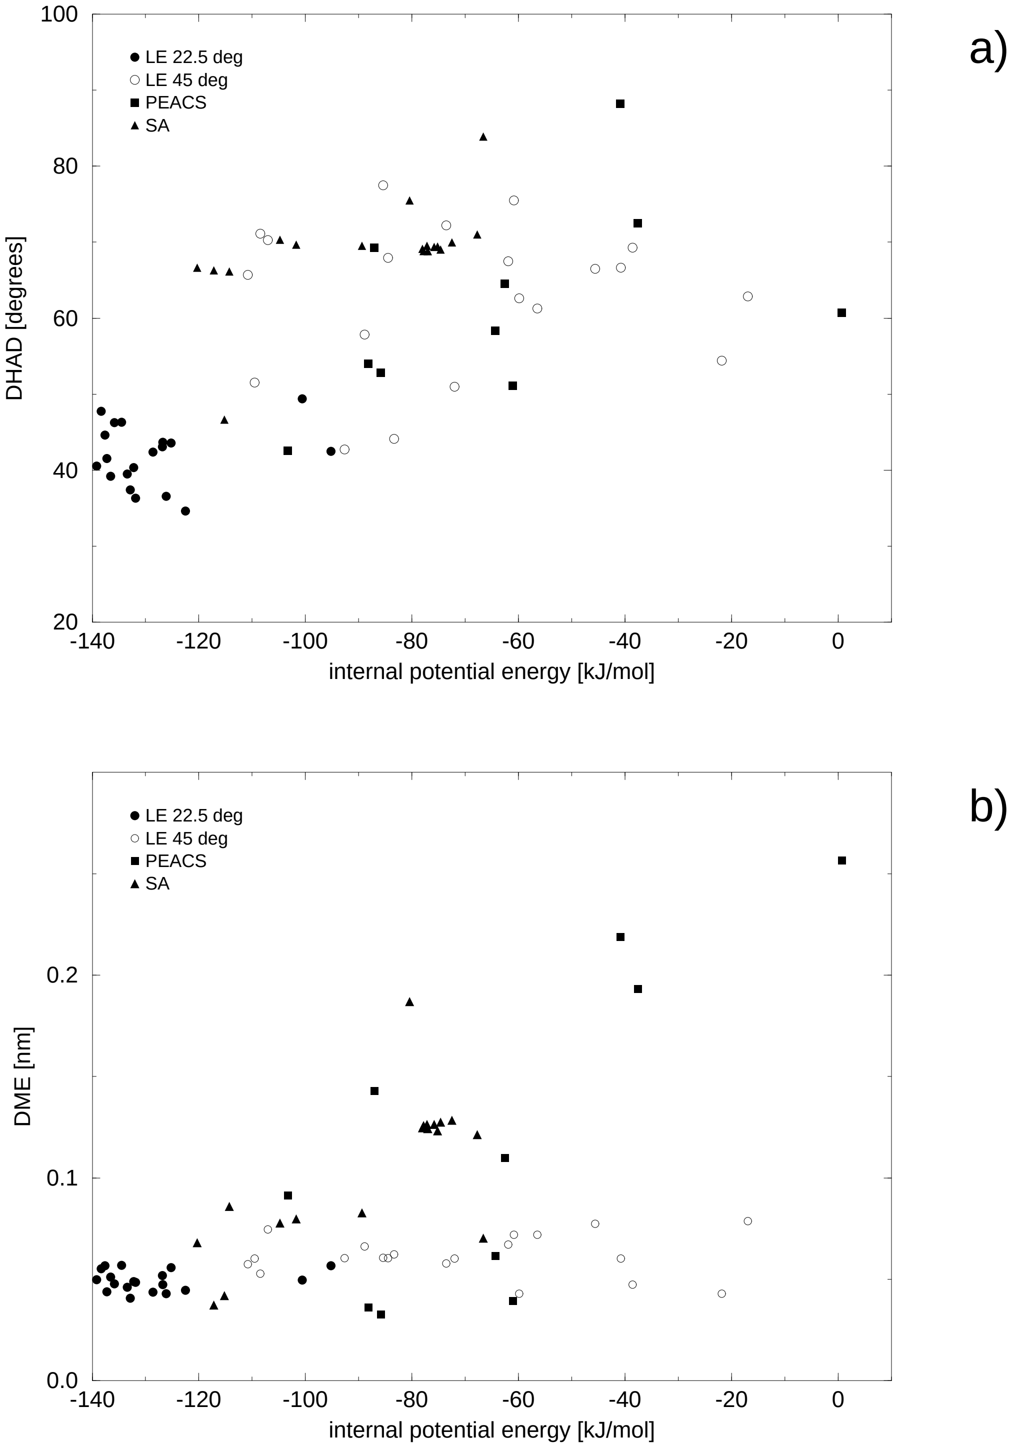

Figure 2.5: Similarity to starting structure versus internal potential energy of selected low energy structures of Cyclosporin A. Structures from local elevation simulation with penalty grid range $w = 22.5°$ are represented by filled circles, from $w = 45°$ by circles, from simulated temperature annealing by filled squares and from simulated potential energy annealing by filled triangles. (a) DHAD versus potential energy, (b) DME versus potential energy.

-80 to $-65 kJ\ mol^{-1}$, which reflect the finite cooling rate of the protocol. These might have reached better minima with slower cooling.

Calculations using potential energy annealing (PEACS) resulted in a group of structures with quite distinct properties. First, as shown in Figure 2.5, the best conformations in terms of potential energy are in the best half of all structures. At the same time the vertical axes in Figure 2.5a and 2.5b suggest a better spread of structures than with simulated annealing and even better than with local elevation. What is not shown in the figure is the relative proximity to the nearest minima. All LE structures converged during energy minimisation in less than 300 steps. Structures from the PEACS group required between 1370 and 6000 steps to converge.

Collectively, the results shown in Figure 2.5 do highlight the ad hoc nature of minimisation protocols. It may well be, that slower cooling with PEACS would have produced the lowest energies. Since the results do not only reflect the design of a method, but also implementation and protocol, we try to classify the different methods by their advantages. Simulated annealing seemed to be a very effective energy minimiser. PEACS (with our protocol) showed excellent searching properties at some cost in terms of potential energy. Local elevation is guaranteed to move from minimum to neighbouring minimum as demonstrated by Cyclosporin A. Unlike the other methods it will try to move away from pathological attractors as shown by the pentane calculations.

2.6 Conclusions

For small systems such as considered here, the results obtained using different search algorithms show the advantages of searching by learning from a molecule's dynamic history. The additional computational effort of about 15% compared to free simulation is justifiable by considerable gains in information about available conformational space. One must, however, note that the dynamics of the system is, of course, non Newtonian. The local elevation method adds energy to the system each time the force constant in the Gaussian potential energy term is increased. Clearly this kind of searching can only be used when some kind of temperature regulation is applied.

When searching for low energy structures, the method is a powerful alternative to normal molecular dynamics simulations. A comparison with other commonly used optimisation procedures shows the advantages and disadvantages of the different methods, which arise from the differences in search techniques. Whereas the local elevation method is designed to search for a set of low energy structures during one simulation, other methods usually refine one structure per simulation. Bearing this in mind, it would appear that the local elevation type of algorithm is an attractive method for situations where one is interested in the family of low energy structures which are adopted by a molecule in solution. The clear disadvantage of the method is that it is limited to small systems where the history can be easily stored and searched and the memory penalty function hypersurface is not of too high dimensionality. We are currently investigating similar techniques which store even less information but should be applicable to larger systems.

References

[1] SCHERAGA, H.A., Searching conformational space, in *Computer simulation of biomolecular systems*, edited by VAN GUNSTEREN, W.F., WEINER, P.K. and WILKINSON, A.J., volume 2, pp. 231–248, ESCOM, Leiden, 1993.

[2] KIRKPATRICK, S., GELATT, C.D.JR. and VECCHI, M.P., Optimization by simulated annealing, *Science* **220**, 671–680 (1983).

[3] VAN SCHAIK, R.C., VAN GUNSTEREN, W.F. and BERENDSEN, H.J.C, Conformational search by potential annealing algorithm and application to Cyclosporin A, *J. Computer Aided Mol. Design* **6**, 97–112 (1992).

[4] CRIPPEN, G.M., Conformational analysis by energy embedding, *J. Comput. Chem.* **3**, 471–476 (1982).

[5] CRIPPEN, G.M., Linearized embedding: A new metric matrix algorithm for calculating molecular conformations subject to geometric restraints, *J. Comput. Chem.* **10**, 896–902 (1989).

[6] CRIPPEN, G.M. and HAVEL, T.F., Global energy minimization by rotational energy embedding, *J. Chem. Inf. Comput. Sci.* **30**, 222–227 (1990).

[7] PURISIMA, E.O. and SCHERAGA, H.A., An approach to the multiple-minima problem by relaxing the dimensionality, *Proc. Natl. Acad. Sci. USA* **83**, 2782–2786 (1986).

[8] VAN SCHAIK, R.C., BERENDSEN, H.J.C., TORDA, A.E. and VAN GUNSTEREN, W.F., A structure refinement method based on molecular dynamics in four spatial dimensions, *J. Mol. Biol.* **234**, 751–762 (1993).

[9] D.D.Beusen and Marshall, G.R., The generation of three-dimensional structures from NMR derived constraints, in *Protein Structure and Engineering*, edited by Jardetzky, O., pp. 97–109, Plenum Press, N.Y., 1989.

[10] Glover, F., Tabu Search – Part 1, *ORSA Journal on Computing* **1**, 190–206 (1989).

[11] Glover, F., Tabu Search – Part 2, *ORSA Journal on Computing* **2**, 4–31 (1989).

[12] van Gunsteren, W.F. and Berendsen, H.J.C., A leap-frog algorithm for stochastic dynamics, *Mol. Simulation* **1**, 173–185 (1988).

[13] Ryckaert, J.-P., Ciccotti, G. and Berendsen, H.J.C., Numerical Integration of the cartesian equation of motion of a system with constraints: molecular dynamics of n-alkanes, *J. Comput. Phys.* **23**, 327–341 (1977).

[14] Berendsen, H.J.C., Postma, J.P.M., van Gunsteren, W.F., DiNola, A. and Haak, J.R., Molecular dynamics with coupling to an external bath, *J. Chem. Phys.* **81**, 3684–3690 (1984).

[15] van Gunsteren, W.F. and Berendsen, H.J.C., *Groningen Molecular Simulation (GROMOS) Library Manual*, Biomos, Groningen, 1987.

[16] Lautz, J., Kessler, H., Kaptein, R. and van Gunsteren, W.F., Molecular dynamics simulations of Cyclosporin A: The crystal structure and dynamic modelling of the solution structure based on NMR data, *J. Computer Aided Mol. Design* **1**, 219–241 (1987).

[17] Brunne, R.M., van Gunsteren, W.F., Brüschweiler, R. and Ernst, R.R., Molecular dynamics simulations of the proline conformational equilibrium and dynamics in antamanide using the GROMOS force field, *J. Am. Chem. Soc.* **115**, 4764–4768 (1993).

[18] Havel, T.F., The sampling properties of some distance geometry algorithms applied to unconstrained polypeptide chains: A study of 1830 independently computed conformations, *Biopolymers* **29**, 1565–1585 (1990).