JOURNAL OF CHEMICAL PHYSICS VOLUME 118, NUMBER 8 22 FEBRUARY 2003

Efficient Monte Carlo sampling for long molecular chains using local moves, tested on a solvated lipid bilayer

Mihaly Mezei[a)]
Department of Physiology and Biophysics, Mount Sinai School of Medicine, New York University,
New York, New York 10029

(Received 20 February 2002; accepted 27 November 2002)

An improved acceptance criterion for local move Monte Carlo method (in which trial steps change only seven consecutive torsion angles of a polymer) is introduced. In contrast to previous methods, not all possible loop closures are considered, but only the most structurally conservative one. By filtering such moves with the ''reverse proximity criterion'' introduced here, detailed balance is satisfied. Furthermore, the new method is ergodic, and is shown to be significantly more efficient than previous methods when applied to a fully solvated hydrocarbon chain with bulky sidechains as well as a fully solvated lipid bilayer. © *2003 American Institute of Physics.*
[DOI: 10.1063/1.1539839]

I. INTRODUCTION

Computer simulations of biological systems face the difficulty imposed by the time scale of biological processes. This is particularly true to the widely used molecular dynamics methodology that is limited to time steps in the femtosecond range. The other major alternative, Monte Carlo, has the potential to overcome this limitation, but it has proven fairly difficult to realize this potential.

Inherent to the Monte Carlo methodology are several potential windfalls: (a) it is very easy to control the degrees of freedom that are active in a particular simulation; (b) it is very easy to emphasize sampling of a particular region of the configuration space; (c) in aqueous systems (as most biological systems are) it allows an efficient implementation of grand-canonical ensemble sampling, thereby solving the crucial problem of solvating systems with potential pockets of water.

A major obstacle in applying the Monte Carlo methodology for polymers and biomacromolecules in the otherwise advantageous torsion angle space has been eliminated by the introduction of the so-called ''local move'' (also referred to as ''window move''). Local moves start with changing one torsion angle (called the driver torsion) followed by the adjustment of the six subsequent torsions to allow the rest of the chain to remain in its original position (while preserving all bond lengths and bond angles). It has already been observed by $\overline{\text{Go}}$ and Scheraga[1] that their algorithm calculating the torsion angles that close a loop can be applied to generate Monte Carlo moves in torsion space where only a segment of a polymer chain is moved.[2] A Jacobian correction to the acceptance criterion (which is essential for unbiased sampling) was added by Dodd *et al.*[3] Their formalism has been used successfully by Hoffmann and Knapp,[4] who tested it on the folding of a small polypeptide, by Wu and Deem,[5] who applied it to proline-containing peptides, and by Dinner, who tested it on proteins and nucleic acids.[6] Moves that are restricted to a segment of the polymer backbone allow torsion angle changes of reasonable magnitude even when the torsion is in the middle of a long chain, raising the possibility of performing efficient Monte Carlo simulations of biopolymers.

The present paper introduces an improved variant of local moves and compares its performance with the previous version and with the recently introduced extension biased torsions[7]—a method found to be of comparable efficiency to molecular dynamics.

II. METHODS: LOCAL-MOVE MONTE CARLO

Local-move Monte Carlo is based on simulations of a polymer in its torsion-angle space where the attempted change of a torsion angle is accompanied by changes in six subsequent torsion angles, allowing one to leave the rest of the molecule unchanged. As a result, the probability of accepting a torsion change would increase significantly for long polymer chains.

Limiting the change in the polymer backbone to only six torsions requires finding the right combination of torsions that achieve this—a nontrivial geometry problem, referred to as the loop-closing (or rebridging) problem. Examples for solutions to the loop-closing problem (and further references) can be found in Refs. 1, 4–6, 8.

While calculations aiming at a search of the conformational space (e.g., performing simulated annealing) can be performed relying on finding any solution of the loop-closing problem (even a solution to a weaker form of the loop-closing problem, such as the efficient torsional relaxation procedure of Wong *et al.*,[9] suffices), for calculations where correct Boltzmann-weighted averages are sought, i.e., algorithms based on the METROPOLIS method, microscopic reversibility has to be ensured as well.

The improved method presented in this paper allows the choice of the most advantageous solution of the loop-closing problem while maintaining Boltzmann sampling. Further-

---

[a)]Electronic mail: mezei@inka.mssm.edu

more, it is possible to implement it in a way that bypasses the full exploration of the loop-closing problem's solution space.

A. Ensuring microscopic reversibility with local moves

Most Markov-chain based Monte Carlo methods rely on maintaining microscopic reversibility (a sufficient condition) to ensure that the calculation samples the Boltzmann distribution. The difficulty in maintaining microscopic reversibility with local moves has two sources: the number of solutions to the loop-closing problem depends on the conformation and the torsion angle changes during such a local move are correlated, thus the volume element in the Boltzmann factor $\exp(-E/kT)d\phi_i d\phi_{i+1}...d\phi_{i+6}$ is not volume preserving.

Correct solutions to these problem were presented by Dodd *et al.*[3] followed by work of Hoffmann and Knapp[4] and Dinner.[6] All of these solutions require the determination of the full solution set to the loop-closing problem and the introduction of the Jacobian J corresponding to the transformation between the torsion angles and the independent constraint variables:[3]

$$P_{m\rightarrow n}^{\mathrm{acc}}=\min\left\{1,\frac{p_{n\rightarrow m}^{\mathrm{sel}}\exp(-E_n/kT)J(n)}{p_{m\rightarrow n}^{\mathrm{sel}}\exp(-E_m/kT)J(m)}\right\}, \tag{1}$$

where $P_{m\rightarrow n}^{\mathrm{sel}}$ and $P_{m\rightarrow n}^{\mathrm{acc}}$ are the probabilities of selecting and of accepting the trial move $m\rightarrow n$, respectively. The formula for the Jacobian can be found in Refs. 3, 4, and 6. Dodd *et al.* selected the trial conformation from the loop-closing problem's solution set with equal probability[3] while Hoffmann and Knapp selected from the solution set with probability proportional to the Jacobian[4]—the latter was shown by Dinner to perform better than the uniform selection.[6] Either solution introduces the ratio of the probabilities of making that selection in the forward and reverse moves into the acceptance expression.

The need to consider all solutions for generating the new conformation, however, is a drawback. It means that no matter how small the driver torsion is changed, the change in the remaining torsions can be quite large. When the affected polymer segment has side chains attached, this results in a significantly reduced likelihood of accepting such moves. This translates into a disproportionate amount of move attempts that will end up rejected. The problem with this is that this is only found out *after* the energy calculation and that is the costly part of the procedure. One possibility to alleviate this problem is to precede the energy calculation with a check for collisions and reject the attempt when a collision is found. The drawbacks of this approach are that (a) a special neighbor list has to be maintained during the simulation; (b) exactly for the case of the bulkier side chains there are an increasing number of neighbors to check; and (c) unless the check is conservative there is the danger of rejecting a move that would have ended up with favorable enough energy to be accepted (again, this possibility is more likely with the bulky side chain case) but the more conservative of the check, the less useful it is in preventing unnecessary energy calculations.

In the present work, a different solution to this problem is proposed. The new method, in contrast to the randomly chosen solution, always selects the loop-closing conformation after an attempted torsion angle change that is *the closest to the original conformation*. However, attempted moves are rejected outright *whenever the original conformation is not the closest conformation* for the reverse rotation applied to the selected conformation—this will ensure that $P_{n\rightarrow m}^{\mathrm{sel}}=P_{m\rightarrow n}^{\mathrm{sel}}$. This filter will be referred to as the "reverse proximity criterion." As a result, microscopic reversibility is maintained. Thus, this method eliminates those attempts that are unlikely to succeed, at the expense of only having to occasionally solve an additional loop-closing problem. As the cost of solving the loop-closing problem is independent of the system size, this additional expense becomes negligible for large systems. Thus, in the limit of dense systems and large system size the expected speedup approaches a factor equal to the mean number of solutions to the loop-closing problem.

An additional problem is that the range of valid torsion angles is not known in advance, thus any given choice of driver torsion can lead to an unsolvable loop-closing problem. The simplest and computationally least expensive solution to this is to reject outright any attempted torsion change when no solution is found to the loop-closing problem.[3]

B. Ergodicity

As the proposed method limits the space of acceptable moves compared to earlier versions, it is important to examine the possibility that it affects the ergodicity of the simulation. A simulation is ergodic if for any pair of states there is a sequence of moves, each with positive acceptance probability, that transforms one into the other.

First, there is a simple way to eliminate (at least, in theory) the problem: for a certain fraction of the moves use one of the previously introduced local moves. If the fraction is small then the loss of efficiency will also be small. The arguments in the following, however, indicate that it is unlikely that one has to resort to this device.

For systems without closed cycles (linear or branched polymers) it is easy to see that ergodicity is maintained with the new method. The argument relies on the fact that there is a finite probability that the moves are made in such an order that each driver torsion change leaves all previous driver torsions unchanged. This way any change in the conformation can be obtained by considering the changes in the driver torsion and ignoring the additional changes required for each local move (since those torsions will be adjusted later by the driver torsion of another local move). Whenever a required driver torsion change leads to an unsolvable loop-closing problem, additional driver torsion changes can be assumed to have been made (to torsions that limit this particular solution space). Since the system was assumed not to have cycles, these additional driver torsions will be closer to the end of the polymer chain than the driver torsion that was limited. Therefore, this process of applying driver torsions to enlarge solution spaces will always terminate, resulting in the required conformation.

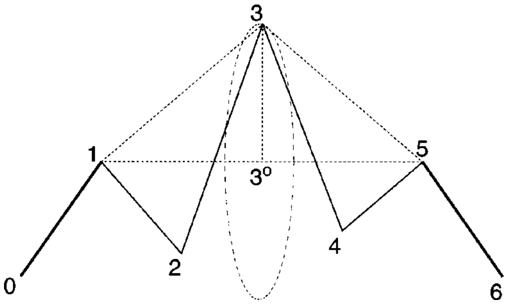

FIG. 1. Schematic description of the loop-closing problem and its solution. Thick lines represent fixed bonds. The dashed circle (whose plane is perpendicular to the line connecting the fixed positions of atoms 1 and 5) shows the possible placements of atom 3.

For system containing cycles (e.g., a protein with disulfide bridges) the above-given argument fails. In fact, rings with six atoms or less cannot be transformed from one allowed conformation to another with a sequence of small moves—just consider the ''boat'' and ''chair'' conformations of cyclohexane. On the other hand, the more atoms are in the ring, the less likely that the simple constraint of having to close the cycle can separate the conformation space into regions that are not connected by infinitesimal local moves. However, when the high likelihood of ergodocity is not enough, local moves suggested by Hoffman and Knapp[4] can be interspersed with the new local moves. These can be limited to segments of the cycle with no or only short side chains for maximum efficiency.

C. Finding the "nearest" solutions of the loop-closing problem

The solution to the loop-closing problem nearest to the initial configuration can, of course, always be found by obtaining the full set of solutions (*vide supra*) first. In the present work an iterative approach was chosen. The iteration is based on the fact that, of the three atoms whose coordinates are to be determined, the middle one should lie on a circle.[4,8] Figure 1 shows this circle in relation to the fixed atoms $\mathbf{r}_0$, $\mathbf{r}_1$, $\mathbf{r}_5$, $\mathbf{r}_6$ and the atoms whose coordinates the algorithm has to determine, $\mathbf{r}_2$, $\mathbf{r}_3$, $\mathbf{r}_4$. The center, $\mathbf{r}_{3^o}$, and radius, r, of the circle is given as

$$\mathbf{r}_{3^o}=\mathbf{r}_1+(0.5+c)(\mathbf{r}_5-\mathbf{r}_1), \tag{2}$$

$$r=\sqrt{d_{13}^2-c^2 d_{15}^2}, \tag{3}$$

where d_{ij} is the distance between $\mathbf{r}_i$ and $\mathbf{r}_j$ and

$$c=\frac{d_{13}^2-d_{35}^2}{2d_{15}^2}. \tag{4}$$

Thus, for any putative value of $\mathbf{r}_3$, putative coordinates of $\mathbf{r}_2$, $\mathbf{r}_4$ can readily obtained since knowledge of the distance from three known points determines the coordinates of that point:[10] the constraints result in two independent quadratic equations thus the solution of these two equations provide four pairs of putative solutions for $\mathbf{r}_2$ and $\mathbf{r}_4$.

The algorithm proceeds by calculating these putative solutions at evenly spaced points on this circle and comparing the resulting $|\mathbf{r}_4-\mathbf{r}_2|$ values with the reference distance d_{24}. Whenever the comparison shows a change of sign, the search continues by bisecting the interval over which the sign change is observed, until convergence is reached.

The nearest solution only can be obtained by a scan of the circle that starts from the point that is closest to the $\mathbf{r}_3$ of the reference conformation and is extended in both directions until a solution (if exists) is found. Extending the search to the full circle will lead to the complete set of solutions (simple geometric arguments can be used to exclude segments of the circle *a priori*). Note that the solution found nearest with the restricted search can occasionally be different from the solution selected from the full solution set using a least-squares criterion. This, however, should not affect the microscopic reversibility as long as the same definition of nearness is used in both directions.

The loop-closing algorithm was tested by solving it for known conformations. The algorithm missed the original conformation in 0.2% of the calculations when the full range was scanned and in 0.8% when the scan was directed to search for the nearest solution only.

While the iterative approach can sometimes miss solutions, it has several attractive features: ease of programming, the avoidance of computing solutions that will not be used, and its numerical stability. The numerical stability follows from several facts: (1) the roots of quadratics only are needed and not of polynomials of higher order (8 in Ref. 5 and 16 in Ref. 8); (2) the numerical errors in $\mathbf{r_2}$ are not affecting the numerical errors in $\mathbf{r_4}$; (3) unlike the methods using higher order polynomials, this iteration directly yields the coordinates of the backbone atoms. Since the formulas leading from the roots of the polynomial to the coordinates of the loop-closing atoms are rather complex (as seen in Appendices D and F of Ref. 8), it can lead to a concomitant loss of precision which can, on occasion, be quite significant (in spite of the fact that the roots of a polynomial can be ''polished'' to arbitrary precision). However, it is to be stressed that a detailed comparison of the speed and precision of the various loop-closing algorithms, including the one described here, should be performed to see if the above-listed attractive features translate into real advantage.

D. Numerical considerations

Numerical procedures repeatedly solving the loop-closing problem are generally prone to accumulation of round-off errors, especially at conformations that have torsion angles near 180°. This necessitated the use of double-precision arithmetics. Another issue to be considered carefully is the question of whether one should use the actual or the ideal (i.e., the original) bond lengths and 1–3 distances as they become slightly altered after each accepted move. It was found that the propagation of errors in the bond lengths is minimized if the actual values are used for d_{01} and d_{56}; d_{02} and d_{46} are calculated using the law of cosines using the actual values of d_{01} and d_{56}, the ideal values of d_{12}, d_{45} and of the angles α_{012} and α_{456}; and the ideal values for the rest. Here d_{ij} refers to the distance between atoms labeled i and j, α_{ijk} refers to the angle between the bonds $i-j$ and $j-k$, and the labels 0–6 refer to consecutive atoms in the chain, with atoms 2, 3, and 4 being the ones whose coordinates are to be

TABLE I. Comparison of the extent of sampling of the diphenyl substituted hexadecane with the different MC methods.

LPX					LJW			
f_1[a]	f_t[b]	P_{acc}[c]	P_{rr}[d]	$\mathrm{RMSD}_{\mathrm{bb}}$[e]	f_t	P_{acc}	$P_{\mathrm{acc/nn}}$[f]	$\mathrm{RMSD}_{\mathrm{bb}}$
2	30	0.75	0.05	1.95	18	0.27	0.0010	0.55
4	18	0.71	0.06	1.17	18	0.17	0.0015	1.05
6	14	0.57	0.07	1.82	22	0.21	0.0022	0.80
8	22	0.48	0.12	1.28	22	0.15	0.0007	0.74
10	22	0.44	0.12	1.43	14	0.15	0.0022	0.56
12	18	0.43	0.16	1.65	30	0.14	0.0009	1.17
14	18	0.39	0.15	1.19	22	0.11	0.0006	0.69
16	18	0.44	0.08	1.83	18	0.10	0.0015	0.88
18	22	0.35	0.15	1.32	18	0.10	0.0015	0.77
20	18	0.30	0.17	1.24	14	0.08	0.0016	1.06
40	30	0.21	0.19	1.04	18	0.04	0.0003	0.80
60	18	0.12	0.24	1.02	30	0.03	0.0008	1.04
80	22	0.08	0.31	1.02	14	0.03	0.0003	0.53

[a] f_1: the overall scale factor applied to the local move ranges.
[b] f_t: the best overall scale factor applied to the regular torsion ranges.
[c] P_{acc}: probability of acceptance.
[d] P_{rr}: probability of reverse rejection.
[e] $\mathrm{RMSD}_{\mathrm{bb}}$: root-mean square deviation of the driver torsion atoms.
[f] $P_{\mathrm{acc/nn}}$: probability of acceptance for local move that did not choose the nearest solution.

determined. A further reduction of the error propagation was obtained by computing the new torsion angles from the calculated positions of the three middle atoms, and recalculating these positions from the new torsion angles using the standard bond lengths and angles. Any side chain was added to these newly recomputed backbone atoms.

When the full range was scanned for solutions, the reverse proximity criterion was considered satisfied when the solution nearest to the original conformation from the solution set obtained from the original conformation was also the closest to the trial conformation. This avoids the need of selecting a threshold value to decide if the solution nearest the trial conformation is the same as the original conformation. However, when the scan is stopped after the nearest solution is found, its comparison with the original conformation does require such a threshold value—this is an additional drawback since for any choice there will be occasional false positives or false negatives.

It is also of interest to note that the concept of looking for the solution that is the least perturbed form the original conformation was also used by Gō and Scheraga.[2] However, their technique did not address the issue of detailed balance and, unlike the method presented here, also involved shifting the affected backbone segment.

III. RESULTS: TEST OF THE LOCAL-MOVE MONTE CARLO

First, the correctness of the sampling was tested on a polymer of 11 dummy atoms (i.e., no interactions, zero torsion potential) and on a single lipid molecule (DMPC, *vide infra*) in a dielectric continuum with $\epsilon = 80$. The freely rotating polymer yielded uniform distributions for all torsion angles; and both the torsion and the nonbonded energies of the simulated lipid molecule agreed with the results of simulating the same system with simple torsion-angle Monte-Carlo (MC).

Next, a polymer with a bulky side chain was simulated in aqueous solution to model the condensed phase environment. The aim of this exercise was to show that the limiting case referred to in estimating the expected improvement can be reached in realistic systems. The polymer chosen was hexadecane and it had biphenyl groups attached to it at the 3, 7, 10, and 13 positions. It was solvated with 770 water molecules in a cube of 29.01 Å edge. The solute was modeled with the CHARMM force field[11] and the TIP3P[12] model was used for water (sampled preferentially[13] to move the ones near the solute more frequently).

First, the ranges of torsion angle changes for local moves were adjusted individually to produce reasonably uniform acceptance rates. Next, runs of 5×10^6 MC steps were performed where one torsion was changed on the hexadecane at every 30th MC step, using different overall scale factors on the local torsion angle step range and on the range of the regular torsions (at the tail end of the polymer). The extent of sampling was characterized by the root-mean square deviation (RMSD) of the backbone atoms that were moved by the driver torsions of the local moves. Table I shows the results for runs with progressively larger scale factor on the local moves. For each local move factor, the result with the best regular torsion scale factor is shown. Besides the backbone RMSDs, Table I shows the acceptance rates obtained using the Jacobian-based selection strategy of Hoffmann and Knapp (labeled LJW) and using the method proposed here, based on full scans (labeled LPX). For the LJW runs, the acceptance probabilities of steps when not the nearest solution was chosen is also given and for LPX runs the probability of the reverse rejection is shown as well.

The results verify the suggestion that the acceptance

TABLE II. Comparison of the extent of sampling of the lipid bilayer with the different MC methods.

	N_{MC}[a]	$N_{\mathrm{MC}}^{\mathrm{eff}}$	D_{α_x}[b]	D_{α_y}	D_{α_z}	D_L[c]	D_T[d]	$D_{T_{\mathrm{HG}}}$	$D_{T_{\mathrm{HC1}}}$	$D_{T_{\mathrm{HC2}}}$
EXB	1	1.0	3.6	3.9	4.4	0.28	11.8	5.6	6.8	7.4
EXB	5	5.0	5.6	6.2	6.5	0.46	15.1	6.1	9.3	9.7
EXB	10	10.0	6.6	7.5	7.7	0.59	16.9	7.0	10.1	11.1
LPX	1	0.5	2.5	2.5	2.5	0.27	9.5	5.3	5.3	5.5
LPX	5	2.5	4.6	5.1	5.2	0.48	13.2	6.1	7.6	8.2
LPX	10	5.0	5.6	5.6	5.9	0.49	15.2	6.5	8.8	10.0
LJW	1	0.6	2.8	2.5	3.0	0.26	9.9	5.2	5.6	5.7
LJW	5	2.8	4.6	4.7	5.2	0.43	13.5	5.8	8.3	8.2
LJW	10	5.6	5.5	5.7	5.7	0.48	14.9	6.3	9.5	8.9

[a] N_{MC}: number of Monte Carlo steps/10^6.

[b] D_{α_x}, D_{α_y} D_{α_z}: the mean overall rotation of the molecule around the space-fixed x, y, and z axes (the bilayer is in the $y-z$ plane).

[c] D_L: the overall mean displacement of the center of the lipid molecules during the run.

[d] D_T, $D_{T_{\mathrm{HG}}}$, $D_{T_{\mathrm{HC1}}}$, $D_{T_{\mathrm{HC2}}}$: the mean overall displacement of the entire molecule, the headgroup chain, and the two hydrocarbon chains, respectively, due to torsion angle changes only.

probability of non-nearest solution is extremely low. They also show that the conformational changes obtained by the new method in the same number of steps (even including the steps that did not require energy calculations due to the new method's filter) is significantly higher than those obtained with the method of Hoffman and Knapp—on the whole about 50% more. Since in general one order of magnitude improvement in the precision of a Monte Carlo method requires two orders of magnitude increase in the computational effort required, this translates into a factor of 2 improvement in efficiency.

For the final test, a bilayer of 50 dimyristoylphosphatidylcholine (DMPC) molecules (i.e., 25 in each layer) was chosen, modeled also with the CHARMM force field.[11] The bilayer was solvated by a total of 2033 TIP3P[12] waters in a periodic hexagonal cell of 72.0 Å length and of 24.8 Å hexagon edge. The temperature of the system was 37 °C, well above the gel/liquid crystalline transition temperature of DMPC. This system was used earlier in this Laboratory to test the Extension-Biased Monte Carlo[7] that was found to produce conformational changes of comparable magnitude to a molecular dynamics run using the same CPU time.

There are conflicting arguments for the choice of a test system. The advantage of the method proposed here will best manifest itself in dense systems with longer side chains. These systems, however, converge significantly slower, thus making comparisons both more costly and less reliable. The latter argument led to the choice of the previously studied lipid system, even though the expected improvements are much less striking. Besides the two kinds of local moves tested earlier (LJW and LPX) the lipid test also included extension biased torsion angle sampling for all torsions (labeled EXB, that was found to be of comparable efficiency to molecular dynamics in Ref. 7) to provide a link with our earlier work.

The ranges of torsion angle changes for local moves were adjusted individually to produce reasonably uniform acceptance rates. An overall scale factor for these ranges was tuned to maximize the conformational changes of the lipids. The optimal acceptance rate for the LPX run turned out to be 0.20%, while the optimal acceptance rate for the LJW run was found to be only 0.10% (using the same step size as for the LPX run, as it turned out to be optimal here as well). This difference is just a manifestation of the fact that having to choose randomly from the solution set all too frequently results in rather drastic conformational change.

Using these optimized parameters, 10^7 MC step long canonical ensemble simulations were performed from the same starting configuration. Table II shows the progression of sampling for these runs. Since the acceptance strategies are quite different for the three techniques, they require different CPU times to run. This fact is indicated by the ''effective'' number of MC steps in Table II giving the number of MC steps the EXB algorithm would make during the same time.

The LPX results in Table II were obtained using full scans. Using the limited scan approach gave similar energies and conformational changes. The probabilities for the outright rejection of a move (i.e., when no solution is found to the loop-closing problem) were 27% and 28% for the full and limited scans, respectively, and the probabilities of failing the reverse proximity test were 24% and 29% for the full and partial scans, respectively. The latter difference is a manifestation of the different nearness definition used in the limited and full search.

Comparison of the new local move strategy (LPX) with the original technique (LJW) shows that even on this sidechainless system the new local moves perform slightly better (i.e., produce larger change in the system with a given amount of computer time) than the original since it achieved a consistently larger change in the various conformational indicators with 12% less computational effort. The real power of the new local moves will be apparent for dense polymers with side chains since in comparing two different solutions to the same loop-closing problem (e.g., the boat and chair conformations of cyclohexane) the displacement of the side chain atoms will be significantly larger than the main chain's displacement.

The extent of sampling with the local move runs was also found to be of comparable magnitude that of the extension-biased torsion-angle sampling runs. Even though for polymers of such moderate length the move of the full chain still provides somewhat better sampling than the local moves, the important thing is that the efficiency of local moves is independent of the chain length while the efficiency of any regular torsion move decreases steadily with the increase of the polymer length.

The various versions of the local-move MC have been implemented into the program MMC. It is available through the WWW.[14]

IV. DISCUSSION

The work presented demonstrated that local-move Monte Carlo can be applied to molecules with long chains in condensed phase aqueous environment with efficiency commensurate to molecular dynamics. The power of local moves results from the fact that only the minimal number of torsions are moved and the rest of the chain remains unchanged. Such moves have been introduced only relatively recently as they not only require the solution of the so-called loop-closing problem but pose the nontrivial problem of maintaining detailed balance.

Tuning/optimizing Monte Carlo moves is generally made relatively simple when a smooth decrease in the step size results in a smooth increase in the change in the system resulting in a smooth increase in the acceptance rate. This, however, is not the case for the previous versions of local move but the sampling technique introduced in this work eliminates this problem.

Furthermore, the new method is not only better tuneable, but it also has the potential to eliminate attempts that are unlikely to succeed, at the expense of only having to solve occasionally an additional loop-closing problem. This capability has been demonstrated in our test case using bulky side chains.

The possibility of efficient Monte Carlo simulations on proteins and nucleic acids opens the way to the application of a large arsenal of sampling techniques. Perhaps the best example is the cavity-biased grand-canonical ensemble simulation[15] that provides for the insertion and deletion of the solvent (i.e., water), thereby eliminating the uncertainty of the number of waters required in the interior of a protein or a complex.

ACKNOWLEDGMENTS

Conversations with Dr. Frank Guarnieri and Dr. Richard Lavery at various times have helped this project to reach its current state and are gratefully acknowledged. The author also acknowledges access to the computer facilities at the Institute of Computational Biomedicine (ICB) of the Mount Sinai Medical Center. Partial support of the work by NIH Grants Nos. P01 DA-12923 and GM-66531 are gratefully acknowledged.

[1] N. Gō and H.A. Scheraga, Macromolecules **3**, 178 (1970).
[2] N. Gō and H.A. Scheraga, Macromolecules **11**, 552 (1978).
[3] L.R. Dodd, T.D. Boone, and D.N. Theodorou, Mol. Phys. **78**, 961 (1993).
[4] D. Hoffmann and E.-W. Knapp, Eur. Biophys. J. **24**, 387 (1996).
[5] M.G. Wu and M.W. Deem, J. Chem. Phys. **111**, 6625 (1999).
[6] A.R. Dinner, J. Comput. Chem. **21**, 1132 (1999).
[7] P. Jedlovszky and M. Mezei, J. Chem. Phys. **111**, 10770 (1999).
[8] W.J. Wedemeyer and H.A. Scheraga, J. Comput. Chem. **20**, 819–844 (1999).
[9] W.H. Wong, Y. Cui, and R.S. Chen, J. Comput. Biol. **5**, 655–665 (1998).
[10] I.N. Brohnshtein and K.A. Semendyayev, *Handbook of Mathematics* (Van Nostrand Reinhold, New York, 1985).
[11] M. Schlenkirch, J. Brickmann, A.D. MacKerell, Jr., and M. Karplus, in *Biological Membranes*, edited by K.M. Merz, and B. Roux (Birkhäuser, Boston, 1996), p. 31.
[12] W.L. Jorgensen, J. Chandrashekar, J.D. Madura, R. Impey, and M.L. Klein, J. Chem. Phys. **79**, 926 (1983).
[13] J.C. Owicki, ''Optimization of sampling algorithms in Monte Carlo calculations of fluids,'' in *Computer Modeling of Matter*, edited by P.G. Lykos (American Chemical Society, Washington, DC, 1978).
[14] M. Mezei, MMC: Monte Carlo program for simulation of molecular assemblies. URL: http://inka.mssm.edu/~mezei/mmc
[15] M. Mezei, Mol. Phys. **61**, 565 (1987); **57**, 1207 (1989).