Parallel excluded volume tempering for polymer melts

Alex Bunker and Burkhard Dünweg
Max Planck Institute for Polymer Research, Ackermannweg 10, D-55128 Mainz, Germany
(Received 29 June 2000; published 22 December 2000)

We have developed a technique to accelerate the acquisition of effectively uncorrelated configurations for off-lattice models of dense polymer melts that makes use of both parallel tempering and large-scale Monte Carlo moves. The method is based upon simulating a set of systems in parallel, each of which has a slightly different repulsive core potential, such that a thermodynamic path from full excluded volume to an ideal gas of random walks is generated. While each system is run with standard stochastic dynamics, resulting in an *NVT* ensemble, we implement the parallel tempering through stochastic swaps between the configurations of adjacent potentials, and the large-scale Monte Carlo moves through attempted pivot and translation moves that reach a realistic acceptance probability as the limit of the ideal gas of random walks is approached. Compared to pure stochastic dynamics, this results in an increased efficiency even for a system of chains as short as $N=60$ monomers, however at this chain length the large-scale Monte Carlo moves were ineffective. For even longer chains, the speedup becomes substantial, as observed from preliminary data for $N=200$. We also compare our scheme to the end bridging algorithm of Theodorou *et al.* For $N=60$, end bridging must allow a polydispersity of more than 10% in order to relax the end-to-end vector more quickly than our method. The comparison is, however, hampered by the fact that the end-to-end vector becomes a somewhat artificial quantity when one implements end bridging, and is perhaps no longer the slowest dynamic variable.

DOI: 10.1103/PhysRevE.63.016701 PACS number(s): 05.10.Ln, 61.20.Ja, 61.25.Hq, 61.41.+e

I. INTRODUCTION

Computer simulations of dense polymer systems that make use of off-lattice models have been successful in the determination of both static properties, such as phase equilibria [1–4] and rubber elasticity [5], and dynamic properties, such as the details of single-chain and collective relaxation [6–9]. For a melt of polymers of length N, the relaxation time τ, the time taken for the polymer to assume a new configuration, scales as $\tau \propto N^2$ for small N (Rouse dynamics), while for larger values of N reptation behavior, $\tau \propto N^3$ [6,10] sets in. If the computer simulation of a polymer melt is performed in such a way that the dynamic properties are realistically reproduced, then this scaling will be directly related to the computational effort needed to effectively sample phase space and obtain meaningful results for the static properties [11]. As a result, progress in the simulation of systems involving polymers with large N has been severely hampered.

A modern trend in Monte Carlo simulations in statistical physics is to strictly distinguish between simulations that only aim to determine the static properties of a given system and those that also set out to determine the dynamic behavior. In the latter case, one has to follow the natural motion of the system confined to local dynamics constrained by topology and/or barriers. If one is, however, only interested in generating uncorrelated configurations as quickly as possible, one can use an artificial dynamics that is able to reach new effectively uncorrelated configurations much more quickly than the physical dynamics would allow.

The Rouse scaling law $\tau \propto N^2$ is a direct consequence of only allowing local motions to occur. It is independent of any constraint to motion and holds even for phantom chains with no interaction whatsoever except connectivity. Clearly, violating locality is an important step if one wishes to accelerate the acquisition of uncorrelated configurations. Particularly successful examples of schemes that achieve this include cluster algorithms [12] for critical phenomena, and the pivot algorithm [13,14] for isolated polymer chains, which collectively rotates a large part of the chain at once, thus allowing one to study the static properties for $N=10^5$ and above [14].

In a dense polymer system, however, such an approach will clearly fail, since practically any attempted large-scale move will be rejected due to overlap with other monomers. The effective constraints to motion that cause these techniques to fail are of physical importance—this is the mechanism that, for sufficiently long chains, gives rise to the onset of the considerably slower reptation dynamics. As a result, previous attempts to speed up simulations of polymer melts by only lifting locality are unable to alleviate the problem. For example, the continuum configurational biased Monte Carlo method (CCB [15–18]) and its variants [19] remove a chain (partly), and attempt to regrow it into the existing matrix. This can in principle be seen as a nonlocal approach like the pivot algorithm. However, in a simulation of a dense polymer melt the chain will grow preferentially back into the cavity from which it was previously removed. This effect becomes more pronounced with increasing chain length.

A simulation algorithm geared at only generating uncorrelated equilibrium configurations should thus not only find a way to violate locality but also the constraints resulting from the topology of the system and/or barriers. Fortunately, techniques have been developed to achieve this. The multicanonical ensemble and its variants (also called ''umbrella sampling,'' ''entropic sampling,'' or ''1/k sampling'') [20–25] try to identify barriers and then introduce a suitable bias in order to remove them (i.e., to allow the system to easily enter these unfavorable states). Simulated tempering (also called ''expanded ensemble'') [26–28] tries to systematically soften the constraints to motion by giving the system access to different parameter values where the barriers are weaker.

While the original implementations of simulated tempering [26,27] focused on rather obvious intensive variables such as temperature or chemical potential, it is evident that the formalism is applicable to any parameter that appears in the (effective) Hamiltonian. One useful parameter for polymer systems is the length of a particular tagged chain; this is the approach followed in the work by dePablo *et al.* [28]. Later this was combined with additional parameters (temperature, chemical potential) [29]. Another possible control parameter is the strength of the excluded volume interaction, and this is the route we follow in this paper. Since this term in the Hamiltonian directly generates the topological constraints that ultimately give rise to reptation-like slowing down, it is reasonable to expect that a systematic reduction of its strength should bear the potential of significant speedups. This latter manipulation has already successfully been applied to lattice polymers to measure chemical potentials [30,31] and within the framework of a multicanonical ensemble [32], while for continuum polymers it has so far only been used in an *ad hoc* fashion for equilibration purposes [6,8]. The present work should thus be viewed as an alternative approach to the existing method of dePablo *et al.* [29]. While their simulations introduce nonlocal moves by chain growth and removal, and effective constraint reduction by linking the system to lower chemical potentials, and thus lower densities, we rather remove the constraints by reducing the excluded volume interaction within the framework of parallel tempering and introduce nonlocality via pivot moves [13,14].

Parallel tempering, also called ''multiple Markov chains'' [33,34] or ''exchange Monte Carlo'' [35], is very similar in spirit to simulated tempering [26,27], but offers a number of both conceptual and technical advantages. Both approaches are based on studying a whole family of Hamiltonians $\mathcal{H}_i, i=1,\ldots,n$, each of which defines a standard Boltzmann weight $\exp(-\mathcal{H}_i)$, where, for convenience, the temperature has been absorbed into the definition of the Hamiltonian. This family of Hamiltonians will form a sequence in the one-dimensional space of the control parameter. Along this line, the Hamiltonians must be located close enough to each other such that the distribution of equilibrium states resulting from the Boltzmann weight $\exp(-\mathcal{H}_i)$ has significant overlap with the distributions given by the Boltzmann weights $\exp(-\mathcal{H}_{i-1})$ and $\exp(-\mathcal{H}_{i+1})$. A typical configuration for Hamiltonian $\mathcal{H}_i$ should be within the thermal fluctuations for both Hamiltonians, $\mathcal{H}_{i-1}$ and $\mathcal{H}_{i+1}$. As system size (volume V) increases, the distributions become sharper and sharper. As a result, more and more Hamiltonians will be required for this condition to still be satisfied. One should thus, in principle, study a system that is as small as possible.

If we denote the control parameter by ϕ, the above condition can be expressed as follows: For the averages of a given extensive variable A in two adjacent ensembles characterized by ϕ and $\phi+\Delta\phi$, the relation

$$|\langle A\rangle_{\phi+\Delta\phi}-\langle A\rangle_{\phi}|\approx\left|\frac{\partial\langle A\rangle}{\partial\phi}\right||\Delta\phi|\lesssim(\langle A^2\rangle_{\phi}-\langle A\rangle_{\phi}^2)^{1/2} \quad (1)$$

should hold. Since

$$(\langle A^2\rangle_{\phi}-\langle A\rangle_{\phi}^2)^{1/2}\propto V^{1/2} \quad (2)$$

for reasons of Gaussian statistics and $\partial A/\partial\phi\propto V$, one finds $\Delta\phi\propto V^{-1/2}$ or n, the number of Hamiltonians in the sequence, $\propto V^{1/2}$.

Given a family of Hamiltonians with the above condition satisfied, the tempering procedure consists of allowing a given system to make stochastic switches to neighboring Hamiltonians on the sequence in parameter space at fixed system configuration. Ideally, this results in a diffusion process with respect to the Hamiltonians. In particular, a configuration that was originally subject to a ''hard'' Hamiltonian (with constraints) can diffuse to a ''soft'' Hamiltonian (without), relax there quickly, and return to the original hard Hamiltonian. This should, ideally, accelerate the rate at which the system traverses phase space.

For a dense three-dimensional melt of flexible polymers, one expects the following scaling: The time to diffuse along the path of Hamiltonians and back is proportional to $n^2\propto V$. Assuming that the soft Hamiltonian does not provide any constraints, and that a suitable algorithm is able to generate there a completely new configuration in practically zero relaxation time, one finds altogether $\tau\propto V$. Furthermore, the smallest system one can study is given by equating the linear box size to the mean end-to-end distance $R\propto N^{1/2}$ (in a melt, the conformations are random walks [10]). Thus $\tau\propto V\propto N^{3/2}$, which is somewhat better than plain Rouse relaxation, $\tau\propto N^2$, and considerably faster than reptation, $\tau\propto N^3$. Nevertheless, it should be noted that the well-known slithering-snake algorithm [14] scales as $\tau\propto N^{\approx 1}$, i.e., is expected to be asymptotically even better than our procedure. For very dense systems, the prefactor in this law will, however, be large, due to small acceptance rates of the slithering-snake moves, such that one might need unrealistically long chains in order to actually observe the superiority. Where we expect the biggest payoff for our algorithm, however, is in systems where the (true physical) dynamics is governed by an activated process, such as star polymers [36], where

$$\tau\propto\exp(\text{const}\,N), \quad (3)$$

and for which the slithering-snake algorithm is not applicable. Regardless of these considerations, our first tests have deliberately focused on melts of linear chains, since this is the system that is characterized best with respect to both statics and dynamics.

The difference between simulated tempering and parallel tempering originates in how this idea is put into practice. Standard simulated tempering [26–28] considers only one system, whose configurations we denote by $\vec{x}$, and simply adds the parameter ϕ as an additional degree of freedom, which is treated via a standard Monte Carlo algorithm in that expanded state space. This procedure is governed by the Hamiltonian $\mathcal{H}(\phi,\vec{x})-\eta(\phi)$, where η is a suitable preweighting factor, to be determined self-consistently in order to prevent the simulation from getting trapped in the softest Hamiltonian. The partition function of the resulting expanded ensemble is given by

$$Z=\sum_i \exp(\eta_i)\int d\vec{x}\exp(-\mathcal{H}_i)=\sum_i \exp(\eta_i-F_i), \tag{4}$$

where F is the free energy (temperature is again absorbed in the definition). Since the arguments of the exponentials are extensive, the sum will always be strongly dominated by the largest term, unless all of them are practically identical. This means that unless $\eta_i \approx F_i \forall i$, the system will not be able to traverse the full extent of the available parameter space as one or more parameter values will become highly improbable.

In parallel tempering, n systems are run in parallel, each of which is assigned one of the Hamiltonians $\mathcal{H}_i$. Diffusion in Hamiltonian space is then facilitated by simple swaps of the configurations of adjacent Hamiltonians. Since each Hamiltonian will always be occupied, there is no problem of the simulation not visiting any particular ''unfavorable'' Hamiltonian, and thus it is no longer necessary to determine preweighting factors. Furthermore, the scaling considerations from above remain valid; the increased CPU effort by a factor of n is rewarded by the fact that we now have n random walkers available to produce data. The series of n systems can be seen as one extended ensemble with the partition function

$$\begin{aligned} Z&=\int d\vec{x}_1\cdots d\vec{x}_n \frac{1}{n!}\sum_p \exp(-\mathcal{H}_{p(1)}(\vec{x}_1))\cdots \\ &\quad\times\exp(-\mathcal{H}_{p(n)}(\vec{x}_n)) \\ &=\prod_i Z_i, \end{aligned} \tag{5}$$

where p denotes the possible permutations of the index set $i=1,...,n$, and we have made use of the arbitrariness in labeling. Thus the method just simulates n statistically independent systems.

The detailed balance condition for the swap is derived in a straightforward manner: If we denote two systems in which we attempt to switch the Hamiltonians by $\vec{x}$ (governed initially by Hamiltonian $\mathcal{H}_1$) and $\vec{y}$ (governed initially by Hamiltonian $\mathcal{H}_2$), then the transition probabilities w must satisfy

$$\begin{aligned} \frac{w[(\vec{x},\vec{y})\to(\vec{y},\vec{x})]}{w[(\vec{y},\vec{x})\to(\vec{x},\vec{y})]}&=\frac{P_{\rm eq}(\vec{y},\vec{x})}{P_{\rm eq}(\vec{x},\vec{y})} \\ &=\exp[-\mathcal{H}_1(\vec{y})-\mathcal{H}_2(\vec{x})+\mathcal{H}_1(\vec{x})+\mathcal{H}_2(\vec{y})] \\ &=:B \end{aligned} \tag{6}$$

(the partition functions cancel out in the ratio of equilibrium distributions). Using the standard Metropolis rule, the attempted swap $(\vec{x},\vec{y})\to(\vec{y},\vec{x})$ is accepted with probability $\min(1,B)$.

Given the simplicity of the method, and its potential, one should expect that its popularity will increase substantially in the future. So far, its use has not been very widespread, partly due to the fact that access to massively parallel computing facilities (for which the approach is ideally suited) is still somewhat limited. Applications up to now have included spin glasses [35,37], structural glasses [37,38], liquid-vapor phase coexistence in both simple [39] and polymeric [29] fluids, and several studies on the theta collapse of single polymer chains and related issues [34,40–45]. The fact that the strength of the excluded volume interaction could be used as a parameter in parallel tempering is mentioned in Ref. [32]; however, no actual run data were presented.

The remainder of this paper is organized as follows. In Sec. II, we describe the details of the model and the algorithm. Section II A defines the standard Kremer-Grest model of a polymer melt [6] and its simulation by means of stochastic Langevin dynamics. Section II B then describes the most important ingredients of our parallel tempering procedure, which is based upon altering the functional form of the repulsive core potential and replacing it by a nondivergent ''soft-core'' potential, until the limit of phantom chains is reached. This allows the chains to pass through each other, thus eliminating the slow reptation dynamics. When the repulsive core potential is soft enough, we will be able to perform pivot and whole polymer translation moves in the melt for which $\tau\propto N^{\approx 0}$, as described in more detail in Sec. II C. We also perform a comparison with end bridging [46,47], a very fast Monte Carlo algorithm that, however, does not conserve the chain lengths, and whose basic features are outlined in Sec. II D. Section III reports our numerical results. Section III A describes how we found the parameters for our procedure, while important time correlation functions to measure the efficiency of our algorithm are defined and presented in Sec. III B, resulting in our conclusion (Sec. III C).

II. MODEL AND ALGORITHM

A. Kremer-Grest model and Langevin dynamics

The Kremer-Grest model [6] is one of several off-lattice models for polymer melts which are commonly known as ''bead-spring'' models. All particles have purely repulsive Lennard-Jones (LJ) cores of the form

$$\begin{aligned} U_{\rm LJ}(r)&=4\epsilon\left[\left(\frac{\sigma}{r}\right)^{12}-\left(\frac{\sigma}{r}\right)^{6}+\frac{1}{4}\right], \quad r\leqslant 2^{1/6}\sigma, \\ U_{\rm LJ}(r)&=0, \quad r\geqslant 2^{1/6}\sigma, \end{aligned} \tag{7}$$

where ϵ and σ, as well as the bead mass, are set to unity such that time is in Lennard-Jones units. The finitely extensible nonlinear elastic attraction between the neighboring monomers on the chains is given by

$$U_{\rm ch}(r)=-\frac{k}{2}R_0^2\ln\left(1-\frac{r^2}{R_0^2}\right), \tag{8}$$

where $R_0=1.5$ is the maximum extension of the nonlinear spring, and $k=30$ is the spring constant. The spring constant is set to be strong enough to prohibit two polymer chains from crossing each other. We consider a system of M chains of length N in a cubic box with periodic boundary conditions at constant volume with density $\rho=0.85$.

We have simulated an *NVT* ensemble of this system through the use of Langevin (stochastic) dynamics [6,11], fixing the temperature at $k_BT=1.0$, where k_B denotes Boltzmann's constant. This involves the addition of a random force and a friction term, resulting in the following equations of motion in terms of particle positions $\vec{r}_i$ and momenta $\vec{p}_i$:

$$\dot{\vec{p}}_i=\dot{F}_i-\frac{\gamma}{m_i}\vec{p}_i+\vec{f}_i,$$
$$\dot{\vec{r}}_i=\frac{\vec{p}_i}{m_i}, \tag{9}$$

where $\vec{F}_i$ is the force due to the interactions with other monomers, m_i the particle mass, γ the friction constant, and $\vec{f}_i$ the stochastic force that satisfies the standard fluctuation-dissipation relation,

$$\langle f_{i\alpha}(t)f_{j\beta}(t')\rangle=2\gamma k_BT\delta_{ij}\delta_{\alpha\beta}\delta(t-t') \tag{10}$$

(i.e., uncorrelated with respect to both particle indices *i* and Cartesian indices α). These equations were solved using the standard velocity Verlet integrator [11,48], with friction coefficient $\gamma=0.5$ and time step $\Delta t=0.0125$.

Both static and dynamic properties of this model are very well known [6,11]. In particular, its slow Rouse- or reptation-like dynamics serves as a reference for the speedup obtained from our new Monte Carlo procedure.

B. Parallel tempering

We have performed parallel tempering by connecting a series of systems to the Kremer-Grest potential using successively softer repulsive core potentials. The *n* systems are simulated in parallel, and each system is on a separate processor of a massively parallel system (Cray T3E). Once an initial locally equilibrated configuration (in real and momentum space) is obtained for each of the systems, the potentials are allowed to switch between systems through Metropolis Monte Carlo steps, as described above. It should be noted that the kinetic energies cancel out in the Metropolis criterion. The swaps are implemented in a checkerboard fashion, where either the odd-even pairs or the even-odd pairs are tried. Between these swaps each system is run for a few stochastic dynamics steps; it is known that this procedure is quite efficient for equilibrating local degrees of freedom. For example, if we were to use eight processors we would first attempt to switch the systems 1–2, 3–4, 5–6, and 7–8, then run some stochastic dynamics, then attempt the switches 2–3, 4–5, and 6–7, then run more stochastic dynamics before attempting the first set of switches again. Figure 1 shows how this aspect of the algorithm is implemented. A reasonable duration for the Langevin runs between the swaps is obtained from studying the potential-energy relaxation, as described later.

In order to achieve large acceptance rates for the swaps, it is necessary to choose the form of the ''softened core'' potential such that the bond length *b* and the chain stiffness C_∞ are approximately maintained. The core repulsion of the

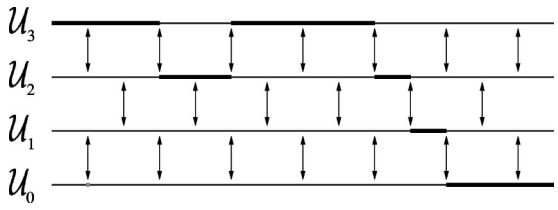

FIG. 1. Schematic representation of our parallelization scheme. The double arrows represent attempted switches and the thick line represents the path through the Hamiltonians followed by one of the systems.

neighbors and next-nearest neighbors on chains was thus kept intact, while all other repulsive Lennard-Jones potentials were replaced by the following ''softened core'' potential:

$$U_{\rm SC}(r)=A-Br^2, \quad r\leqslant r_t,$$
$$U_{\rm SC}(r)=4\epsilon\left[\left(\frac{\sigma}{r}\right)^{12}-\left(\frac{\sigma}{r}\right)^{6}+\frac{1}{4}\right], \quad r_t\leqslant r\leqslant 2^{1/6}\sigma, \tag{11}$$
$$U_{\rm SC}(r)=0, \quad r\geqslant 2^{1/6}\sigma,$$

where *A* and *B* are fixed by the continuity of $U(r)$ and dU/dr, leaving r_t as the only free parameter. For the Kremer-Grest potential $r_t=0$, and r_t is successively larger for each softer potential until the final potential in the series has $r_t=r_c=2^{1/6}\sigma$, the cutoff radius, which is the case for phantom chains. A graph of such a family of potentials is shown in Fig. 2. We will refer to the ratio r_t/r_c as the ''soft-core parameter.'' Observing Fig. 2, it becomes quite apparent why our tempering parameter is a superior choice to temperature for the Kremer-Grest model. Tempering in

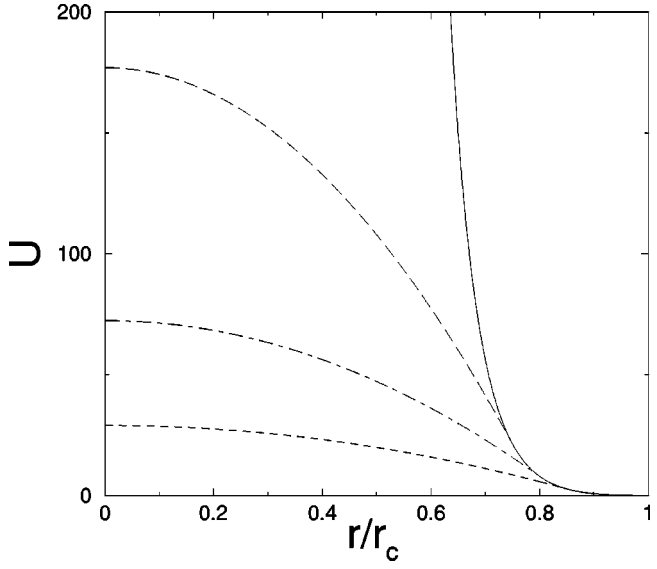

FIG. 2. A set of successively softer repulsive core potentials. In the simulation, we connect the system with the purely repulsive Lennard-Jones potential, shown as the solid line, through a series of such softened core repulsion potentials to the limit of phantom chains ($U=0$).

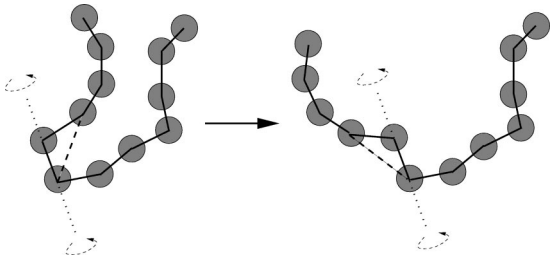

FIG. 3. Schematic representation of our implementation of pivot moves: Neither the nearest-neighbor nor the next-nearest-neighbor distances on the polymer chain are affected. As a result, none of the potential interactions that are kept at full strength are affected.

temperature would be the equivalent of altering the potential by a constant multiple. No matter how high a temperature reached, the r^{-12} divergence of the core would not be alleviated.

C. Pivot and translation moves

When we reach systems with extremely soft repulsive core potentials, then large-scale motions within the systems will have observable transition probabilities. We must ensure that these large-scale motions through phase space are such that only the parts of the potential that have been softened are affected. Thus the large-scale motions should not affect any bond lengths or angles within the polymers. Large-scale moves that fulfill this criterion are pivot and translation moves.

The pivot move involves rotating part of the polymer around the axis of a given bond. As shown in Fig. 3, all of the interactions that have not been softened are unchanged in this move. The translation move involves taking the entire polymer and shifting it a random distance in a random direction. It is in reaching systems where these kinds of moves are possible before returning to the Kremer-Grest Hamiltonian where we expect our algorithm to pay off. We have implemented the pivot and translation moves together in a single move where the whole chain is simultaneously translated and every bond is rotated, thus relaxing all the degrees of freedom of the chain with the exception of the bond lengths and angles, which are relaxed by the Langevin dynamics. We attempt these moves also quite frequently, as discussed later in the paper.

D. End bridging

In order to compare our algorithm with an established Monte Carlo method for equilibrating dense polymer systems, we have also implemented an end-bridging procedure combined with Langevin dynamics. End bridging, developed by Theodorou *et al.*[46,47], is a very efficient algorithm; however, it gains its speed only by giving up monodispersity. Instead, a fixed number of monomers and a fixed number of chains are simulated; the length, however, is allowed to fluctuate within predefined limits. In practice, these limits are defined by allowing all chain lengths between $N(1-f)$ and $N(1+f)$, where f is typically of order $\frac{1}{2}$. The algorithm involves allowing bonds within polymers to break and reattach

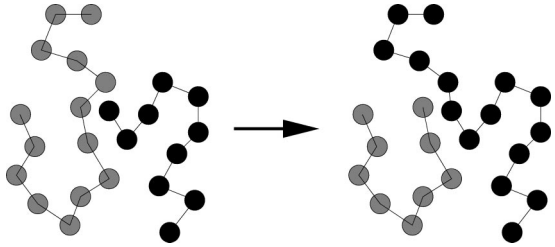

FIG. 4. End-bridging move for the Kremer-Grest model. If the end-bridging move in question involves an energy change ΔE, then the probability of the move is given by $P=\min(1,We^{-\Delta E})$, where W is a weight factor given by the ratio between the number of possible monomers to which the initial end can bridge, and the number of monomers to which the newly created end could bridge back.

to the ends of different polymers. It was originally devised for atomistic simulations with fixed bond lengths and involved an intricate procedure. Implementing this algorithm on the Kremer-Grest model, however, is far simpler since our bond lengths are able to fluctuate. The end of a chain searches for a possible chain to bridge to. This search is performed by finding all the monomers within the cutoff radius that are not on the same chain as the chain end in question. One out of these is selected at random. If a possible end bridge is found, then the move is accepted according to a Metropolis function, where the Boltzmann factor is multiplied with a weight factor. This weight factor is given by the number of monomers to which the end could possibly bridge divided by the number of monomers to which the newly created end could bridge back. This is necessary to satisfy detailed balance, i.e., to correct for the different probabilities to select the original reaction and the backreaction. A diagram of how end bridging works for the Kremer-Grest model is shown in Fig. 4. In our implementation, we run the Langevin dynamics for 10 LJ time units, followed by n_{br} bridging attempts, where n_{br} is 20 times the number of chains.

III. RESULTS AND DISCUSSION

A. Construction of simulation procedure

We first tested the algorithm with a system of 20 chains of length 60. For further simulation parameters ($\rho=0.85$, $k_BT=1.0$, etc., which were not varied), see Sec. II A. The effect of softening the potential is clearly seen in the standard pair correlation function $g(r)$, which is the probability to find a particle pair with distance r, normalized by the ideal gas value. This function is shown in Fig. 5, excluding the nearest and next-nearest neighbors on the polymer chains where the core repulsions are maintained at full strength. Compared to the fully repulsive system, there is a considerable probability for very short distances as soon as $r_t/r_c \gtrsim 0.95$, reflecting the ability of chains to pass through each other. This removes the topological constrains for chains of arbitrary length. Thus even without pivot moves, the dynamics would not be slower than Rouse relaxation.

Furthermore, we measured the single-chain static structure factor $S(q)=N^{-1}\langle|\Sigma_i \exp(i\vec{q}\cdot\vec{r}_i)|^2\rangle$ for both the Kremer-Grest model and the phantom chains as shown in Fig. 6. As

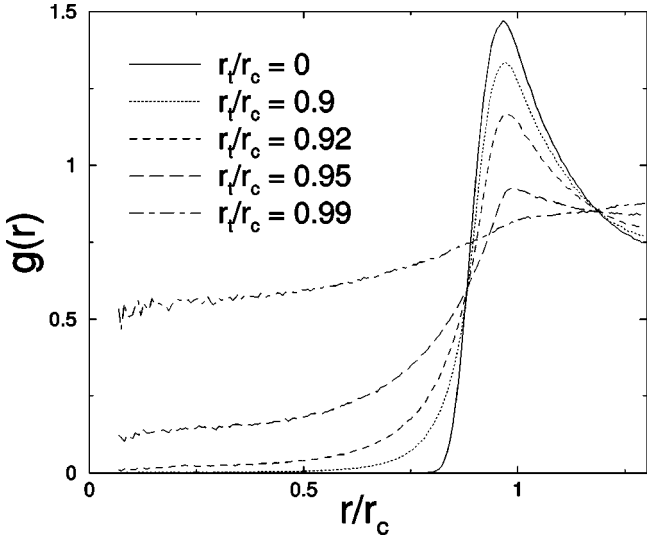

FIG. 5. Pair correlation function of an $N=60$ polymer system at density $\rho=0.85$ for several values of r_t/r_c, excluding nearest- and next-nearest-neighbor pairs along the chain backbone. Note that starting at about $r_t/r_c=0.95$, the chains are effectively able to pass through each other.

expected, the Kremer-Grest model reproduces the random-walk exponent $\nu=0.5$, observed from the decay $S(q) \propto q^{-2}$. The result for the phantom chains is very similar to the Kremer-Grest model, indicating that the overall structure of the chain does not change very much as we soften the repulsive core potentials. This in turn means that no major chain rearrangements are necessary along the thermodynamic path, such that the transitions should be quite easy.

Nevertheless, it turned out that for our number of monomers one needs of the order of 10^2 (due to computer restrictions we used 128) systems in order to connect from the fully repulsive potential to the phantom chain limit, requiring that the swap acceptance rates are of order $\frac{1}{2}$. The soft-core parameters were adjusted by hand, in essence via a trial-and-error procedure. It was found that the r_t step size to the next softer potential had to be reduced drastically as the phantom chain limit was approached.

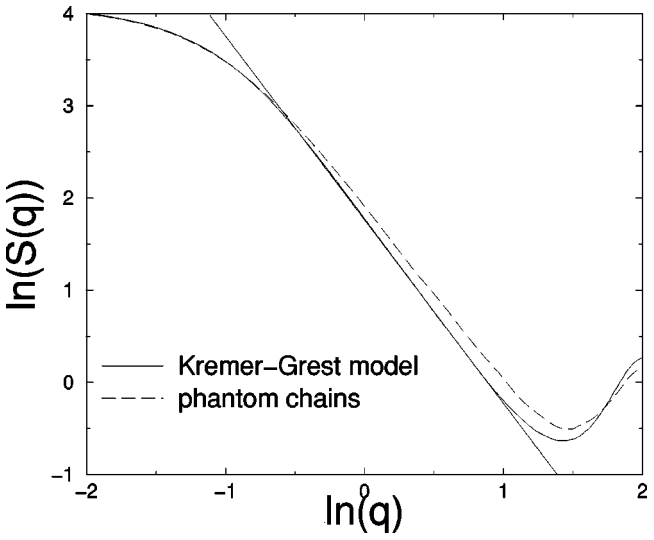

FIG. 6. Single-chain static structure factor for both the Kremer-Grest model (chain length $N=60$, density $\rho=0.85$) and phantom chains where the nearest- and next-nearest-neighbor interaction along the chain was left intact. The similarity in the static structure factor indicates that the large-scale structure of the chains is nearly invariant as the repulsive core potentials are softened. The straight line represents a slope of -2 corresponding to $\nu=0.5$.

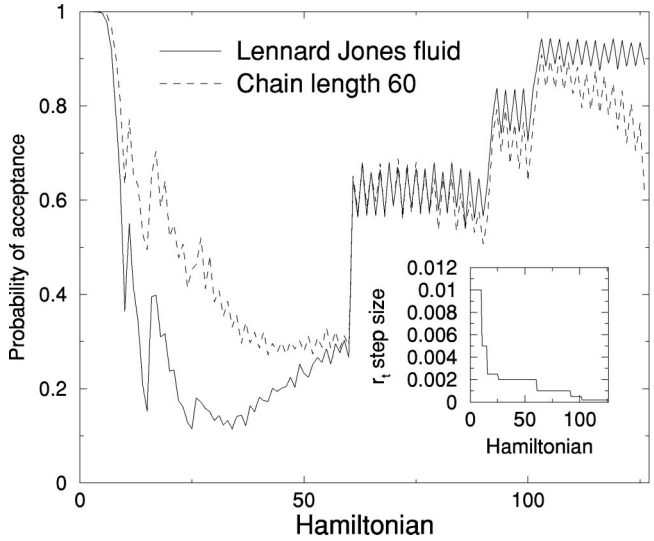

FIG. 7. Probability that a system would transfer to the next softer potential, as a function of its own potential index, which increases with softness. Although the low molecular weight fluid and the polymer system have the same density $\rho=0.85$ and temperature $k_BT=1.0$, there are marked differences in behavior, as discussed in the text. The inset graph shows the steps in soft-core parameter values for each of the potentials.

A graph of the resulting transition probabilities for each of the 128 potentials for both 20 chains of length 60 and 1200 purely repulsive Lennard-Jones particles is shown in Fig. 7. The first potential with a softened core was chosen at $r_t/r_c=0.74$. From then on, we picked further r_t values with uniform spacing; however, occasionally this spacing had to be reduced along the path, in order to prevent the transition rates from dropping too strongly. This procedure thus introduced a number of steps in the r_t spacing, as shown in the inset of Fig. 7. These steps in turn give rise to jumps in the transition rates, as clearly seen in Fig. 7. Furthermore, there are very rapid oscillations, which we believe to be a consequence of the checkerboard algorithm.

Comparing the low molecular weight case with the polymer system, one observes that in the regime of small $r_t/r_c=0,...,0.95$, where the pair correlation function exhibits clear core repulsion (see Fig. 5), the transition probabilities increase with increasing chain length. In the intermediate regime, the transition probabilities are roughly independent of chain length and as the limit of phantom chains is approached, $r_t/r_c \gtrsim 0.99$, the transition probability drops off drastically with increasing chain length.

We believe this phenomenon can be explained as follows. The transition probability is governed by the overlap in the energy distributions. The degree of this overlap is dependent on the difference between the mean energies and the width of

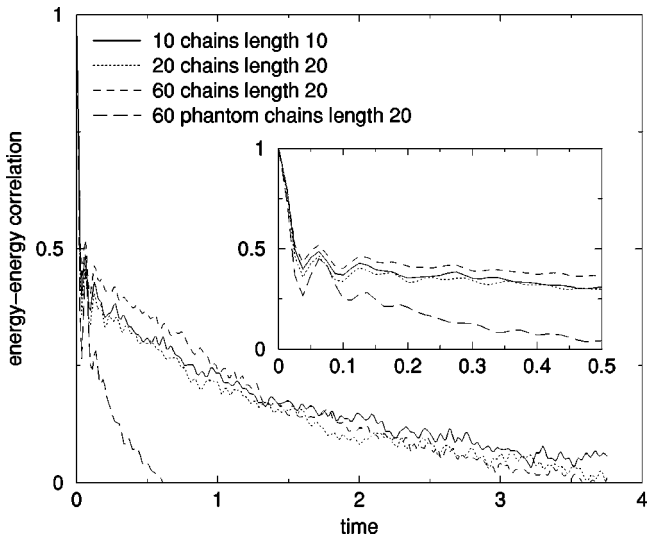

FIG. 8. Normalized energy-energy correlation function as a function of time in Lennard-Jones time units, at density $\rho=0.85$. Note the very fast decay on short time scales followed by a much slower decay that is dependent on the soft-core parameter.

the energy distributions. In the low soft-core parameter regime, the dominant factor is the difference in average energy, since large interaction energies are present. As chain length increases, the fraction of interactions which belong to nearest or next-nearest neighbors along a chain also increases. As these interactions are not softened, the difference in mean energy between adjacent potentials is reduced upon increasing the degree of polymerization. Thus the overlap, and the transition probability, is raised. Conversely, as the limit of phantom chains is approached, the interaction energies involved become very weak, and the behavior rather becomes entropy-driven. Since polymerization reduces the translational entropy, one should expect that also the *variation* of the free energy with r_t is decreased upon polymerization. As the width of the energy distributions should be directly related to the ''specific heat,'' i.e., the second derivative of the free energy with respect to r_t (which can be viewed as the equivalent of temperature for our system), one expects an entropic narrowing of the distributions, while the mean values remain largely unaffected. Thus the overlap and the transition rate are reduced.

We feel that an analytical relationship between the chain length and the transition probability in the phantom chain limit can probably be developed, maybe by a perturbation expansion around the ideal gas of random walks. If this is so, then, from a determination of a set of soft-core parameters that have a constant transition probability for a certain number of Lennard-Jones particles, a general relationship could be developed that would provide a set of soft-core parameters for an arbitrary number of chains of an arbitrary length. The development of such a procedure is, however, beyond the scope of this paper.

From measuring the time autocorrelation function of the potential energy for the untempered system (i.e., without swaps between potentials), shown in Fig. 8, we have found that the correlations decay very quickly after only a few time steps of Langevin dynamics and then cross over to a much slower decay. Our interpretation of the very fast initial decay is that it is a direct consequence of the local bond oscillations, which happen on roughly this time scale. This is also in accord with the observation that it occurs independently of the degree of softening, see Fig. 8, since the potentials between bond neighbors remained unchanged for all potentials. Conversely, the long-time behavior is quite strongly affected by the softening, again in agreement with the expectation that the function should decay much faster for a softer system. Furthermore, the energy-energy autocorrelation function is independent of chain length, as expected.

These results suggest that it is most efficient to attempt the switches frequently, on the time scale of the bond oscillations. We thus constructed the following algorithm: On every system we perform the large-scale chain reorientation attempt on enough chains so that attempts are made to move at least 5% of the monomers. This is followed by four Langevin dynamics steps, after which the Hamiltonian swaps (either even-odd or odd-even) are performed. Then the procedure is repeated. For reasons of simplicity, we have applied the identical moves to all systems. It should be stated, in the interest of future development of this method, that this condition is, however, not necessary. Any algorithm that leaves the Boltzmann distributions of the different systems invariant will be valid. For our algorithm, however, as currently implemented, not to attempt the large-scale moves on the hard systems would only generate idle CPU time since the processors all have to wait for the slowest system to finish before attempting the next swap. A further optimization could, however, involve eliminating the large-scale moves on the hard systems and replacing them by more Langevin steps. By fine-tuning this, one should be able to reduce synchronization overhead to a minimum. We also found that as the soft-core parameter approaches the phantom chain limit, local bond oscillations can become unstable. This can be easily remedied by increasing the friction in the Langevin dynamics as the soft-core parameter increases.

B. Correlation functions

An appropriate way to benchmark the program is to determine the CPU time needed per relaxed chain. Since every system periodically passes through the Hamiltonian with the full Lennard-Jones repulsive hard-core potential, each system can be seen as a Kremer-Grest model that is sampled every time this Hamiltonian happens to lie on it. Useful single-chain quantities to measure are the normalized end-to-end vector autocorrelation function, $\langle \vec{R}(t)\cdot\vec{R}(0)\rangle/\langle R^2\rangle$ with $\vec{R}=\vec{r}_N-\vec{r}_1$, and the autocorrelation function of the lowest five Rouse modes, $\langle \vec{X}_p(t)\cdot\vec{X}_p(0)\rangle$ with $p=1,...,5$ and [8]

$$\vec{X}_p=\sqrt{2}N^{-1/2}\sum_{i=1}^{N}\vec{r}_i\cos\left[\frac{p\pi}{N}(i-1/2)\right]. \tag{12}$$

These quantities must be measured in such a way that only correlations between configurations where the Kremer-Grest Hamiltonian is present are counted. Thus we define the following procedure to measure autocorrelation functions:

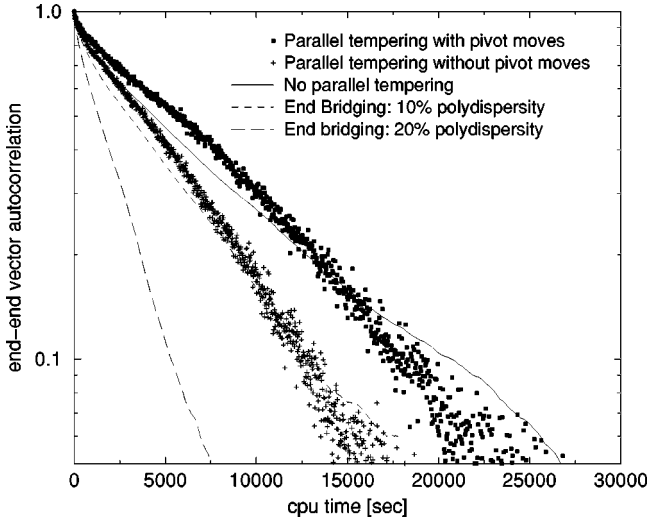

FIG. 9. Normalized end-to-end vector autocorrelation function, for chain length $N=60$ at density $\rho=0.85$, using standard Langevin dynamics, our parallel tempering procedure with and without pivot moves, and end-bridging simulations for various degrees of polydispersity.

$$C(t)=\frac{\Sigma_{t'}u(t'+t)u(t')\delta(U(t'+t),U_0)\delta(U(t'),U_0)}{\Sigma_{t'}\delta(U(t'+t),U_0)\delta(U(t'),U_0)}, \tag{13}$$

where $u(t)$ is the quantity whose autocorrelation function is determined, U_0 is the Kremer-Grest potential, and $U(t)$ is the potential at time t. The Kronecker $\delta(U(t),U_0)$ vanishes unless the potential is U_0, where it is unity.

Figure 9 displays the normalized end-to-end vector autocorrelation function for (i) end bridging with degree of polydispersity set to 10% and 20% of the chain length, (ii) our parallel tempering procedure, and (iii) standard Langevin (Rouse-like) dynamics, for our system of 20 chains of length $N=60$. It is seen that end bridging can achieve better efficiency in the relaxation of $\vec{R}$, but only if the polydispersity exceeds 10%, which must be considered a large value if one is mainly interested in the properties of an approximately monodisperse sample. Furthermore, it is most likely that the end-to-end vector does not describe the slowest relaxation in the system for this type of algorithm. Since for the end-bridging algorithm what constitutes a polymer chain becomes an ill-defined quantity, what actually does become the slowest mode is unclear. It is probably a collective quantity like the stress or similar; however, since the fluctuations in such quantities are hard to measure with good statistical accuracy, due to lack of self-averaging, we did not study this point further.

The parallel tempering procedure turns out to be somewhat faster than plain Rouse relaxation, in particular in the long-time limit. In terms of integrated autocorrelation time, the speedup amounts to roughly 10% with the pivot moves present and 30% without. The reason why performing pivot moves actually slows the simulation down is explained by Figs. 10 and 11, which show the normalized end-to-end vector autocorrelation function and the autocorrelation of the

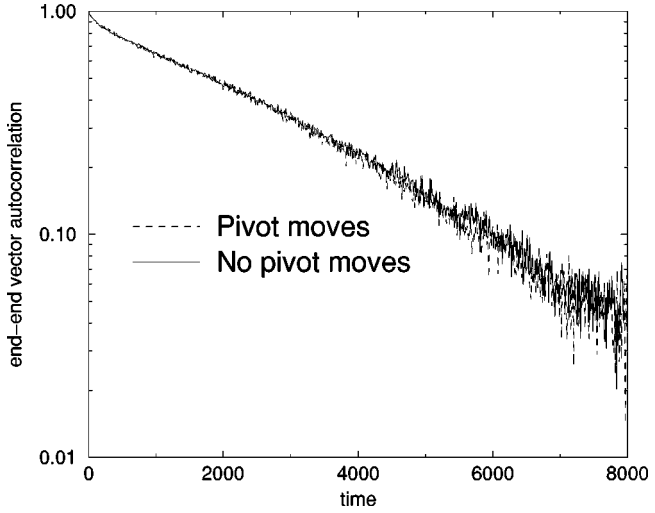

FIG. 10. Normalized end-to-end vector autocorrelation function, for chain length $N=60$ at density $\rho=0.85$, using our parallel tempering algorithm with and without pivot moves.

first Rouse mode (in terms of ''physical'' time, not CPU time). The parallel tempering alone already needs such a long time for traversing Hamiltonian space from the ideal gas to the full repulsive interaction that on this time scale the chains are already fully relaxed. After a ''diffusive loop'' through Hamiltonian space, the configuration is thus already fully decorrelated, even without pivot moves. Therefore, the pivot moves just generate additional CPU overhead and cause a slowdown.

In Fig. 12, the autocorrelation of the lowest five Rouse modes is shown as a function of the Rouse scaling variable $t\sin^2(p\pi/2N)$, such that for pure Rouse dynamics all curves would collapse onto a single line. As is known from older simulations [6,8], $N=60$ is already slightly in the crossover regime to reptation, where ultimately the lower modes are slowed down. Nevertheless, $N=60$ is still too short for this effect to become visible, such that Rouse behavior for the

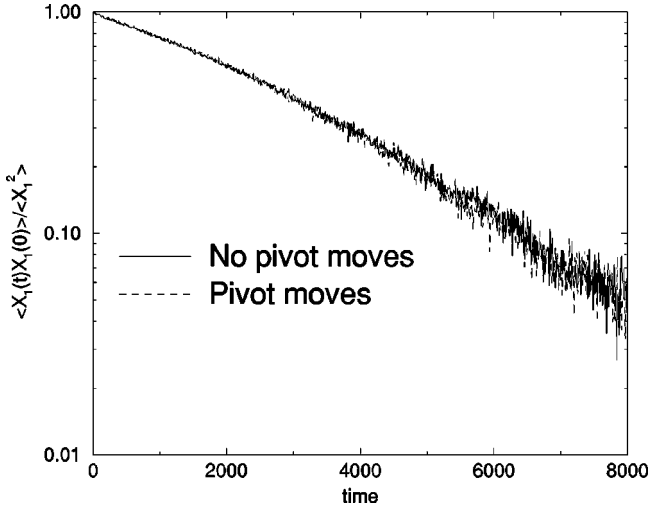

FIG. 11. Normalized autocorrelation of the first Rouse mode, for chain length $N=60$ at density $\rho=0.85$, using our parallel tempering algorithm with and without pivot moves.

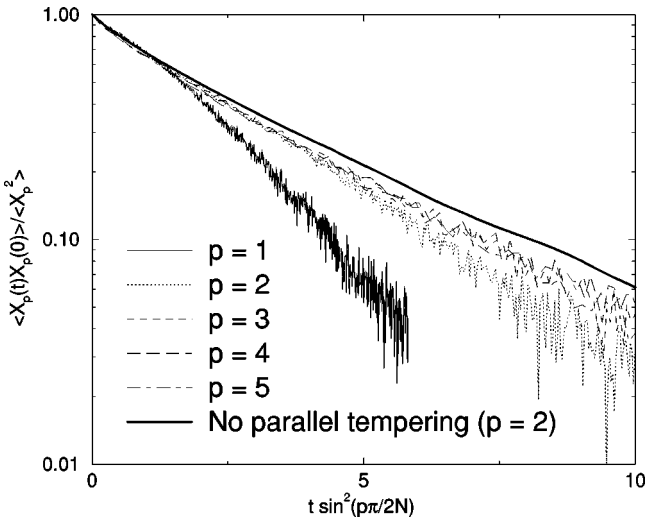

FIG. 12. Rouse mode analysis for the lowest five Rouse modes, for our parallel tempering procedure, and compared to the case of pure Langevin dynamics, for chain length $N=60$ and density $\rho=0.85$. Though we have shown only the result for the second Rouse mode, all Rouse modes fall onto the same line for pure Langevin dynamics.

case of pure Langevin dynamics can still be assumed. Our tempering procedure, on the other hand, produces a disproportionate acceleration of the lower modes. They are the only modes that are slow enough to be able to capitalize on the excursions to the softer interactions. They first relax exponentially in accord with the pure Rouse dynamics of the hard system, while at later times the decay is significantly steeper.

We were also able to obtain results, shown in Fig. 13, for a system of 32 chains of length 200, using 256 Hamiltonians. Limitations in the CPU time available to us have prevented us from performing a comparison between simulations with and without pivot moves and from measuring the correlation function until full decay. Nevertheless, our preliminary data clearly show a steep drop off in the correlation functions, which is much more pronounced than for the case of $N=60$. These data were obtained for the case *with* pivot moves; we believe that they actually did help to accelerate the equilibration of this system. The data for pure Langevin dynamics without tempering were taken from Ref. [49]. From an approximate fit to our data, we see a speedup (in terms of integrated autocorrelation time for the first Rouse mode) of greater than a factor of 8.

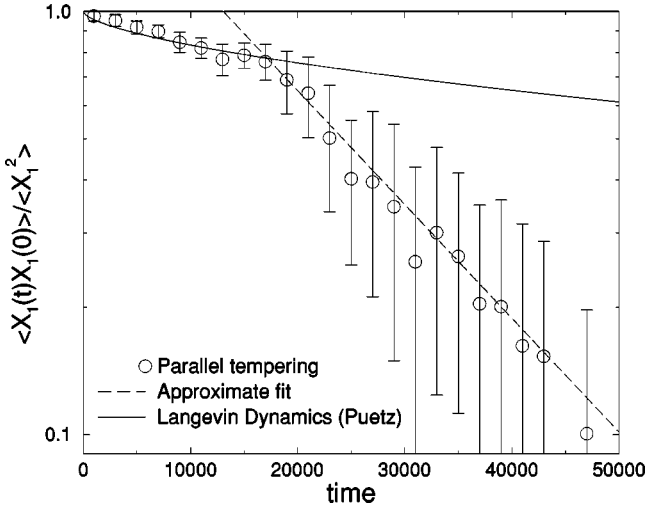

FIG. 13. Normalized autocorrelation of the first Rouse mode for our parallel tempering procedure for chain length 200 at density $\rho=0.85$. We have compared our result to a fit to results previously obtained by Pütz [49].

C. Conclusions

Our results indicate that parallel excluded volume tempering combined with large-scale chain moves is a viable route to speeding up simulations of dense polymer systems. It is expected that the method will become more and more useful as the chain length increases, as indicated from our preliminary results for chain length 200, in particular when compared to the results for shorter chains. For our initial attempt, which is most likely not the optimal choice of all simulation parameters, it seems that $N=60$ is rather close to the crossover length, while $N=200$ is significantly above. Even more dramatic speedups are expected for more complicated molecular architectures such as stars. These issues will be the subject of future investigations. Current trends in the development of computational facilities indicate that over the next decade we will see an increase in the availability of massively parallel computers with more and more processors running at approximately the speed of today's processors. With the advent of such facilities, we expect the full potential of this algorithm to be realized.

There are several directions in which this algorithm can be further developed. Through the development of an analytical understanding of the effect of chain length on transition probabilities as the phantom chain limit is approached, one could realize a general scheme to generate an optimal set of transfer radii for a particular system. Another possible development is performing the parallel excluded volume tempering by only softening a limited set of interactions in the Hamiltonian, if the original system is based upon a more fine-grained or even atomistic model. In such models, it is quite typical that a particular term in the Hamiltonian creates a significantly greater energy barrier than any other term. Since such ''hard'' interactions pertain only to a subset of the overall system, the effective system size V on which the tempering acts is reduced. This results in a smaller number of necessary processors.

One could also consider the softening of only a subvolume of the system, for example a single polymer or all sites in a certain region of the simulation box. However, such approaches need careful testing, since one must expect that the softening perturbation is not strictly local, due to long-range elastic stresses. Combining parallel excluded volume tempering with CCB may also yield a more efficient algorithm.

In its current form, our algorithm should be seen as complementary to the ''hyperparallel tempering'' algorithm of dePablo *et al.* [29]. Both approaches are geared at con-

'straint softening (via the density in the ''hyperparallel tempering,'' via direct manipulation of the interaction in our case), combined with nonlocal chain moves (chain growing versus pivot moves). Chain insertion then becomes feasible, and thus both methods are, in principle, well suited for calculating phase equilibria. The results presented in Ref. [29] for the gas-liquid transition of dense polymer systems look very encouraging. At the present stage, it seems an open question which algorithm (or perhaps a combination) will prove the most feasible and successful for such problems, in particular when dealing with very dense systems.

ACKNOWLEDGMENTS

We thank J. J. de Pablo, R. Everaers, and A. Khokhlov for helpful suggestions and stimulating discussions. A.B. thanks D. Theodorou and V. I. Mavrantzas for hospitality at the Institute of Chemical Engineering and High Temperature Chemical Processes, FORTH, Patras, Greece, where part of this work was done. This research was supported by the EU TMR network ''NEWRUP,'' Contract No. ERB-FMRX-CT98-0176.

[1] G. S. Grest, M.-D. Lacasse, K. Kremer, and A. Gupta, J. Chem. Phys. **105**, 10 583 (1996).
[2] F. A. Escobedo and J. J. dePablo, J. Chem. Phys. **105**, 4391 (1996).
[3] F. A. Escobedo and J. J. dePablo, J. Chem. Phys. **106**, 2911 (1997).
[4] F. A. Escobedo and J. J. dePablo, J. Chem. Phys. **106**, 9858 (1997).
[5] R. Everaers, New J. Phys. **1**, 12.1 (1999).
[6] K. Kremer and G. S. Grest, J. Chem. Phys. **92**, 5057 (1990).
[7] M. Kröger, W. Loose, and S. Hess, J. Rheol. **37**, 1057 (1993).
[8] A. Kopf, B. Dünweg, and W. Paul, J. Chem. Phys. **107**, 6945 (1997).
[9] W. Paul *et al.*, Phys. Rev. Lett. **80**, 2346 (1998).
[10] M. Doi and S. F. Edwards, *The Theory of Polymer Dynamics* (Clarendon Press, Oxford, 1986).
[11] B. Dünweg, G. S. Grest, and K. Kremer, in *Numerical Methods for Polymeric Systems*, edited by S. G. Whittington (Springer-Verlag, Heidelberg, 1998), p. 159.
[12] R. H. Swendsen and J. S. Wang, Phys. Rev. Lett. **58**, 86 (1987).
[13] N. Madras and A. D. Sokal, J. Stat. Phys. **50**, 109 (1988).
[14] A. D. Sokal, in *Monte Carlo and Molecular Dynamics Simulations in Polymer Science*, edited by K. Binder (Oxford University Press, New York, 1995), p. 47.
[15] J. I. Siepmann and D. Frenkel, Mol. Phys. **75**, 59 (1992).
[16] J. J. dePablo, M. Laso, and U. W. Suter, J. Chem. Phys. **96**, 2395 (1992).
[17] F. A. Escobedo and J. J. dePablo, J. Chem. Phys. **102**, 2636 (1995).
[18] D. Frenkel and G. Mooij, in *Monte Carlo and Molecular Dynamics of Condensed Matter Systems*, edited by K. Binder and G. Ciccotti (Societa Italiana di Fisica, Bologna, 1996), p. 163.
[19] S. Consta, N. B. Wilding, D. Frenkel, and Z. Alexandrowicz, J. Chem. Phys. **110**, 3220 (1999).
[20] J. P. Valleau and G. M. Torrie, in *Statistical Mechanics Part A: Equilibrium Techniques*, edited by B. J. Berne (Plenum, New York, 1977).
[21] G. M. Torrie and J. P. Valleau, J. Comput. Phys. **23**, 187 (1977).
[22] B. A. Berg and T. Neuhaus, Phys. Rev. Lett. **68**, 9 (1992).
[23] B. A. Berg, Int. J. Mod. Phys. C **3**, 1083 (1992).
[24] J. Lee, Phys. Rev. Lett. **71**, 211 (1993).
[25] B. Hesselbo and R. B. Stinchcombe, Phys. Rev. Lett. **74**, 2151 (1995).
[26] E. Marinari and G. Parisi, Europhys. Lett. **19**, 451 (1992).
[27] A. P. Lyubartsev, A. A. Martinovski, S. V. Shevkunov, and P. N. Vorontsov-Velyaminov, J. Chem. Phys. **96**, 1776 (1992).
[28] F. A. Escobedo and J. J. dePablo, J. Chem. Phys. **103**, 2703 (1995).
[29] Q. Yan and J. J. dePablo, J. Chem. Phys. **113**, 1276 (2000).
[30] M. Müller and W. Paul, J. Chem. Phys. **100**, 719 (1993).
[31] N. B. Wilding and M. Müller, J. Chem. Phys. **101**, 4324 (1994).
[32] Y. Iba, G. Chikenji, and M. Kikuchi, J. Phys. Soc. Jpn. **67**, 3327 (1998).
[33] C. J. Geyer and E. A. Thompson, J. Am. Stat. Assoc. **90**, 909 (1995).
[34] E. Orlandini, in *Numerical Methods for Polymeric Systems*, edited by S. G. Whittington (Springer-Verlag, Heidelberg, 1998), p. 33.
[35] K. Hukushima, H. Takayama, and K. Nemoto, J. Phys. Soc. Jpn. **65**, 1604 (1996).
[36] T. C. B. McLeish and S. T. Milner, Adv. Polym. Sci. **143**, 195 (1999).
[37] W. Kob, C. Brangian, T. Stühn, and R. Yamamoto, e-print cond-mat/0003282 (unpublished).
[38] R. Yamamoto and W. Kob, Phys. Rev. E **61**, 5473 (2000).
[39] Q. L. Yan and J. J. dePablo, J. Chem. Phys. **111**, 9509 (1999).
[40] M. C. Tesi, E. J. J. van Rensburg, E. Orlandini, and S. G. Whittington, J. Stat. Phys. **82**, 155 (1996).
[41] P. P. Nidras and R. Brak, J. Phys. A **30**, 1457 (1997).
[42] U. H. E. Hansmann, Chem. Phys. Lett. **281**, 140 (1997).
[43] E. Orlandini and T. Garel, Int. J. Mod. Phys. C **9**, 1459 (1998).
[44] E. J. J. van Rensburg, E. Orlandini, and M. C. Tesi, J. Phys. A **32**, 1567 (1999).
[45] A. Irbäck and E. Sandelin, J. Chem. Phys. **110**, 12256 (1999).
[46] P. V. K. Pant and D. N. Theodorou, Macromolecules **28**, 7224 (1995).
[47] V. G. Mavrantzas, T. D. Boone, E. Zervopoulou, and D. N. Theodorou, Macromolecules **32**, 5072 (1999).
[48] M. P. Allen and D. J. Tildesley, *Computer Simulations of Liquids* (Clarendon Press, Oxford, 1989).
[49] M. Pütz, Ph.D. thesis, University of Mainz, 1999 (unpublished).