High-order sampling schemes for path integrals and Gaussian chain simulations of polymers

Martin H. Müser[1,2] and Marcus Müller[3]

[1]Jülich Supercomputing Centre, Institute for Advanced Simulation, FZ Jülich, Jülich, Germany
[2]Dept. of Materials Science and Engineering, Universität des Saarlandes, Saarbrücken, Germany
[3]Institut für Theoretische Physik, Georg-August Universität, Göttingen, Germany

In this work, we demonstrate that path-integral schemes, derived in the context of many-body quantum systems, benefit the simulation of Gaussian chains representing polymers. Specifically, we show how to decrease discretization corrections with little extra computation from the usual $O(1/P^2)$ to $O(1/P^4)$, where P is the number of beads representing the chains. As a consequence, high-order integrators necessitate much smaller P than those commonly used. Particular emphasis is placed on the questions of how to maintain this rate of convergence for open polymers and for polymers confined by a hard wall as well as how to ensure efficient sampling. The advantages of the higher-order sampling schemes are illustrated by studying the surface tension of a polymer melt and the interface tension in a binary homopolymers blend.

I. INTRODUCTION

The mathematical isomorphism between the partition function of a Gaussian chain representing a polymer and the density matrix of a quantum-mechanical point particle has been known at least since Feynman formulated quantum mechanics in terms of path integrals.[1–4] To cast the partition function of Gaussian chains or quantum-mechanical point particles as path integrals, bead-spring chains comprised of P beads are embedded into an external potential and configurations occur with a probability following the usual laws of statistical physics. Differences between polymer and quantum calculations usually arise in the nature of the external potential. In the context of quantum mechanics, beads of different chains only interact if they carry the same index, at least as long as all degrees of freedom in the system are considered explicitly.[5,6] In contrast, beads representing parts of a homopolymer are commonly assumed to interact with all others irrespectively of their position along the chain. Within self-consistent field theory, these interactions are replaced by an external potential that represents the average mean interaction of a bead with its neighborhood[7,8]. In single-chain-in-mean-field (SCMF) simulations the interactions among beads are temporarily replaced by an external potential that approximates the interaction of a bead with its instantaneous surrounding[9,10].

Despite the just-mentioned differences between quantum and polymer chains, their simulations pursue similar goals and face related difficulties.[11–13] To obtain converged results or results that mimic the universal behaviour of long polymer chains, the discretization of the chain contour, P, should always be as large as possible. However, increasing P implies more computing time per chain. Even worse, the stiffness of springs connecting adjacent beads increases with P, which, in turn, gives rise to a time-scale separation between stiff bonded and weak non-bonded interactions[14] and affects the sampling efficiency of many algorithms. The ensuing impediments can

be quite dramatic when a complex task is handled with naive, yet, frequently used algorithms. As an example, we discuss by how much the necessary computing time increases if one wants to reduce a given systematic "finite-P" error by a factor of two when simulating bead-spring chains in the presence of a hard wall. We assume that the goal is to keep the stochastic error constant and that no tricks of the trade are used — single-bead moves in either Monte-Carlo (MC) or molecular dynamics (MD). For hard walls, results converge only with $1/\sqrt{P}$[15] so that P needs to increase by a factor of four to halve the error. Using only single-bead moves, the slowest mode in the system has a correlation time proportional to P^2 in terms of CPU sweeps according to the Rouse model,[16] which means we need to run 16-times more global sweeps, even when dynamics are underdamped.[17] In total, this means 64 times the computing effort, i.e., the numerical effort scales proportional to the inverse sixth power of the acceptable error.

In both the quantum and the polymer communities, important advances have been made to improve convergence of path integrals. For example, the hard-wall problem has been solved in a similar fashion for quantum particles as for Gaussian chains by finding solutions for the free chain in the presence of one or two walls.[15,18,19] However, some advances have been made for path integrals in the quantum community that do not yet have appeared in the context of polymers. Most notably, easy-to-code decompositions of the high-temperature density matrix have been derived that lead to discretization errors, which disappear as $1/P^4$ in the case of cyclic chains.[20,21] In numerical self-consistent field theory, a fourth-order accuracy can be achieved in solving the modified diffusion equation that describes the single chain in an external field by combining a second-order pseudospectral algorithm with Richardson extrapolation.[22,23] This technique, however, cannot be carried over in a straightforward fashion to a particle-based description that is employed in partial enumeration techniques[24,25] and SCMF simulations.[10]

In addition, a sampling algorithm called "staging" was proposed 30 years ago to be used in path-integral Monte Carlo simulations.[26] It allows one to resample a given fraction of a bead-spring chain in an external potential at a roughly constant acceptance ratio, no matter how large P. Slowing down due to Rouse dynamics is thereby completely eliminated. If both concepts could be combined, one would only need $2^{1/4}$ rather than 2^6 times the computing time, when decreasing the error of our chain in front of a hard wall by a factor of two. However, convergence can also be dramatically improved when no hard walls are present.

In this work, we study to what extent the advances on the higher-order density-matrix decomposition benefit the simulation of linear Gaussian chains, including polymers adjacent to a hard wall and narrow interfaces in strongly segregated polymer blends. The good convergence for closed chains is not automatically maintained for open chains, because, as we demonstrate in this paper, particular caution has to be taken for the acquisition of averages over an open chain. Moreover, we describe how to adopt the staging algorithm in the presence of a hard wall. The advantage of a hard-wall adjusted staging move is that no more trial moves can be rejected. The latter frequently happens when the effect of the hard wall on the finite-P chain is reflected by the presence of an effective (P-dependent) potential. As a consequence of the hard-wall staging algorithm, large fractions of a chain can be sampled in a single step with good acceptance ratio.

Our manuscript is arranged as follows: In section II, we describe high-order density-matrix decompositions and the staging algorithm including our extensions of these schemes. Results are presented in section III and conclusions are drawn in section IV.

II. THEORY AND METHODS

A. Approximations to path integrals

We consider the one-dimensional path-integral

$$I(\tau, x_0, x_1) \equiv C \int_{x(0)=x_0}^{x(\tau)=x_1} \mathcal{D}[x(t)] \exp\{-\mathcal{H}[x(t)]\} \quad (1)$$

with

$$\mathcal{H}[x(t)] = \frac{1}{\tau} \int_0^\tau dt \left[\frac{1}{2} \left\{ \frac{\tau}{\lambda} \frac{\partial x(t)}{\partial t} \right\}^2 + v\{x(t)\} \right]. \quad (2)$$

Here, $x(t)$ is a path having $x(t = 0) = x_0$ as starting point and $x(t = \tau) = x_1$ as end point. t is the variable on which the paths depend. It runs from 0 to τ, while λ is a parameter controlling to what extent paths want to localize. The prefactor C is a normalization constant, which can depend on τ and λ but not on the realization of the paths. It is irrelevant for the purpose of this study.

We abstain from giving a detailed account of path-integrals in the context of quantum mechanics or polymer physics and instead refer to pertinent literature.[10,27] It shall suffice to state the interpretation of various terms. $0 \le t \le \tau = \hbar/k_B T$ is called imaginary time in quantum physics, where $\hbar$ is the reduced Planck constant and $k_B T$ the thermal energy, while in polymer physics t tends to be chosen as a unitless variable indexing the path of a linear polymer with a dimensionless number $0 \le t \le \tau = 1$. The parameter λ corresponds to $\hbar/\sqrt{m k_B T}$ in quantum physics, which is the thermal de Broglie wavelength divided by $\sqrt{2\pi}$, m being the mass of the particle, or $\lambda = R_{EE}/\sqrt{D}$ in polymer physics, where R_{EE} is the root-mean-square end-to-end radius of a free polymer embedded in D-dimensional space. Moreover, $v(x)$ is the external potential energy per length measured in units of the thermal energy $k_B T$. In the following we measure all energies in units of $k_B T$. Lastly, $I(\tau, x_0, x_1)$ corresponds to either the real-space representation of the density matrix of a quantum point particle or to the probability density of a polymer to have its head and tail monomer be located, respectively, at x_0 and x_1.

In the following, we call the integrand on the r.h.s of equation (1) the propagator in imaginary time, or simply the propagator. Numerical evaluations of path integrals require one to factorize the propagator into short-time propagators and to find approximations for them. These approximations then translate into discretization recipes for the r.h.s. of equation (2). Factorization can be achieved, for example, by (recursively) making use of the identity

$$I(\tau, x_0, x_1) \propto \int_{-\infty}^\infty dx' \, I\left(\frac{\tau}{2}, x_0, x'\right) I\left(\frac{\tau}{2}, x', x_1\right), \quad (3)$$

or by factorizing the original propagator into P terms rather than into two terms as done implicitly in equation (3). As a brief side remark, we note that the proper choice of the normalization constant C allows one to replace the proportionality sign in equation (3) or related factorizations with an equality sign.

In Dirac notation, one can write

$$I(\Delta\tau, x_t, x_{t+\Delta\tau}) = \langle x_t | e^{-\hat{\mathcal{H}}/P} | x_{t+\Delta\tau} \rangle, \quad (4)$$

with an effective Hamiltonian $\hat{\mathcal{H}}$

$$\hat{\mathcal{H}} = \frac{\lambda^2 \partial_x^2}{2} + \hat{v}, \quad (5)$$

where ∂_x denotes a partial derivative with respect to x and $\hat{v}$ the potential energy operator. Moreover, $\Delta\tau = \tau/P$. So far, no approximations have been made. However, to evaluate the r.h.s. of equation (4), which one may want to call the short-time propagator, one can write

$$e^{-\hat{H}/P} = e^{-\hat{v}/2P} e^{-\lambda^2 \partial_x^2/2P} e^{-\hat{v}/2P} + O(1/P^3), \quad (6)$$

where the $O(1/P^3)$ error is valid for potentials whose second-order derivative is not divergent.[28,29] Inserting

this approximation into equation (4) and evaluating it yields,

$$I\left(\Delta\tau, x_t, x_{t+\Delta\tau}\right) \approx Ce^{-\{v(x_t)+v(x_{t+\Delta\tau})\}/2P}$$
$$\times e^{-P(x_t-x_{t+\Delta\tau})^2/2\lambda^2} \qquad (7)$$

which can be used to identify a finite-difference approximation of summands in equation (2) as

$$\mathcal{H}[x(t)] = H(x_0, \cdots, x_P) + O(1/P^2)$$
$$\approx \frac{P}{2\lambda^2}\sum_{n=1}^{P}(x_n - x_{n-1})^2 + \frac{1}{P}\sum_{n=0}^{P}w_n v(x_n) \quad (8)$$

where we now write the end-point of the path as x_P. The weights for the energies are $w_n = 1$ except $w_0 = w_P = 1/2$, i.e., the potential energy is integrated along the path according to the trapezoidal rule. In the following, the term "primitive propagator (PP)" relates to a path integral whose integrand is discretized in the just-described way, i.e., according to Eq. (8).

This equation corresponds to a discretization of the Edwards Hamiltonian of a Gaussian chain, i.e., a chain consists $P + 1$ beads that are connected via P harmonic springs. Full beads are placed between two adjacent springs, whereas the terminal beads are only "half" beads. In the context of polymer physics, the standard discretization of the Edwards Hamiltonian (DEH) is slightly different, i.e., one gives each bead an identical weight, which would translate into $w_n = P/(P + 1)$.

More accurate decompositions of the density operator than that in equation (6) have been proposed in the past few decades,[29,30] most notably by Takashi and Imada.[31] A disadvantage of their and related approaches is that taking mean values of observables turns out more elaborate than when using the primitive propagator. Recently, Chin[20] proposed a new decomposition of the short-time propagator, which, like Takahshi-Imada propagator,[31] has leading errors that scale as $O(1/P^4)$ but does not necessitate elaborate redefinitions of observables, as pointed out by Jang et al.[21] Defining

$$\hat{t} = \lambda^2 \partial_x^2/2, \qquad (9)$$

the decomposition reads

$$e^{-2\hat{H}/P} = e^{-\hat{v}/3P}e^{-\hat{t}/P}e^{-4\hat{v}_{\mathrm{g}}/3P}e^{-\hat{t}/P}e^{-\hat{v}/3P}, \quad (10)$$

where the operator $\hat{v}_{\mathrm{g}}$ is the operator with the eigenvalue

$$v_{\mathrm{g}}(x) = v(x) + \frac{1}{12P^2}(\partial_x v)^2. \qquad (11)$$

for a given $|x\rangle$. The pertinent expression for the discretized version of $\mathcal{H}[x(t)]$ can be obtained in a similar fashion as for the primitive propagator. They are summarized in section II B

For ring polymers or cyclic chains, $x_0 = x_P$, observables can be defined only on "even beads" (given that we start enumerating monomers with $n = 0$) according to

$$\langle O\rangle_{\mathrm{CC}} = \frac{2}{P}\sum_{n=0,2,..}^{P-2}\langle O(x_n)\rangle_{\mathrm{HOA}} + O(1/P^4), \qquad (12)$$

where the index CC stands for cyclic chains and HOA for an estimate based on the higher-order approximation propagator of Eq. (10). As we discuss in the following, Eq. (12) is not valid for open chains.

To substantiate this claim, it is easiest to consider an example. Fig. 1 shows how the mean external potential energy of individual beads depends on their index n, or to be more precise, on the ratio n/P. In the given example, we consider a chain with $\lambda = \tau = 1$ and an external harmonic potential $v(x) = kx^2/2$ with $k = 64$. One can see that the even-numbered beads of the HOA calculation with contour discretization $P = 16$ correspond quite closely to the converged results. In fact, further analysis confirms that errors on even-number beads are $O(1/P^4)$. In contrast, odd-numbered beads as well as the calculations using the primitive-propagator with the same chain discretization, $P = 16$, deviate from the converged results in a clearly visible fashion.
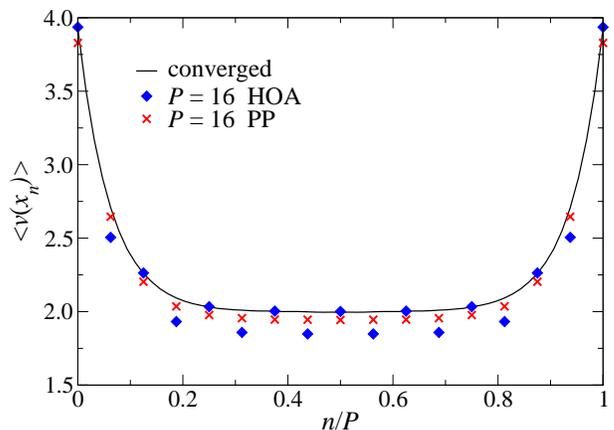

FIG. 1: Mean external potential energy of beads as a function of their index n for $\lambda = \tau = 1$ and $v(x) = kx^2/2$ with $k = 64$. The full line shows converged results for $P \gg 1$ while symbols show data for $P = 16$. Blue diamonds represent averages obtained with a higher-order approximation (HOA) to the propagator, while red crosses refer to those obtained with the primitive propagator (PP).

To obtain accurate estimates for the mean external potential $\langle v\rangle$, whose error only increases proportional to $1/P^4$, one needs to sum over the $\langle v\{x(t = n/P)\}\rangle$ curve with an algorithm that converges sufficiently fast. Since the trapezoidal rule leads to $O(1/P^2)$ errors, we propose to resort to (a composite) Simpson's rule, which is sufficiently accurate to allow for an $O(1/P^4)$ convergence. This leads to expressions that differ from Eq. (12), which are also summarized in section II B

Finally, we note that the review of higher-order approximations and the proposition to use Simpson's rule

$n\%4$	w_n	f_n^{c}	w_n^{m}
0	2/3	0	4/3
1	4/3	1	0
2	2/3	0	8/3
3	4/3	1	0
except			
$n = 0$	1/3	0	2/3
$n = P$	1/3	0	2/3

TABLE I: Value of weighting coefficients as a function of the bead index n. The modulo function is denoted as $n\%4$. The enumeration starts at $n = 0$ and ends at P, which must be an integer multiple of 4. w_n gives the weight of the energy in the Boltzmann factor. f_n^{c} states if the square-gradient corrections needs to be added to $v(x)$ to yield $v_{\mathrm{g}}(x)$. Lastly, w_n^{m} states the weight of an observable in a measurement.

for the measurement of observables defined on open chains only relate to non-grafted polymers. Once chains are tethered with one head-group to a wall, one may have to simulate them differently than non-grafted chains in order to maintain $O(1/P^4)$ convergence. This issue will be addressed in a separate study.

B. Implementation of higher-order approximations

While the derivation of higher-order approximations to path integrals may be seen as somewhat cumbersome, its implementation can be achieved in a rather straightforward fashion. In this section, we explain what needs to be done if a conventional code, e.g., one based on the primitive propagator, is present.

First, no differences appear as to the handling of the springs connecting two beads. Second, for each bead, we need to evaluate the potential $v(x_n)$ as before. Third, we need to evaluate the square-gradient correction to $v(x_n)$ only on beads with an odd index given that we start enumerating beads with $n = 0$. For self-consistent theory calculation on polymers this should not be an expensive procedure, because derivatives can be evaluated numerically at moderate expense. Fourth, measurements are only taken on even beads and the weights of an observable evaluated at x_n follow from the composite Simpson's rule. In all cases, except for the head and the tail monomer, every bead with an index that is an integer multiple of four is both a starting point as well as an end point of Simpson rule. Thus, the value of P needs to be an integer multiple of four, as Simpson's rule necessitates two starting point, a mid point, and an endpoint. We summarize the results in table I.

The resulting expression for observables defined as integrals over open chains (OC) is

$$\langle O \rangle_{\mathrm{OC}} = \frac{1}{P} \sum_{n=0}^{P} w_n^{\mathrm{m}} \langle O(x_n) \rangle_{\mathrm{HOA}} + O(1/P^4), \qquad (13)$$

where the weights are taken from Tab. I. The rule to obtain averages for observables is also valid for the external potential, although it differs from the expression entering the Boltzmann factor. The latter is

$$v[x(t)] = \frac{1}{P} \sum_{n=0}^{P} w_n \left\{ v(x_n) + f_n^{\mathrm{c}} \Delta v_{\mathrm{g}}(x_n) \right\}, \qquad (14)$$

which replaces the second summand of Eq. (8) in calculations based on the higher-order approximation to the propagator. Here, $\Delta v_{\mathrm{g}}(x)$ refers to the square-gradient correction $v_{\mathrm{g}}(x) - v(x) = \{\partial v(x)/\partial x\}^2/12P^2$.

C. Free-particle staging

As the continuum approximation of the path is increasingly better realized for larger P, numerical approaches solely based on single-bead dynamics quickly slow down. The reason is that the correlation time, as measured in MD time steps or in number of MC sweeps, increases with P^2 (Rouse dynamics). A similar problem is known from numerical evaluations of path integrals, for which an efficient sampling algorithm, called staging, was proposed three decades ago.[26] It overcomes the problem that springs connecting two adjacent beads become stiffer thereby making single-bead MC trial moves more inefficient with an increasing number of beads representing the path.

The idea of the staging technique is to make a trial move of a chain segment such that a single trial move has the (exact) stochastic properties of a segment that is not coupled to external potentials. The relative probability of the trial path $x_{\mathrm{trial}}(t)$ and the existing path $x(t)$ then becomes

$$\mathrm{Pr_r} = \exp\left(-\int_{\tau_1}^{\tau_2} dt \left\{ v[x_{\mathrm{trial}}(t)] - v[x(t)] \right\} \right), \qquad (15)$$

which in turn can be used as a transition probability in the Metropolis algorithm.

The sampling of free segments proceeds as follows: Assume that the fixed end points t_1 and t_3 both lie within the Gaussian chain, i.e., within the interval $[0, \tau]$, see equation (1). The effective interaction of the midpoint bead at $t_2 = (t_1 + t_3)/2$ is a harmonic coupling to the beads at t_1 and t_3, each time with an effective inverse spring constant of $k_{12}^{-1} = k_{23}^{-1} = (t_3 - t_1)\lambda/2$. Since the two springs act in parallel, their action adds to a net inverse spring constant of

$$k^{-1} = (t_3 - t_1)\lambda/4. \qquad (16)$$

Thus, the center bead will fluctuate around the center of mass of the two fixed outer beads with a normal distribution whose second moment is k^{-1}, or $k_{\mathrm{B}}T/k$ if $k_{\mathrm{B}}T$ is not chosen as unit for the thermal energy. One can realize the correct distribution by assigning the following

value to $x_{\text{trial}}(t_2)$

$$x_{\text{trial}}(t_2) = \frac{x(t_1) + x(t_3)}{2} + \frac{u_{\text{G}}}{\sqrt{k}}, \qquad (17)$$

where u_{G} is a Gaussian random variable of zero mean and a standard variance of one. The same procedure also holds for arbitrary physical dimension D. Using variables of the Gaussian chain description, trial coordinates are assigned through

$$x_{\text{trial}}^{\alpha}(t_2) = \frac{x^{\alpha}(t_1) + x^{\alpha}(t_3)}{2} + \sqrt{\frac{t_3 - t_1}{D\tau}} \frac{u_{\text{G}}^{\alpha}}{2R_{\text{EE}}}, \quad (18)$$

where for each Cartesian component α an independent Gaussian random variable u_{G}^{α} must be drawn.

We now have a new fixed point at $t = t_2$ and two new midpoints located at, respectively, $(t_1 + t_2)/2$ and $(t_2 + t_3)/2$. These new midpoints can be assigned trial coordinates with the just described scheme, in which, however, trial coordinates need to be used as fixed points on the r.h.s. of Eq. (17). Remaining points in the interval $t_1 < t < t_3$ can be assigned by iteration.

If a moved chain segment includes an open end of a polymer, the trial move of the end bead can be done similarly as before. The only difference is that the end bead only couples to one starting bead with one effective spring. Thus, for end polymers, we need to modify equation (18) to

$$x_{\text{trial}}^{\alpha}(\tau) = x^{\alpha}(t_1) + \sqrt{\frac{\tau - t_1}{D\tau}} \frac{u_{\text{G}}^{\alpha}}{R_{\text{EE}}}, \qquad (19)$$

In practice, the just given recipe can be implemented in a straightforward fashion, when choosing the indices of the two outer beads, $n_{1,3} = t_{1,3}P$, such that they differ by an integer power of two. Note that this does not mean that P itself has to be an integer power of two. It is certainly possible to make staging moves such that the end points are not separated by an integer power of two through an appropriate generalization of the algorithm. However, its description is slightly more involved and its action does not speed up simulations to any significant degree.

The implementation of the staging technique — with or without hard-wall corrections — for the current problem is done as follows: We first pick a monomer in the chain with $n_1 = \text{int}(uP)$, where $0 < u \leq 1$ is a uniform random number. (Monomers are numerated from 0 to P.) We then pick n_3 by adding or alternatively subtracting $2^{s_{\max}}$ from n_1, where $s_{\max}$ is a positive integer. If n_3 lies outside the chain, we shift the chosen segment such that it includes the ends, for example, for $n_3 > P$, we move the segment such that the new $n_3 = P$ corresponds to an end-bead, which then gets resampled as well. Otherwise, we proceed as described above. In one sweep, $P/2^{s_{\max}}$ such moves are done, which means that we evaluate, on average, the coupling to an external potential once for each monomer. The integer $s_{\max}$ is usually chosen as large as possible so that an acceptance rate of greater 10% is still achieved. Depending on the nature of the problem, e.g., when the curvature of the external or self-consistent potential varies quickly, it might be beneficial to run a "V"-cycle, in which $s_{\max}$ is varied between one and an appropriately chosen maximum.[36]

D. Hard-wall adopted staging

Path-integral simulations can be made not only more efficient but also more quickly converging by introducing appropriate trial moves. For example, the propagator for a particle in a box can be corrected systematically through approximations in terms of winding numbers, see, e.g., Eqs. (11)–(12) in Ref.[15]. Similarly, one can correct the path integrals for Gaussian chains in front of a hard wall by solving the diffusion equation in the presence of a reflecting wall.[18] The latter can be obtained within a few lines by proceeding as in Ref.[15], i.e., the (short-time) free-particle propagator follows from

$$I(\Delta\tau, x_1, x2) \equiv \langle x_1 | \exp\left(-\hat{t}/P\right) | x_2 \rangle \qquad (20)$$
$$= \int_{-\infty}^{\infty} dk \Psi_k(x_1) \Psi_k(x_2) \exp(-e_k \Delta\tau/P),$$

where the kinetic energy operator (still expressed in units of $k_{\text{B}}T$) was defined in equation (9). Here, the $\Psi_k(x)$ are eigenfunctions of a free particle, $e_k = k^2/2$ being their eigenenergies in the appropriate unit system.

To turn the free-particle propagator into a hard-wall propagator, we may only integrate over those eigenfunctions that satisfy the boundary condition on the walls, i.e., $\sin(kx)$ rather than $\exp(ikx)$. Inserting the $\sin(kx)$ into equation (21) and integrating over non-negative values of k yields

$$\rho(x_1, x_2, \sigma) \propto \sinh\left(\frac{x_1 x_2}{\sigma^2}\right) \exp\left\{-\frac{x_1^2 + x_2^2}{2\sigma^2}\right\}, \quad (21)$$

with

$$\sigma = \frac{\lambda}{\sqrt{P}}. \qquad (22)$$

The proportionality in Eq. (21) can also be cast in terms of the product of a hard-wall correction factor and the free-particle propagator:

$$\rho(x_1, x_2, \sigma) \propto \left(1 - e^{-2x_1 x_2/\sigma^2}\right) e^{-(x_1 - x_2)^2/2\sigma^2}. \quad (23)$$

In the context of a particle diffusing in front of a reflective wall located at $x = 0$, Eq. (21) gives the (unnormalized) probability (density) of a random walk to end up at x_2 given that it started at x_1. We can therefore use it to sample the head or tail-monomer of a Gaussian chain, i.e., we use it to replace equation (19) whenever x_1 does not distinctly exceed σ. Details of how this can be achieved are given further below.

In addition to the end-point propagator, we also need the one for "mid-point beads." The probability of a mid-point bead to end up at $x_2 = x(t_2)$ is proportional to $\rho(x_1, x_2)\rho(x_2, x_3)$ given that $x_1 = x(t_1)$ and $x_3 = x(t_3)$ are fixed and $t_2 - t_1 = t_3 - t_2$, i.e.,

$$\rho_{\mathrm{m}}(x_1, x_2, x_3, \sigma) \propto \sinh\left(\frac{x_1 x_2}{\sigma^2}\right) \sinh\left(\frac{x_2 x_3}{\sigma^2}\right)$$
$$\times \exp\left\{-\frac{x_1^2 + 2x_2^2 + x_3^2}{2\sigma^2}\right\}. \quad (24)$$

The probability distribution ρ_m deviates from a simple Gaussian if neither x_1 nor x_3 are much greater than σ. We note in passing that when comparing equation (24) with equations from the previous section, one needs to keep in mind that σ relates to $t_2 - t_1$, whereas k related to $t_3 - t_1$. This explains why there is a prefactor of two in front of x_2^2 in the exponential on the r.h.s. of equation (24) rather than a prefactor of four.

Once the relative probabilities for mid-point beads and end-point beads are known, they can be used to generate trial moves that are distributed exactly according to the distribution of a polymer in front of a hard wall. Unfortunately, only a selected number of distributions can be drawn directly. One possibility to overcome this problem is to incorporate the exact probabilities in terms of an effective, external potential as done, for example, in references[18] and[19]. In this case, the effect of the hard walls on the propagator is discarded in the generation of the trial path and instead moved to the Metropolis step. The disadvantage of this solution is that it risks to decrease the fraction of the chain that can be resampled with a reasonable acceptance rate. For short chains, this may not have to be a large drawback. In the context of quantum mechanics, where values of P beyond 1,000 are not exceptional, this would constitute a serious drawback. However, even for chains as short as $P = 16$, significant gains can be made if one can resample the full chain, i.e., 16 monomers, instead of making single or double monomer moves. This is why we choose to implement hard-wall adopted staging moves.

We first discuss how to draw head or tail monomers. As already mentioned, once $x_1 \gg \sigma$, one can simply draw from a Gaussian distribution. In our implementation, we considered $x_1 \geq 6\sigma$ to satisfy that criterion for all practical purposes. At the other end, once $x_1 \ll \sigma$, one can approximate the expression $\sinh(x_1 x_2/\sigma^2)$ linearly without affecting the probability distribution in a noticeable way. This time, we see $x_1 \lesssim \sigma/10$ as being sufficient for x_1 to be called very small. One can then draw x_2 via the heat-bath algorithm, i.e., by solving the equation

$$u = \int_0^{x_2} dx_2' \Pr(x_2'), \quad (25)$$

where $\Pr(x_2')$ is the normalized probability distribution and u is a uniform random variable on $]0, 1[$. In the limit $x_1 \ll \sigma$, where $\Pr(x_2) \approx (x_2/\sigma^2)\exp(-x_2^2/2\sigma^2)$, one then obtains

$$x_2 = \sigma\sqrt{-2\ln u}. \quad (26)$$

For $x_1/\sigma = O(1)$, i.e., outside the two regimes just discussed, one must resort to alternative strategies. One possibility would be to solve the heat bath equation numerically, e.g., by determining the probability distribution and its integral numerically. However, we found it more effective to run a small MC simulation with which to generate the trial coordinate. The idea is that the true distribution function is approached extremely fast if one can draw (at low CPU-time cost) from a distribution that closely resembles the exact one. We moved the description of this procedure into appendix A, because it might distract from the central content of this section and more importantly because the underlying idea may be interesting in its own right.

To generate trial coordinates for mid-point beads, we proceed in a similar fashion as for end-points, i.e., we draw trial coordinates from the distribution that is exact in the absence of an external potential. This time, we first identify the location x_{c} of the maximum of the true distribution function, equation (24). If $x_{\mathrm{c}} > 6\sigma$, we can safely sample from a Gaussian, while we resort to drawing x_{c} from a short MC simulation otherwise. Technical details are presented in appendix A, where we also demonstrate that exact distributions are reproduced to high accuracy after only one or two MC time steps.

III. RESULTS

A. Convergence tests

1. Open chain in an harmonic potential

To test the convergence of different integration schemes for typical situations, we consider the generic example of an open chain in an harmonic potential. For this purpose, we keep the example from section II A, in which the curvature of the external potential is 64 times the end-to-end stiffness of the (one-dimensional) polymer. Fig. 2 reveals that the simulation result are consistent with our expectations about finite-P discretization errors: The observed deviations appear to scale with $O(1/P^4)$ for our HOA-based algorithm while PP-based approaches only achieve $O(1/P^2)$. In addition, we find that HOA-based simulations reduce to $O(1/P^2)$ if measurements of the energy are performed with the trapezoidal rule rather than with a higher-order integration scheme, i.e., when assigning the measuring weight of zero on odd beads, one on start and end bead, and two on all other even beads.

The next convergence test is motivated by applications, in which a hard wall is present. Hard walls induce a polymer depletion close to their surfaces, which in turn leads to a strong, effective attraction towards the wall. The interaction can usually be approximated by a function $v(x) = v_0\{1 - \tanh^2(x/\zeta)\}$, where v_0 and ζ represent energy and length scales, respectively. In practical applications, the curvature of the attractive interaction, v_0/ζ^2 is many times the end-to-end stiffness. We mimic this sit-
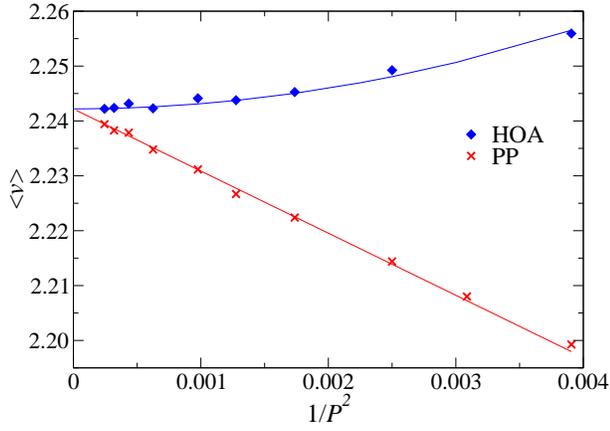
FIG. 2: Convergence of the mean external energy $\langle v \rangle$ of the open chain introduced in Fig. 1 as a function of $1/P^2$, where P is the number of springs of the chain. Blue diamonds and red crosses refer to simulations that are based on, respectively, a higher-order approximation (HOA) or the primitive propagator (PP) algorithm. Lines reflect fits according to $v_0 + c_1/P^4$ (HOA) and $v_0 + c_2/P^2$ (PP), where v_0, c_1, and c_2 are fit coefficients. Symbol size corresponds to error bars.
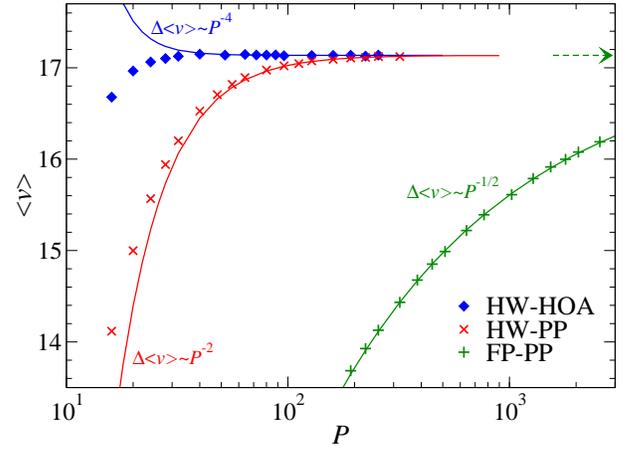

FIG. 3: Convergence of the mean external energy $\langle v \rangle$ of an open chain in a harmonic trap right next to a wall. Blue diamonds and red crosses refer to simulations that are based on, respectively, a higher-order approximation (HOA) or the primitive propagator (PP) algorithm. These two sets of simulations are produced with hard-wall (HW) adopted staging moves. Green plus signs represent free-particle (FP) simulations produced with the primitive propagator (PP). Lines reflect fits according to $v_0 + c_1/P^4$ (HW-HOA), $v_0 + c_2/P^2$ (HW-PP), and $v_0 + c_3/\sqrt{P}$ (FP-PP), where v_0, c_1, c_2, and c_3 are fit coefficients. Fits are conducted for values of P, above which the asymptotic behavior is observed. The green arrow indicates the asymptotic value of the FP-PP fit.

uation by running simulations of a harmonic trap, whose minimum is located right on the surface of the wall. In the new set of calculations, we set the curvature of the harmonic potential to 512 times the end-to-end stiffness.

Fig. 3 reveals that the error analysis is again consistent with the simulations, i.e, the slowly vanishing $O(1/\sqrt{P})$ hard-wall corrections are clearly revealed, as long as no hard-wall corrections are included. Corrections disappear with $1/P^2$ once the appropriately adopted staging moves are employed. The situation is a little less clear cut when adding the higher-order corrections to the hard-wall propagator. The problem is that the asymptotic convergence does not set in earlier than for the less accurate integration scheme. This behavior is generic for both higher-order approximations and integration schemes diagonalizing part of the Hamiltonian prior to the simulation.[32] Part of the reason is that $O(1/P^6)$ corrections need to be small in order for the $O(1/P^4)$ correction to become visible. As a consequence, the leading-order corrections may already be very small and thus within numerical noise, when asymptotic convergence sets in. We note in passing that the combination of hard-wall and HOA propagators only yields a $1/P^4$ scaling if the gradient of the potential disappears at the surface of the wall. Since this should be the case for most practical applications, we do not invest efforts into adopting the HOA scheme for hard-wall propagators beyond their simple combination.

It remains to be discussed if one can ascertain the required value for P prior to a Gaussian chain calculation. This question is answered most easily in the framework of the quantum harmonic oscillator. In that case, one is close to being converged when $k_B T P \gtrsim \hbar \omega_0$, where ω_0 is the eigenfrequency of the oscillator, which is propor-

tional to $\sqrt{k}$, where k is the curvature of the harmonic potential. This is because for thermal energies exceeding $\hbar \omega_0$, the quantum harmonic oscillator resembles its classical counter part quite closely. For the Gaussian chain calculation, one sets $k_B T = 1$ so that choosing P proportional to $\sqrt{k}$ should bring one close to convergence. Thus, it should not be necessary for the springs connecting two adjacent beads to be as stiff as the curvature of the potential. This conclusion is supported from the convergence tests presented in this section. Data shown in Fig. 2 correspond to $k = 64$ and deviations from converged results are already relatively small for $P = 16$, that is, for $P = 2\sqrt{k}$ — assuming our unit system of $\hbar = m = k_B T = 1$ for quantum or $R_{EE}^2/D = k_B T = 1$ for polymer systems. For the hard-wall adopted primitive propagator, we observe that deviations from the exact result are again $\gtrsim 1\%$ for $P = 2\sqrt{k}$ in case of the $k = 512$ system presented in Fig. 3. If adding the higher-order corrections, results are very close to the exact values at much smaller values of P. This, however, is likely a fortuitous cancellation of errors, as higher-order corrections appear to have the opposite sign of the leading term. To conclude, once a reasonable estimate for the maximum curvature of the external or self-consistent potential is known, it is straightforward to choose P.

2. Surface tension of a polymer melt

Using self-consistent field calculation in conjunction with a partial enumeration scheme to compute the single-chain partition function, we study the surface properties of a dense polymer melt in contact with a solid wall in the grandcanonical ensemble.[19] The mean-field interaction in units of $k_B T$ that a bead at a distance x away from the wall experiences is given by

$$\frac{v(x)}{P} = \frac{\kappa N}{P} \left[\phi(x) - 1 \right] \quad (27)$$

where the dimensionless quantity $\phi(x)$ is the local bead number density normalized by its value in the bulk. κN is proportional to the inverse isothermal compressibility and quantifies the repulsion of the chain molecules. In the numerical calculations we use $\kappa N = 50$, which is typical for SCMF simulations.[10]

To compute the single-chain partition function, we generate a large number $N_{\mathrm{conf}} = 10^7$ or $5 \cdot 10^7$ of one-dimensional random walks with independent Gaussian steps of mean-squared length $\langle b_x^2 \rangle = R_{\mathrm{EE}}^2 / (3P)$. The first bead of each chain molecule is randomly placed in the interval $0 \leq x \leq L_x = 2R_{\mathrm{EE}}$. If any bead lays outside this interval the *a priori* statistical weight of this chain configuration c will be zero, $w_{c,0} = 0$; otherwise $w_{c,0} = 1$. The spatial coordinate $0 \leq x \leq L_x/2$ is discretized into $N_x = 128$ uniform intervals (slabs), and reflecting boundary conditions are imposed at $x = L_x/2$.

In the HW-HOA scheme we initially compute for each chain conformation, c, the following fixed single-chain slab occupancies according to Tab. I

$$\rho_{c,0}(i_x) = \frac{1}{P} \sum_{n=0,2,\cdots} w_n \Theta_{i_x}(x_n, i_x) \quad (28)$$

$$\rho_{c,1}(i_x) = \frac{1}{P} \sum_{n=1,3,\cdots} w_n \Theta_{i_x}(x_n, i_x) \quad (29)$$

$$\rho_{c,m}(i_x) = \frac{1}{P} \sum_n w_n^m \Theta_{i_x}(x_n, i_x) \quad (30)$$

where Θ_{i_x}, $i_x = 0, \cdots, N_x - 1$, is the characteristic function of slab i_x. Moreover we compute for each chain its HW-weight[18], see also Eq. (23),

$$\eta_c = w_{c,0} \prod_{n=1}^{P} \left(1 - \exp \left[-\frac{2x_{c,n}x_{c,n-1}}{\langle b_x^2 \rangle} \right] \right) \quad (31)$$

where $x_{c,n}$ denotes the x-coordinate of the n^{th} bead on chain configuration c. Using these fixed single-chain slab occupancies and HW-weights, we calculate the statistical weight, w_c, of chain conformation c, the normalized density profile, $\phi(i_x)$, and the single-chain partition function,

$\mathcal{Q}$, according to

$$w_c = \eta_c e^{-\sum_{i_x} \left[v(i_x)\rho_{c,0}(i_x) + v_g(i_x)\rho_{c,1}(i_x) \right]}$$

$$\phi(i_x) = \frac{zN_x}{N_{\mathrm{conf}}} \sum_c w_c \rho_{c,m}(i_x) \quad (32)$$

$$\mathcal{Q} = \frac{1}{N_{\mathrm{conf}}} \sum_c w_c = \frac{1}{zN_x} \sum_{i_x} \phi(i_x)$$

where $z = 1$ denotes the fugacity, which is adjusted such that $\phi = 1$ in the bulk. The gradient-corrected interaction is given by (cf. Eq. (11))

$$v_g(i_x) = v(i_x) + \frac{1}{12P^2} \frac{R_{\mathrm{EE}}^2 N_x^2}{3L_x^2} \left[v(i_x + 1) - v(i_x - 1) \right]^2 \quad (33)$$

Eqs. (27) and (32) are solved self-consistently using a gradient-free Newton technique.

The surface tension γ is computed according to

$$\omega = -\frac{1}{N_x} \sum_{i_x} \left\{ \phi(i_x) + \frac{\kappa N}{2} \left[\phi^2 - 1 \right] \right\} \quad (34)$$

$$\frac{\gamma R_{\mathrm{EE}}^2}{k_B T \sqrt{\bar{\mathcal{N}}}} = \frac{L_x}{2R_{\mathrm{EE}}} \left[\omega - \omega_{\mathrm{bulk}} \right] \quad (35)$$

where ω denotes the grandcanonical free energy per molecule in units of $k_B T$ in the film and $\omega_{\mathrm{bulk}} = -1$ is the corresponding bulk value. $\bar{\mathcal{N}} = (nR_{\mathrm{EE}}^3/V)^2$, with n and V being the number of chains and the volume, respectively, is the invariant degree of polymerization.

We have compared this HW-HOA scheme to the HW-DEH scheme, where we use the HW-weight according to Eq. (31),[18,19] but all monomers contribute uniformly to the slab occupancies and densities, $w_n = w_n^m = P/(P+1)$ and no gradient correction is applied. Lastly we also compute the density profile and the surface tension of the FP-DEH scheme using the free-particle propagator with $w_n = w_n^m = P/(P+1)$, $\eta_c = w_{c,0}$ and $v_g = v$, which corresponds to the standard discretization of the Edwards Hamiltonian (DEH).

In Fig. 4 depicts the density profile, $\phi(x)$, obtained by the three schemes for $P = 32$ and compares the data to the limit $P \to \infty$. For $\kappa N = 50$ the width of the density profile, $\zeta_\infty/R_{\mathrm{EE}} = \frac{2}{\sqrt{12\kappa N}} \approx 0.08$,[33] is comparable to the bond length $\sqrt{\langle b_x^2 \rangle}/R_{\mathrm{EE}} = 1/\sqrt{3P} \approx 0.10$. In the limit $P \to \infty$ and $\kappa N \to \infty$ the ground-state approximation is accurate and the density profile takes the form $\phi_{\mathrm{GS}}(x) = \tanh(x/\zeta)$. Thus, within self-consistent field theory, a polymer melt at a hard wall is similar to the single-chain problem studied in the previous section.

In the self-consistent field calculations, the FP-DEH propagator (without HW-corrections) results in a finite density at the wall.[18,19] The HW-correction, Eq. (31), already provides a significant improvement over the FP-DEH scheme, and the HW-HOA approximation for $P = 32$ provides indeed very close estimate of the large-P behavior.
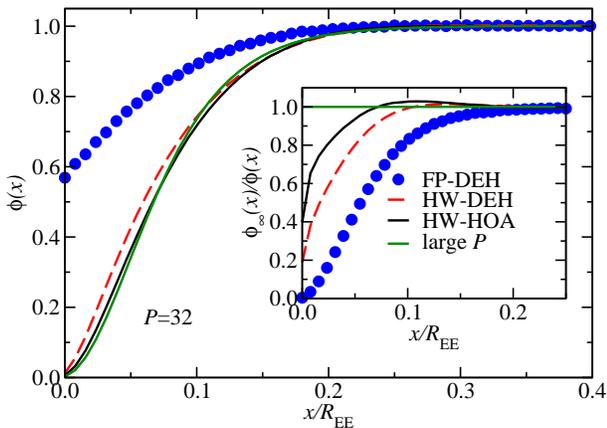
FIG. 4: Comparison of density profiles of a polymer melt with $\kappa N = 50$ at a hard, solid walls for different discretization schemes at fixed $P = 32$ with the large-P limit, $\phi_\infty(x)$. The inset presents the ratio $\phi_\infty(x)/\phi(x)$.
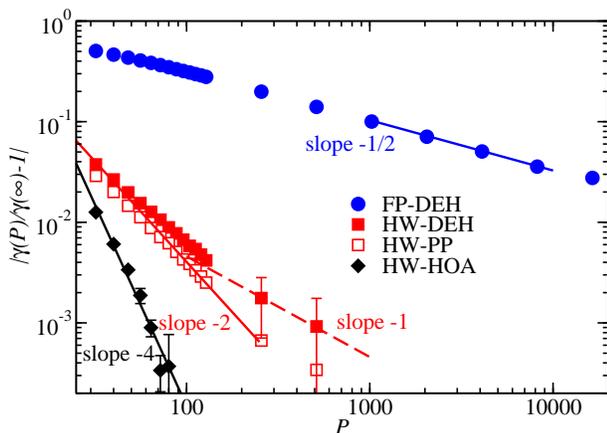

FIG. 5: Relative error of the surface tension of a polymer melt with $\kappa N = 50$ and $\frac{\gamma(\infty)R_{\mathrm{EE}}^2}{k_\mathrm{B}T\sqrt{\mathcal{N}}} = 2.75477$ as a function of the chain discretization P. The FP-DEH, HW-DEH and HW-PP schemes yield $\gamma(P) < \gamma(\infty)$, whereas the HW-HOA scheme overestimates the asymptotic value, $\gamma(\infty)$, of the surface tension. The solid lines indicate the expected power-law behaviour $1/\sqrt{P}$, $1/P^2$ and $1/P^4$ for the FP-DEH, HW-PP, and HW-HOA schemes, respectively. The dashed line indicates $1/P$ scaling.

In the limit of fine discretization and $\kappa N \to \infty$ the surface tension behaves like $\frac{\gamma_{\mathrm{GS}}R_{\mathrm{EE}}^2}{k_\mathrm{B}T\sqrt{\mathcal{N}}} \approx \frac{2}{3}\sqrt{\frac{\kappa N}{3}}\left(1+\frac{0.58}{\kappa N}\right) \approx 2.7532$,[19] and from the extrapolation of the numerical data towards $P \to \infty$ we obtain $\frac{\gamma(\infty)R_{\mathrm{EE}}^2}{k_\mathrm{B}T\sqrt{\mathcal{N}}} = 2.75477$. The relative deviation of the calculated surface tension from $\gamma(\infty)$ are shown in Fig. 5 as a function of chain discretization P. The free-particle (FP) propagator gives rise to large deviations from the Gaussian limit, underestimating $\gamma(\infty)$. As expected, the error decreases like $1/\sqrt{P}$. The HW-DEH, HW-PP, and HW-HOA schemes significantly improve the estimate. For the P-range that

we have investigated, the relative error of the HW-DEH scheme decreases with a power-law but the effective exponent is slightly smaller than the expected value -2. For large P the data are compatible with a $1/P$-behavior within the error bars. Additionally we have used the HW-scheme with a primitive propagator according to the trapezoidal rule (HW-PP), which assigns only half the weight to the end beads, i.e., $w_n = w_n^m = 1$ except for $n = 0$ and P for which $w_n = w_n^m = 1/2$. This HW-PP scheme reduces the relative error compared to the HW-DEH scheme, and the power-law, with which the error decreases, is closer to the expected $1/P^2$ behavior. Combining the HW-correction and the high-order discretization yields the best results, and the numerical data confirm the expected accuracy $O(1/P^4)$ of the HW-HOA scheme. Similar to the behavior of the mean energy of a chain in an harmonic trap next to a wall (cf. Fig. 3) the DEH, HW-DEH, and HW-PP schemes underestimate and the HW-HOA discretization overestimates the value compared to the large-P limit.
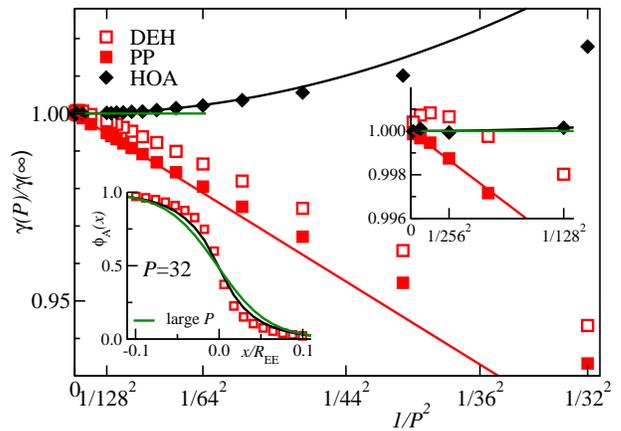
3. Interface tension in a homopolymer blend

FIG. 6: Deviation of the interface tension of the interface tension γ from its asymptotic large-P value, $\frac{\gamma(\infty)R_{\mathrm{EE}}^2}{k_\mathrm{B}T\sqrt{\mathcal{N}}} = 2.68785$, in a symmetric binary polymer blend with $\chi N = 50$ and $\kappa N = 1000$. Symbols present the self-consistent field calculations, whereas lines depict the expected power laws. The right inset highlights the non-monotonic converges of the PP-scheme to the asymptotic result. The left inset compare the density profiles, obtained by DEH and HOA-calculations with $P = 32$, to the results for large P.

The chain discretization is not only important at the narrow surface of a polymer in contact with a solid substrate but sharp gradients of the polymer density also occur at the surface of a polymer melt in contact with air or at interfaces between strongly incompatible homopolymers or in high-χ block copolymer materials. Here we illustrate the advantages of the HOA-scheme by considering the interface in a symmetric binary AB homopolymer blend at rather large incompatibility $\chi N = 50$. Within

the self-consistent field theory or SCMF simulations, A beads at position x experience the interactions

$$\frac{v_A(x)}{P} = \frac{\kappa N}{P}\left[\phi_A(x) + \phi_B(x) - 1\right] + \frac{\chi N}{P}\phi_B(x) \quad (36)$$

and a similar expression holds for v_B acting on B beads. The first term penalizes density fluctuations, and we use $\kappa N = 1000$ to mimic a nearly incompressible blend. The second term quantifies the repulsion between unlike polymers, and in our calculations we employ the value $\chi N = 50$ for the combination of Flory-Huggins parameter χ and number of segments per chain. We consider a symmetric blend, $P_A = P_B = P$, in a system of linear dimension $L_x/2 = 1.5R_{EE}$ with with reflecting boundaries. Space is discretized into $N_x = 128$ slabs. In the limit of fine discretization and large incompatibility the interfacial width is given by $w_\infty/R_{EE} \approx 1/\sqrt{6\chi N} \approx 0.058$.[34]

The calculation of the densities, $\phi_A(x)$ and $\phi_B(x)$, and single-chain partition functions proceed similar to Eq. (32) with $\eta_c = 1$. Three discretization schemes are used: (1) the discretised Edwards Hamiltonian (DEH) with equal weight of all beads, (2) the primitive propagator (PP) where the end beads are assigned half the weight, and (3) the higher-order approximation (HOA).

Once, self-consistency between the densities, Eq. (32) with the external fields, Eq. (36), is achieved, the calculations yield the density profiles across the interface. The grandcanonical free energy per molecule in units of $k_B T$ is computed via

$$\omega = -\frac{1}{N_x}\sum_{i_x}\left\{\phi_A + \phi_B + \frac{\kappa N}{2}\left[\{\phi_A + \phi_B\}^2 - 1\right]\right.$$
$$\left. + \chi N \phi_A(i_x)\phi_B(i_x)\right\} \quad (37)$$

and the interface tension measured by Eq. (35). In the strong segregation limit, the Gaussian chain model predicts $\frac{\gamma(\infty)R_{EE}^2}{k_B T\sqrt{\mathcal{N}}} \approx \sqrt{\frac{\chi N}{6}}\left(1 + \frac{4\ln 2}{\chi N}\right) \approx 2.727$ for an incompressible system.[34]

The inset of Fig. 6 depicts the density profile across the AB interface for $P = 32$ obtained by the DEH-scheme and the HOA-scheme and compare the result to the large-P limit, corresponding to the HOA-result with $P = 1024$. One observes that the discretization effects result in profiles that are too narrow; the deviation is larger for the DEH-scheme than for the HOA-calculation.

The main panel presents the results for the interface tension γ normalized by the large-P limit $\frac{\gamma(\infty)R_{EE}^2}{k_B T\sqrt{\mathcal{N}}} = 2.68785$. The PP calculations underestimate the interface tension, the HOA-calculations yield too large values. The expected scaling of the deviation from the large-P limit is indicated by the solid lines. As expected we observe an error of order $1/P^2$ for the PP-scheme, whereas the HOA-calculations achieve an accuracy of $O(1/P^4)$. The inset presents the convergence of the DEH and PP schemes to $P \to \infty$. The PP scheme nicely converges according to a $1/P^2$ correction; however, the DEH scheme overshoots

the asymptotic value, i.e., $\gamma(P)$ exhibits a maximum at large P. This behaviour indicates that the $1/P^2$-behavior observed at intermediate P values is not the true asymptotics but that there may be an additional, albeit small, $1/P$-corrections that stems from the errors of integrating the interactions along the chain.

B. Efficiency test of staging moves

In order to test the efficiency of staging moves, we computed the time auto-correlation functions (ACF) of the end-to-end radius, which is the slowest coordinate in the system. Specifically, we determined

$$C_{EE}(\Delta t) \equiv \langle R_{EE}(t + \Delta t)R_{EE}(t)\rangle, \quad (38)$$

where t enumerates the MC sweeps through the system and Δt is the number of MC sweeps between two observations. The system consists of a harmonic potential of stiffness $k = 16$, which is centered at the location of a hard wall, while all other parameters of the Hamiltonian and thermal energy are set to unity. In the simulations, we ran staging moves at a fixed level, i.e., by making n trial moves for the n'th fraction of a chain within each single MC sweep. Thus, one sweep costs roughly the same CPU time at a given value of P, irrespective of what fraction of the chain is sampled. The deduced autocorrelation times (ACT) are relatively insensitive to P, which is why we present them in Fig. 7 as a function of the fraction of the moved chain.
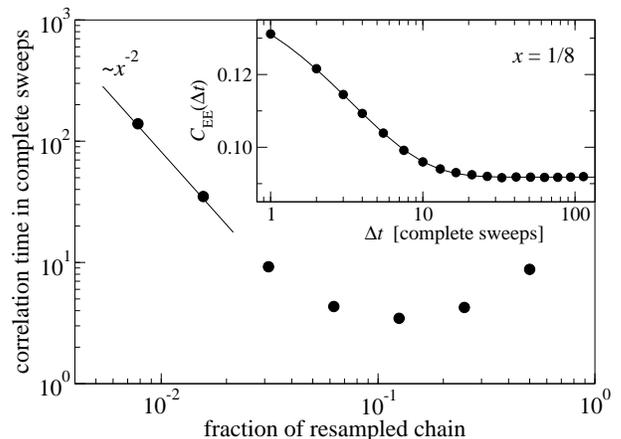

FIG. 7: Inset: End-to-end radius autocorrelation function $C_{EE}(\Delta t)$ if $x = 1/8$ of the chain is resampled with staging. Symbols show simulation results and the full line corresponds to a fit with a stretched exponential. Main graph: Correlation time deduced from $C_{EE}(\Delta t)$ for different fractions x of the resampled chain. At small values of x, the correlation time increases with $1/x^2$ as indicated by the full line.

At very small chain fractions x, the staging moves become very inefficient and autocorrelation times increase with $1/x^2$. Thus, autocorrelation time is largest in the limit of $x = 1/P$ corresponding to single-bead staging

moves. At the other end of very large chain fractions, moves can also become inefficient. The reason is that the trial paths stray too far from the positions associated with a sufficiently small potential energy. The optimum fraction is obtained when the effective spring constant connecting the beginning and the end bead of the sampled segment is close to the curvature of the external potential. In heterogeneous systems with zones of high and small potential curvature, it will therefore be beneficial to employ V-cycles. They ensure efficient sampling everywhere while increasing the cost of one sweep only from $O(P)$ to $O(P \ln P)$.

IV. CONCLUSIONS

In this work, we demonstrated that high-order decompositions of path integrals — introduced for the effective computation of partition functions of quantum-mechanical point particles — likewise benefit Gaussian-chain simulations in the field of polymer physics. Specifically, the rate of convergence can be increased from the usual $1/P^2$ to $1/P^4$, where P is the number of springs in the harmonic chain. For the method to be useful for open chains, averages along the chain have to be performed with specific weights, which differ from those used for closed chains that are common in path integral calculations. These techniques provide a substantial improvement compared to the standard discretization of the Edwards-Hamiltonian (DEH) and allow us to approximate the behavior of Gaussian chain in rapidly varying external fields with a rather crude chain discretization P. We expect that these techniques will be advantageous, for instance, to simulate (i) mixtures of polymers and nanoparticles, which are characterized by a length-scale separation between the narrow interface of the nanoparticle in contact with polymer melt and the length scale of the spatial arrangement of nanoparticles, or (ii) narrow interfaces in high-χ block copolymer materials that enable the fabrication of sub-10nm structures.

In addition, we extended the so-called staging algorithm to situation where a reflecting wall is present so that a given fraction of a chain can be sampled in a Monte Carlo simulation at an essentially constant acceptance ratio, no matter how large the discretization. We expect this method to be useful not only for confined chains subjected to an external potential but also for soft, coarse-grained models or DPD models of polymer films[35]. For external or self-consistent potentials whose derivative disappears on the surface of the wall, the $1/P^4$ convergence remains valid by simply combining hard-wall adopted staging with the higher-order decomposition of the path integral.

Appendix A: Drawing trial coordinates via Monte Carlo from approximate distributions

In this appendix, we discuss in detail how we generate the trial coordinates (for end-point and mid-point beads) in those cases, where we cannot directly relate uniform random numbers to the desired distributions. The idea is that one converges extraordinarily quickly to the exact (free-polymer-in-front-of-a-hard-wall) distribution $\mathrm{Pr}_e(x)$ in a short MC simulation, if one can draw (at small CPU-time cost) from a bias distribution $\mathrm{Pr}_b(x)$ that is a reasonable approximation to the exact distribution.

If x_{old} denotes the present coordinate and x_{trial} a trial coordinate drawn from $\mathrm{Pr}_b(x)$, then the transition probability

$$\mathrm{Pr}_t = \frac{\mathrm{Pr}_e(x_{\mathrm{trial}}) \, \mathrm{Pr}_b(x_{\mathrm{old}})}{\mathrm{Pr}_b(x_{\mathrm{trial}}) \, \mathrm{Pr}_e(x_{\mathrm{old}})} \tag{A1}$$

satisfies detailed balance. Thus, the exact distribution is approached more closely with each Metropolis-algorithm-based MC step if the given transition probability is used. In the following, we demonstrate how quickly it is approached in practice for both end-point and mid-point beads.

We first consider an end-point bead from section II D, for which the starting bead x_1 is neither very far nor very close to the wall. One can assume that for $x_1 > \sigma/\sqrt{2}$, Gaussian random numbers are reasonable approximations, while for $x_1 < \sigma/\sqrt{2}$ random numbers produced with equation (26) are good starting points. According to the ideas just presented, we initialize our random variable by drawing it from the respective bias distribution. We then draw a trial coordinate from the same distribution and compute the acceptance probability according to equation (A1). We then iterate as needed. In this procedure, trial coordinates violating the boundary condition are discarded immediately.

Figure 8 gives an impression of how many MC steps are needed so that the trial coordinates approach the desired distribution function for end-point beads. It shows the "worst-case scenario" for both bias distributions, that is, for $x_1 = \sigma/\sqrt{2}$, which separates the regimes where one or the other bias distribution is taken. One can see that one single time step suffices to produce trial-coordinate distributions that are reasonably close to the exact distribution. Not shown explicitly is that symbols overlap with the exact distribution after one additional move and that another step later, it becomes difficult to ascertain differences between the exact distribution

functions and those produced by Monte Carlo. In practice, we choose $\mathrm{int}(4x_1/\sigma + 0.5)$ steps for $x_1 < \sigma/\sqrt{2}$ and $\mathrm{int}\{4\exp(-x^2/4\sigma^2) + 0.5\}$ steps otherwise, to yield essentially converged distributions.
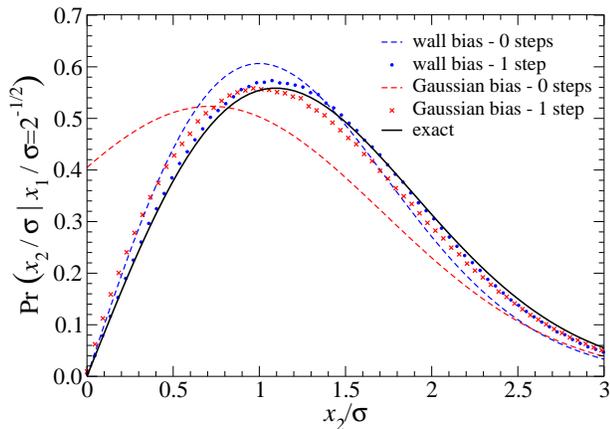

FIG. 8: Bias functions and probability distribution functions used to approximate the exact distribution function assuming that the random walk starts at $x_1 = \sigma/\sqrt{2}$ away from a reflecting wall at $x = 0$, where σ is the variance of the random walk without walls. The blue dotted lines shows the bias function that is based on equation (21) (wall bias) while the red dotted line refers to a normal distribution (Gaussian bias). Blue circles and red crosses show the corresponding distributions after one single MC move. The black line represents the exact distribution function $\Pr(x_2/\sigma|x_1 = \sigma/\sqrt{2})$ for the end monomer coordinate x_2 given that the random walk starts at position $x_1 = \sigma/\sqrt{2}$,

To generate trial coordinates for mid-point beads, we proceed again by using equation (A1). As mentioned in the main text, we first locate (numerically) the point x_c where the exact distribution has its maximum. It satisfies $x_c \geq \sigma$ and $x_c \geq (x_1 + x_3)/2$ and it does not need to be known to very high precision. We then draw Gaussian random numbers with a mean of $x_{\max}$ and the appropriate variance of $\sigma/\sqrt{2}$. To investigate the rate rate of convergence, we consider again the "worst-case scenario", where convergence is slowest, i.e., the case for $x_1 = x_3 \ll \sigma$, see figure 9.

As for the end-bead distribution function, good resemblance to the exact distribution is obtained after a single MC step for the mid-point beads: the boundary condition $\Pr(x_2) = 0$ and $\Pr(x_2) \propto x_2$ (end beads) as well as $\Pr(x_2) \propto x_2^2$ (mid-point beads) are satisfied after one single Monte Carlo step, though the prefactors still have small errors. This time, two MC steps are only almost sufficient to reproduce the exact distribution function to within symbol size of the graphs. In production runs, we set the number of steps to $\mathrm{int}[5\{1.2-\tanh(x_C/4\sigma)\}^2]$ and thereby reach sufficient convergence for all practical purposes, though half the number of steps produce identical results within our stochastic error.

[1] R. P. Feynman, Rev. Mod. Phys. **20**, 367 (1948).
[2] R. P. Feynman and A. R. Hibbs, *Quantum Mechanics and Path Integrals* (McGraw-Hill, New York, 1965).
[3] D. Chandler and P. G. Wolynes, J. Chem. Phys. **74**, 4078 (1981).
[4] H. Kleinert, *Path Integrals in Quantum Mechanics, Statistics, Polymer Physics, and Financial Markets* (World Scientific, Singapore, 2009).
[5] R. P. Feynman, Ann. Phys. **24**, 118 (1963).
[6] A. O. Caldeira and A. J. Leggett, Physica A **121**, 587 (1983).
[7] E. Helfand and Y. Tagami, J. Polym. Sci. B: Polymer Letters **9**, 741 (1971).
[8] E. Helfand, J. Chem. Phys. **62**, 999 (1975).
[9] M. Müller and G. D. Smith, J. Polym. Sci. B: Polymer Physics **43**, 934 (2005).
[10] K. C. Daoulas and M. Müller, J. Chem. Phys. **125**, 184904 (2006).
[11] B. J. Berne, Annu. Rev. Phys. Chem. **37**, 401 (1986).
[12] M. E. Tuckerman, B. J. Berne, G. J. Martyna, and M. L. Klein, J. Chem. Phys. **99**, 2796 (1993).
[13] G. H. Fredrickson, V. Ganesan, and F. Drolet, Macromolecules **35**, 16 (2002).
[14] M. Müller and K. C. Daoulas, Phys. Rev. Lett. **107**, 227801 (2011).
[15] M. H. Müser and B. J. Berne, J. Chem. Phys. **107**, 571 (1997).
[16] P. E. Rouse, J. Chem. Phys. **21**, 1272 (1953).
[17] M. H. Müser, Comp. Phys. Comm. **147**, 83 (2002).
[18] A. Ramirez-Hernandez, F. A. Detcheverry, and J. J. de Pablo, J. Chem. Phys. **133**, 64905 (2010).
[19] M. Müller, B. Steinmüller, K. C. Daoulas, A. Ramirez-Hernandez, and J. J. de Pablo, Phys. Chem. Chem. Phys. **13**, 10491 (2011).
[20] S. A. Chin, Phys. Lett. A **226**, 344 (1997).
[21] S. J. Jang, S. M. Jang, and G. A. Voth, J. Chem. Phys. **115**, 7832 (2001).
[22] A. Ranjan, J. Qin, and D. C. Morse, Macromolecules **41**, 942 (2008).
[23] D. J. Audus, K. T. Delaney, H. D. Ceniceros, and G. H. Fredrickson, Macromolecules **46**, 8383 (2013).
[24] I. Szleifer and M. A. Carignano, Adv. Chem. Phys. **94**, 165 (1996).
[25] M. Müller and A. Werner, J. Chem. Phys. **107**, 10764 (1997).
[26] M. Sprik, M. L. Klein, and D. Chandler, Phys. Rev. B **31**, 4234 (1985).
[27] D. Marx and M. H. Müser, J. Phys.-Condens. Matt. **11**, R117 (1999).
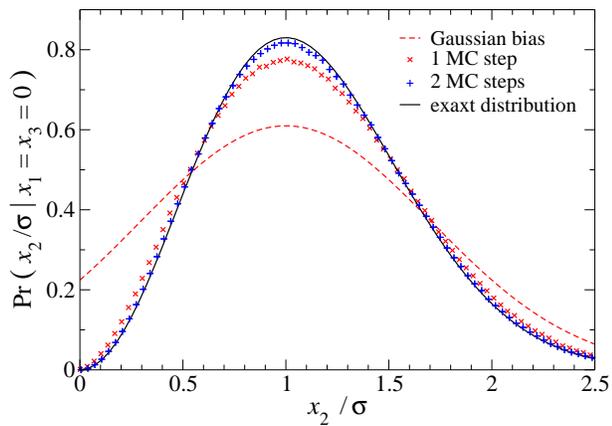[28] H. F. Trotter, Proc. Am. Math. Soc. **10**, 545 (1959).

FIG. 9: Probability distribution for the mid-point bead to be placed at a position x_2/σ. The read dotted lines indicates the bias probability function, which is a truncated Gaussian having its maximum placed at the same value of x_2 as the fully converged distribution function $\Pr(x_2/\sigma|x_1 = x_3 = 0)$. Red crosses and blue plus signs show the distribution obtained after, respectively, one or two MC steps, which are based on equation (A1).

[29] M. Suzuki, Commun. Math. Phys. **51**, 183 (1976).
[30] R. M. Fye, Phys. Rev. B **33**, 6271 (1986).
[31] M. Takahashi and M. Imada, J. Phys. Soc. Jpn. **53**, 3765 (1984).
[32] F. R. Krajewski and M. H. Müser, Phys. Phys. B **65**, 174304 (2002).
[33] D. T. Wu, G. H. Fredrickson, J. P. Carton, A. Adjdari, and L. Leibler, J. Polym. Sci.: Polym. Phys. **33**, 2373 (1995).
[34] A. N. Semenov, J. Phys. (France) II **6**, 1759 (1996).
[35] M. Müller, J. Stat.Phys. **145**, 967 (2011).
[36] The alternative to choose s_{max} on the fly as a function of the (local) curvature of the external potential is not advised as this can violate detailed balance.