

Neighbor List Artifacts in Molecular Dynamics Simulations

Hyuntae Kim,^{†,‡} Balázs Fábián,[†] and Gerhard Hummer^{*,†,¶}

[†]*Department of Theoretical Biophysics, Max Planck Institute of Biophysics, Max-von-Laue Straße 3, 60438, Frankfurt am Main, Germany*

[‡]*International Max Planck Research School on Cellular Biophysics, Max-von-Laue Straße 3, 60438, Frankfurt am Main, Germany*

[¶]*Institute of Biophysics, Goethe University Frankfurt, Frankfurt am Main, Germany*

E-mail: gerhard.hummer@biophys.mpg.de

Abstract

Molecular dynamics (MD) simulations are widely used in biophysical research. To aid non-expert users, most simulation packages provide default values for key input parameters. In MD simulations using the GROMACS package with default parameters, we found large membranes to deform under the action of a semi-isotropically coupled barostat. As the primary cause, we identified overly short outer cutoffs and infrequent neighbor list updates that result in missed long-range attractive Lennard-Jones interactions. Small but systematic imbalances in the apparent pressure tensor then induce unphysical asymmetric box deformations that crumple the membrane. We also observed rapid oscillations in averages of the instantaneous pressure tensor components and traced these to the use of a dual pair list with dynamic pruning. We confirmed that similar effects are present in MD simulations of neat water in atomistic and coarse-grained representations. Whereas the slight pressure imbalances likely have minimal impact in most current MD simulations, we expect their impact to grow in studies of

ever-larger systems, in particular in combination with anisotropic pressure coupling. We present measures to diagnose problems with missed interactions and guidelines for practitioners to avoid them, including estimates for appropriate values for the outer cutoff r_l and the number of time steps `nstlist` between neighbor list updates.

1 Introduction

Molecular dynamics (MD) simulations are a powerful tool to probe molecular processes at a level of detail not currently accessible to experiments.¹ The GROMACS molecular dynamics simulations package² is widely used, in particular for applications in biophysics, chemistry, and soft-matter science. It is computationally efficient² and easy to use with a wide range of atomistic and coarse-grained force fields quantifying the energetics of molecular interactions.^{3,4} Central to its high performance are the nearly linear scaling of the computational cost with system size and its efficient parallelization over multiple computational nodes.^{5,6} A key factor for the computational efficiency is the use of neighbor lists containing the pairs of interacting particles. To avoid costly neighbor list updates at every time step, the Verlet scheme includes a buffer of particles between the actual cutoff distance for pair interactions, r_c , and an outer cutoff $r_l > r_c$. The neighbor list is updated at longer time intervals chosen so that crossing from distances $r > r_l$ to $r < r_c$ by ballistic motion is highly unlikely. For the construction of neighbor lists on single-instruction multiple-data (SIMD) hardware architectures, GROMACS implements the MxN algorithm,^{7,8} which minimizes inter-node communication and memory footprint.^{2,9} The grouping of particles into spatial clusters by the MxN algorithm enables an efficient evaluation of the real-space pair interactions.

Here, we show that the use of default simulation parameters^{10–12} can cause artificial pressure oscillations and a violation of spatial isotropy. As a consequence, large membrane systems can undergo drastic deformations in the form of unrealistic buckling (Figure 1). We analyze the temporal evolution of lipid bilayers in the NPT ensemble with constant particle number N , pressure P , and temperature T ; and in neat solvents in both NPT and NVT

ensembles, the latter fixing the volume V instead of the pressure P . As the primary cause, we identify the infrequent construction of the neighbor list, as a result of a somewhat too large update interval of `nstlist` time steps and a somewhat too short outer cutoff distance r_l . Consequently, long-range cohesive non-bonded interactions are occasionally missed in the force evaluation. The missed interactions cause errors in the elements of the instantaneous pressure tensor. In the NPT ensemble, these errors in the pressure lead to incorrect box rescaling by the barostat, both with the weak-coupling (Berendsen) barostat¹³ and the Parrinello-Rahman (PR) barostat.¹⁴ We conclude by providing tools that practitioners can use to detect such problems and guidance on minimizing their impact or avoiding them altogether.

The problems identified here likely afflicted also earlier simulation studies. Pointedly, several studies of large membrane systems prevented excessive membrane undulations by restraining the vertical movement of certain lipid head groups with harmonic^{15–17} or flat-bottom^{18–20} potentials. One can also restrain the box with a weak harmonic potential for example by using the plumed software package²¹ as a GROMACS plug-in. However, the introduction of such external potentials is unsatisfactory, motivating our efforts to identify and correct the underlying issues.

2 Methods

2.1 Neighbor List and Missed Interactions

In MD simulations, non-bonded pairwise interactions are usually truncated beyond a distance cutoff r_c .²² The remaining long-range interactions are usually estimated analytically²³ assuming a uniform density of particles beyond r_c , but can be evaluated in Fourier space for power-law potentials in a periodic system using the Ewald method as implemented, e.g., in the particle mesh Ewald (PME) algorithm.^{24,25} Without truncation of the real-space interactions, the computational cost of evaluating pair-wise forces would scale with the square

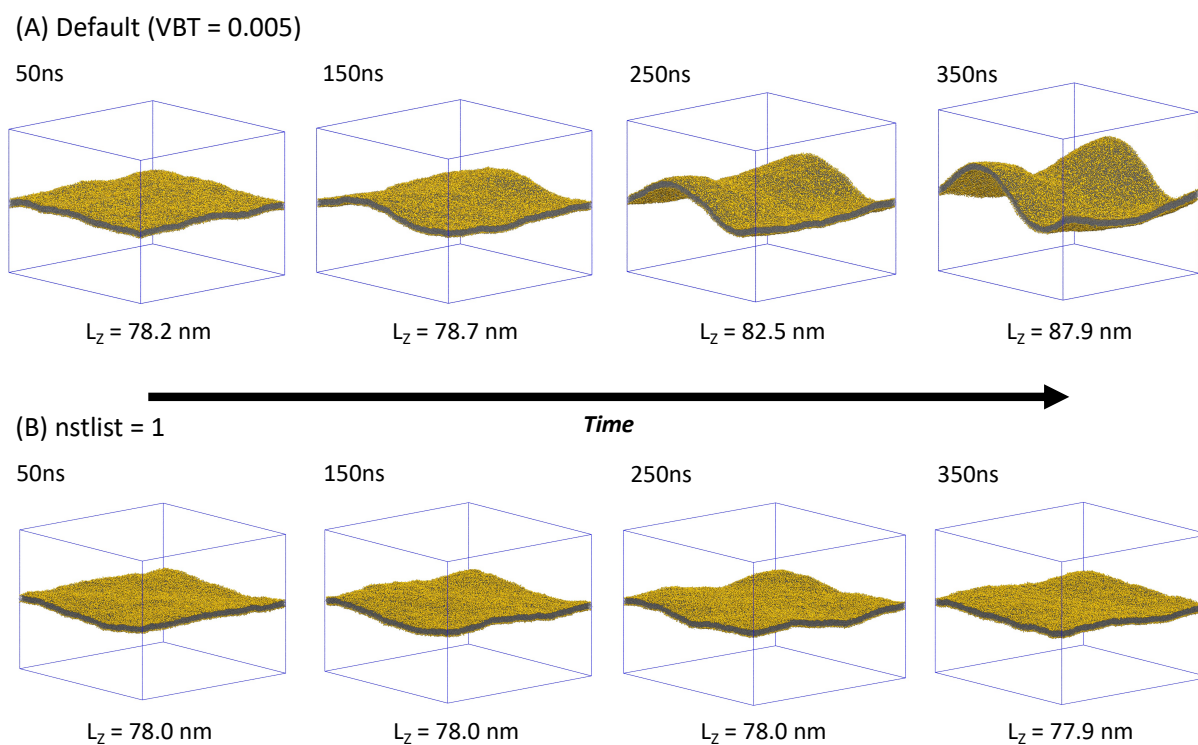


Figure 1: The large Martini POPC bilayer crumples in MD simulations with default simulation parameters, yet stays flat with frequent neighbor list updates. (A) Snapshots of the membrane (phosphate groups in gold) in MD simulation with default parameters ($VBT = 0.005$ $\text{kJ}\cdot\text{mol}^{-1}\cdot\text{ps}^{-1}$; $r_l = 1.269$ nm, `nstlist` = 25 and dual pair list enabled). (B) Snapshots in an MD simulation with neighbor list updates enforced at every time step (`nstlist` = 1). Snapshots are at time points 50, 150, 250, and 350 ns (left to right). Simulation boxes are indicated as blue lines and box heights L_z are listed.

of the particle number N in the system. With fixed cutoff r_c , the cost scales only linearly with N , at least in principle. For the PME algorithm, evaluating the remaining long-range contributions results in $N \log N$ scaling.

The neighbor list of a particle contains the indices of its neighboring particles for which the pairwise interactions are explicitly evaluated in real space. In the Verlet scheme, the neighbor list is constructed by searching for neighbors within a cutoff radius r_l , with $r_l \geq r_c$. The spherical shell between r_l and r_c provides a buffer so that neighbor list updates are not required at every time step. Neighbor search requires an evaluation of the pairwise distances, hence its computational cost scales at least linearly with the system size.²⁶ Furthermore, the neighbor search requires inter-node communications, which can be a major bottleneck for modern hardware architectures.^{10,27} If neighbor list updates are performed only every `nstlist` time steps of length $\Delta t = \text{dt}$, we expect that some pair interactions are missed because particle pairs move from distances $r > r_l$ to $r < r_c$ within the time interval `nstlist` $\times\Delta t$. For particles of mass m moving with Maxwell-Boltzmann velocities in a system of uniform number density ρ , we can estimate (see SI text) the probability that a particular particle misses an interaction as

$$p_{\text{missed}} \approx \frac{\sqrt{2\pi}\rho\sigma^3(r_c^2 + r_l r_c + r_l^2)e^{-\frac{(r_l-r_c)^2}{2\sigma^2}}}{3(r_l - r_c)^2}, \quad (1)$$

where we assumed that $r_l - r_c \gg \sigma = \Delta t\sqrt{2k_B T/m}$ with k_B Boltzmann's constant.

2.2 GROMACS Input Parameters

In this study, we critically examine the following GROMACS input parameters: `nstlist`, `nstenergy`, `nstcalcenergy`, `nstpcouple`, `nsttcouple`, `verlet-buffer-tolerance`, and `rlist`, also denoted as r_l . The parameters `nstlist`, `nstenergy`, `nstcalcenergy`, `nstpcouple`, and `nsttcouple` denote the number of time steps between neighbor list updates, energy sampling, energy evaluation, barostatting, and thermostatting, respectively. The default values

recommended by the developers can be found in the manual:¹² `nstlist` = 10, `nstenergy` = 1000 and `nstcalcenergy` = 100.

The `Verlet-buffer-tolerance` (VBT) denotes the maximally allowed energy drift per particle between neighbor list updates due to missed non-bonded interactions. Its default value is $0.005 \text{ kJ}\cdot\text{mol}^{-1}\cdot\text{ps}^{-1}$.^{2,12} Changes in VBT result in adjustments of r_l and `nstlist`. In standard GROMACS runs, the values of r_l and `nstlist` are therefore not only system-dependent, but there is also no guarantee that the values are constant throughout a trajectory. In particular, the possible values of `nstlist` are 20, 25, 40, 50, 60, 80, and 100. The values of the adjusted r_l and `nstlist` can be found in the output `log` file. To ensure a constant value of `nstlist`, whether it is user-defined or the default value of ten, VBT must be disabled by setting `VBT` = -1.

In GROMACS, the allowed energy drift VBT determines the frequency of neighbor list updates. By contrast, in LAMMPS²⁸ the neighbor lists is updated when any particle travels more than half the buffer thickness. As the system size increases, the time interval between updates shrinks to the point of forcing an update at every time step.²⁷

2.3 The MxN Algorithm and the Dual Pair List

In SIMD hardware architectures, GROMACS employs the MxN algorithm for a grid-based neighbor search.^{7,8} The algorithm clusters a fixed number of particles by gridding the xy plane and binning along the z axis. The clusters with insufficient numbers of particles are filled with dummy particles. The implementation of the algorithm promises a high computational performance. Moreover, the clusters act as another layer of buffer on top of the pre-defined $r_l - r_c$ shell, enabling a further increase in `nstlist`.⁸

The performance can be further improved by implementing a dual pair list algorithm,²⁷ using a long outer and a short inner list cutoff. The inner neighbor lists are generated from a pool of particles within the outer list and hence updated more frequently. The implementation of the dual pair list algorithm reduces the overall computational cost of the

neighbor search. The update frequencies and the cut-off radii for both the outer and inner list are by default controlled by `VBT`, and their exact values can be found in the `log` file. The dual pair list algorithm can be disabled by setting `VBT` to `-1`. When the dual pair list algorithm is enabled, r_l becomes the cutoff radius for the outer neighbor list. For GPUs, dynamic pruning is used to take advantage of their typically large execution width during the neighbor search.²⁷

2.4 Membrane Bending Free Energy

The bending energy E associated with elastic deformations of a fluid and incompressible membrane can be estimated as an integral of the squared local mean curvature H over the membrane surface A ²⁹

$$E_{\text{bend}} = 2\kappa \int dA H^2 \quad (2)$$

where κ is the bending rigidity of the membrane. Here we ignored the contribution of the Gaussian curvature which is invariant for a given topology. We evaluated the local mean curvature H of the membrane systems using the `MemCurv` program³⁰ and then integrated it numerically over the xy plane of the box, thus ignoring curvature corrections to the area element dA .

2.5 Simulation Code

Two versions of GROMACS were examined, namely 2020.3 and 2023. All the numerical analyses were performed using GROMACS 2020.3, which is the version for which the artifacts were initially observed. However, all the system types described in this section were also simulated using GROMACS 2023, the latest version available. All the major artifacts caused by the use of inadequate combinations of r_l and `nstlist` were also present in GROMACS 2023 runs with default parameters. These include the unphysical distortion of the large membrane systems, oscillations in the average instantaneous pressure, and the violation of

spatial isotropy.

2.6 Simulation of Large and Small Martini Membranes

A large membrane system, consisting of 33,282 1-palmitoyl-2-oleoyl-glycero-3-phosphocholine (POPC) lipids and 6,721,594 water molecules, was built using the *insane.py*³¹ script and the Martini force field (version 2.2).⁴ NaCl salt was added at a concentration of 0.15 M and 10% of the water molecules were replaced with the anti-freeze beads (WF).⁴ The initial dimensions of the system were 100 nm × 100 nm × 80 nm. The system was equilibrated first in the NVT ensemble for 150 ns and then in the NPT ensemble with semi-isotropic pressure coupling for another 150 ns, using $r_t = 1.422$ nm and `nstlist` = 20. Production runs of 1 μs length were then performed using the *new-rf*¹¹ simulation parameters with $r_c = 1.1$ nm and a 20 fs time step. The system was coupled to a v-rescale thermostat³² at 310 K, and semi-isotropically coupled to a PR barostat with a target pressure of 1 bar ($\tau_P = 12$ ps). Also, note that `nstcalenergy` = 1 was used for all the simulations in this manuscripts unless specified otherwise.

Similarly, a smaller Martini membrane system was built, consisting of 722 POPC lipids and 10,732 water molecules. The corresponding initial box dimensions were 15 nm × 15 nm × 10 nm. All preparation procedures and input parameters were identical to those of the large membrane system, and it also underwent 1 μs production run.

2.7 Simulation of Water Systems

A system of neat Martini solvent was prepared for MD simulations in both the NVT and NPT ensembles during the production run. Its initial dimensions were 6 nm × 6 nm × 6 nm, consisting of 1,530 Martini water particles. Similarly to the Martini membrane systems, the *new-rf*¹¹ simulation parameters with 20 fs time step were used. Also, the ratio between water particles and antifreeze particles WF was set to 9:1. The system was equilibrated for 200 ns in both the NVT and NPT ensembles using $r_t = 1.422$ nm. Detailed analyses were

performed on a few μs long production runs. It is important to note that there are no dipole moments or charges in the Martini solvents, hence it is effectively a pure LJ solvent.

The distortion of the cubic box with NPT Martini solvent was studied to examine the violation of spatial isotropy due to the use of an inadequate combination of r_l and `nstlist`. A few varying conditions were examined, where the system was coupled to four different barostat types (semi-isotropically coupled PR, Berendsen and C-rescaling,³³ and anisotropically coupled PR) with consistent target pressures of 1 bar. Before the production runs, the equilibrated systems were isotropically scaled by factors of 0.99, 1.00, and 1.01, respectively. This scaling was intended to mimic possible volume artifacts caused by the use of inadequate combinations of r_l and `nstlist`. To account for possible anisotropy in the initial condition, the equilibrated systems were also rotated about the x , y , and z axes, respectively. This procedure was intended to eliminate any bias caused by the initial configuration of the molecules. For the same reason, the initial velocities of the molecules were randomly generated, according to a Maxwell-Boltzmann distribution, for each of the 500 replicates, resulting in a sample size of 4,500 runs for each barostat type. Then, the above procedures were performed using three different combinations: $r_l = 1.9$ nm (`nstlist` = 1), $r_l = 1.9$ nm (`nstlist` = 20), and $r_l = 1.28$ nm (`nstlist` = 25). Hence, 54,000 solvent simulations entered the statistical analysis of possible anisotropy. Finally, hypothesising the asymmetric implementation of the MxN algorithm to be the cause of the potential anisotropy, we disabled the MxN algorithm by recompiling the GROMACS MD engine with `-DGMX_SIMD=None`. We then performed 4500 replicate simulations of the fully anisotropically coupled system with $r_l = 1.28$ nm, `nstlist` = 25, and without MxN algorithm.

Furthermore, the Martini water system was simulated in the NVT ensemble to illustrate that the various observed artifacts (except box rescaling) are not due to the barostat and occur independent of the ensemble type. As for the membrane systems, the respective production runs were 1 μs long. Except for the barostat settings, all input parameters were identical to those used for the NPT Martini water solvent.

Similarly, a system of pure TIP3P water³⁴ was generated via CHARMM-GUI in cubic NVT ensemble,³⁵ with the dimensions of 5 nm \times 5 nm \times 5 nm. The systems was equilibrated for 15 ns in both the NVT and NPT ensembles using $r_l = 1.422$ nm, while the production runs were 200 ns long with 2 fs time step. The non-bonded interaction cutoff, $r_c = 1.2$ nm, was used. The temperature was fixed at 310 K with the Nosé-Hoover thermostat.^{36,37} For the calculation of the power spectral density of the pressure in TIP3P water, we used the v-rescale thermostat to suppress the oscillatory contributions of the Nosé-Hoover thermostat.³² Finally, we used the PME algorithm with default `fourierspacing` (0.12 nm).

2.8 Power Spectral Analysis

The power spectral density (PSD) of the scalar pressure was calculated using Welch's method.³⁸ As in input, we used the time series of the pressure calculated and saved at every time step (`nstenergy = 1`). The resulting PSD was plotted as function of the frequency in units of $1/(\text{nstlist} \times \Delta t)$. The visual inspection focused on peaks in the PSD as a means to identify characteristic time intervals of processes resulting in perturbations of the barostat action.

3 Results

3.1 Unphysical Distortion of Large Martini Membrane

The temporal evolution of the large Martini membrane system simulated with default parameters and PR barostat is shown in Figure 1A. Within the 350 ns of MD, the simulation box contracted in the xy membrane plane and expanded in the z direction. The box dimensions changed from 104.0 nm \times 104.0 nm \times 78.1 nm after equilibration to 97.1 nm \times 97.1 nm \times 89.7 nm at the end of the production run. This change in box shape left the overall volume approximately constant.

During the simulations, the large membrane buckled to form distinct folds in the x and

y directions (Figure 1A). The potential energy and the enthalpy of the system increased by 13,200 and 24,600 $\text{kJ}\cdot\text{mol}^{-1}$, respectively, in 350 ns of MD. These increases are substantial also in relative terms, amounting to changes of $\approx 1.58\%$ and $\approx 1.85\%$, respectively. The membrane bending energy at the end of the MD simulation was estimated using eq 2 as $E_{\text{bend}} \approx 138 k_B T \approx 357 \text{ kJ}\cdot\text{mol}^{-1}$ for a bending modulus of $25 k_B T$ for a Martini POPC bilayer.³⁹ As a test, we also calculated the bending modulus for our setup, obtaining a fully consistent value of $\kappa = 24.7 k_B T$ (see SI Text and Figure S1). These large increases in system energy and enthalpy, and in the membrane bending energy, strongly indicate that the observed deformation is unphysical.

Pressure imbalances, not the barostats per se, appear to drive the box deformations. In MD simulations with semi-isotropically coupled PR and Berendsen barostats, we observed similar box deformations (Figures 1 and S2). Having thus ruled out an effect due to a specific barostat, we examined the components of the pressure tensor driving the barostat action. The running averages of the diagonal pressure tensor elements for the default setup ($\text{VBT} = 0.005 \text{ kJ}\cdot\text{mol}^{-1}\cdot\text{ps}^{-1}$) are plotted as dashed lines in Figure 2A. Results are shown for the early phase of the simulation when the membrane is still flat (see Figure 1A). We observed that the three diagonal pressure components deviate from the target pressure of 1 bar and from each other.

3.2 Cutoff Handling is Responsible for Membrane and Box Deformations

Differences in the apparent pressure averages as function of the frequency of averaging point to the underlying cause. The parameter `nstenergy` is the number of time steps between time points entering the pressure (and energy) averages. It should thus have no effect on the value of the average. However, for `nstenergy` = 1, we overestimated the pressure, and for `nstenergy` = 100 we underestimated it. These differences were significant and reproducible. Moreover, MD simulations of the small membrane system produced similar

results (Figure 2B), pointing to the fact that we are dealing with a generic issue. As a possible explanation, we hypothesized that the pressure values calculated between neighbor list updates, which occur at intervals of `nstlist` = 20 or 25, differ from those right after neighbor list updates, the former dominating the average for `nstenergy` = 1 and the latter for `nstenergy` = 100.

To test this hypothesis, we first performed MD simulations with `nstlist` = 1 to enforce neighbor list updates after every time step. As shown in Figure 1B, this eliminated the box and membrane deformations and made it possible to simulate a stable large membrane system without additional restraints. As a second test, we performed MD simulations in which we set a large outer cutoff distance of $r_l = 1.422$ nm together with neighbor list updates every `nstlist` = 20 time steps. As shown by the solid lines in Figure 2A, the pressure components then converged consistently to the target pressure of 1 bar for `nstenergy` = 1 and 100, respectively. Further support came from MD simulations of the small membrane system, where we also observed convergence to the target pressures (Figure 2B). In addition, MD simulations of a box of Martini water in the NVT ensemble, i.e., without membrane and barostat, showed in essence the same effect of apparent deviations between the pressure averages with default cutoff settings and `nstenergy` = 1 and 100, and consistent averages with $r_l = 1.422$ nm and `nstlist` = 20 (Figure 2C). Infrequent neighbor list updates thus emerged as a likely cause of pressure imbalances and associated box deformations.

3.3 Instantaneous Pressure Oscillates Between Neighbor List Updates

As a further test of the hypothesis that unresolved cutoff violations are at the heart of the observed problems, we examined the pressure as a function of time between neighbor list updates (Figure 3). For the large and small membrane system in NPT ensembles, and for the Martini water system in an NVT ensemble (top to bottom), we saved the time series of the pressure tensor components and averaged them over blocks of length `nstlist` starting

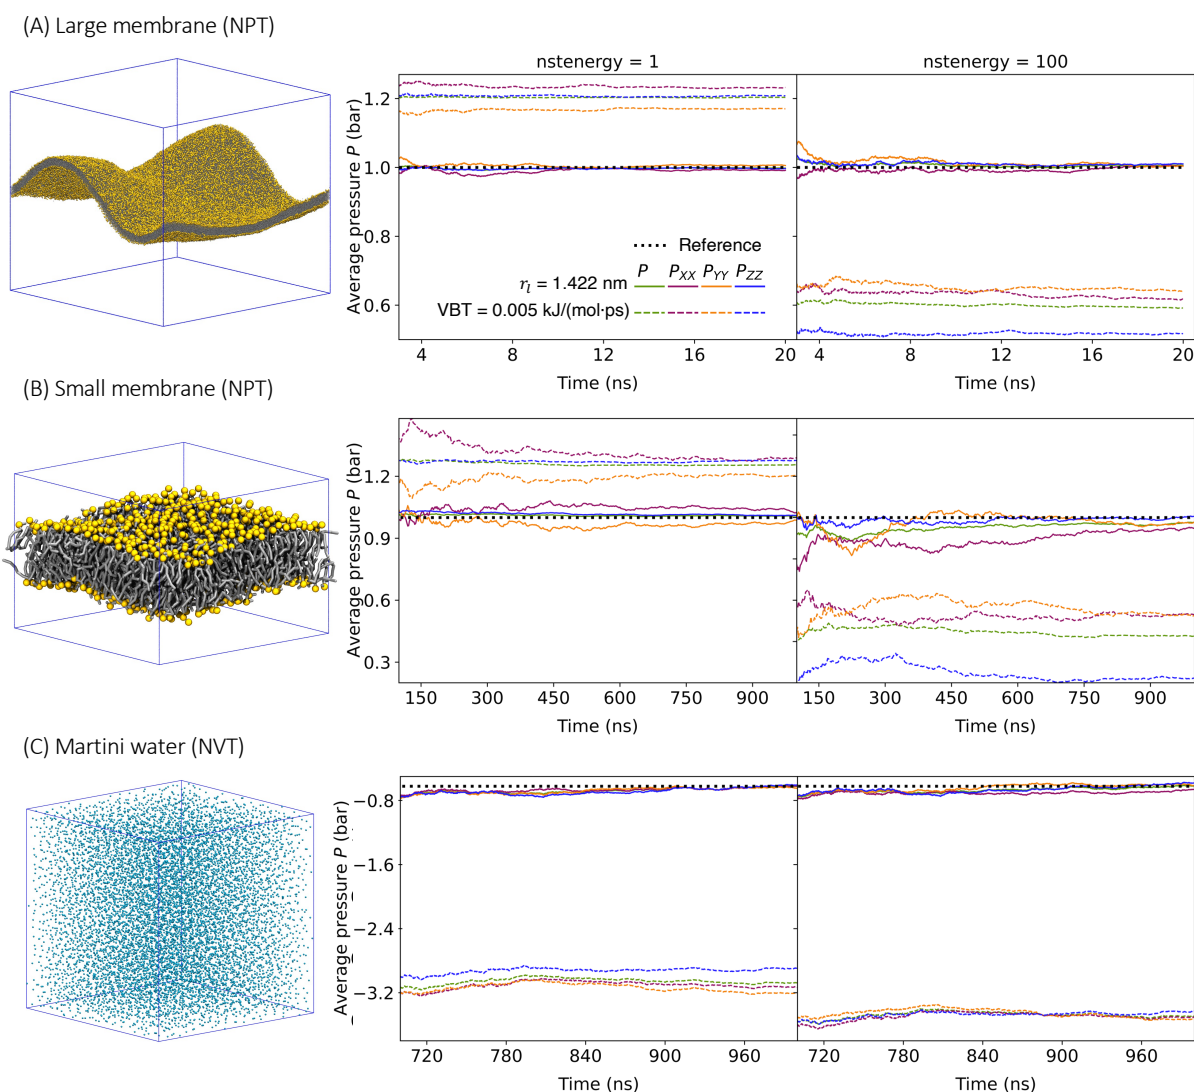


Figure 2: The pressure tensor elements deviate from target pressure in MD simulations with default cutoff handling. Running averages of the diagonal elements of the pressure tensor are shown for (A) the large and (B) the small Martini membrane systems in NPT MD simulations, and (C) for the Martini water system in an NVT simulation. The left column shows snapshots of the systems. The center and right column show the running averages evaluated every $nstenergy = 1$ and 100 steps, respectively. Results obtained with the default simulation parameters for cutoff handling ($VBT = 0.005 \text{ kJ}\cdot\text{mol}^{-1}\cdot\text{ps}^{-1}$) are shown as dashed lines (see legend for color). The solid lines show results for a larger outer cutoff $r_l = 1.422$ nm with $nstlist = 20$ fixed and dual pair list disabled. In (A,B), the target pressure of 1 bar in the NPT simulations is indicated by a dashed black line. For the NVT simulation in (C), the dashed black line indicates the consistent average obtained with $r_l = 1.422$ nm and $nstlist = 20$.




Figure 3: The average pressure tensor components deviate from target pressure between neighbor list updates. Averages were performed over blocks of `nstlist` time steps starting immediately after a neighbor list update (time step 0). Results are shown for the large (top) and small membrane system (center) in NPT simulations, and for the NVT Martini water system (bottom). Results include the default cutoff setting with `nstlist` = 25 (three left columns) and the setting with $r_l = 1.422$ nm and `nstlist` = 20 (right column). The default r_l for the NPT large and small membranes and the NVT solvent were 1.269, 1.267 and 1.28 nm, respectively. For the runs in column 3, the dual pair list was disabled. We averaged the lateral pressure components $P_{||} = (P_{xx} + P_{yy})/2$ (orange curves) and show the normal pressure as $P_{\perp} = P_{zz}$ (blue curves). Dashed horizontal lines of matching color denote the corresponding overall averages. Black horizontal dotted lines represent the reference pressure values. Cyan circles (left column) indicate deviations of the pressure averages from the target. Green circles (column 3) indicate deviations building up just before the neighbor list update.

immediately after a neighbor list update (time step 0). Results are shown for the default cutoff setting with `nstlist` = 25 (three left columns) and for the setting with $r_l = 1.422$ nm and `nstlist` = 20 (right column). We averaged the in-plane pressure components $P_{\parallel} = (P_{xx} + P_{yy})/2$ and show the normal pressure as $P_{\perp} = P_{zz}$.

Consistent with our expectations, we found that with default cutoff settings the average pressures immediately after a neighbor list update (i.e., at time step zero in Figure 3) are systematically lower than the pressures calculated between neighbor list updates (time steps > 0). However, we were initially puzzled by the observation that even at time step zero, the average pressure deviated from the target pressure (even though this is consistent with the findings in Figure 2 for `nstenergy` = 100). As a possible explanation, we considered that neighbor list updates came after `nstlist` = 25 time steps, yet thermostat and barostat actions after `nsttcouple` = `nstpcouple` = 20 time steps, and thus asynchronously (circles in left column of Figure 3). Indeed, by setting `nsttcouple` = `nstpcouple` = `nstlist` = 25 time steps, the target and average pressures at time step 0 are consistent.

Counter to our expectations, we found the instantaneous pressure values to rise rapidly with time and to exhibit distinct oscillations. For missed interactions because of inadequate r_l , we had expected a delayed and monotonic rise with time. As a source of this unexpected behavior, we identified the use of a dual pair list. When we disabled the dual pair list evaluation by setting `VBT` = -1, both the rapid initial rise and the oscillations in the average pressures disappeared (column 3 in Figure 3). With dual pair list enabled, an outer and inner neighbor list are maintained with two distinct intervals of 25 and 4 integration time steps, respectively. The zigzag oscillations in Figure 3 appear to be a superposition of two curves with these two periods. This argument is supported by simulations of Martini water in an NVT ensemble using a different hardware architecture, where the outer and the inner lists were updated every 25 and 5 time steps, respectively (Figure S3). With the update intervals of inner and outer list being multiples of five, the dominant oscillation period is five.

3.4 Pressure Deviations Correlate with Missed Particle Interactions

In Figure 3, the average pressure values right after neighbor list updates are close to the target values. However, at the time step just before the neighbor list update (at time step 24 in Figure 3, column 3), we noticed significant positive deviations ΔP from the pressure at time step 0, as indicated by green circles. In Figure 4, we plot these deviations for lateral (ΔP_{\parallel}) and normal pressures (ΔP_{\perp}) as a function of the outer cutoff r_l with the dual pair list disabled. Results are shown for Martini and TIP3P water. The absolute magnitude of the pressure deviations are plotted for the latter because the electrostatic force contribution is not guaranteed to be systematically cohesive near r_l . For reference, we also show n_{missed} , the expected mean number of unique cutoff violations calculated by multiplying the probability p_{missed} in eq 1 with $N/2$, where N is the number of water molecules. Except for the shortest $r_l \approx r_c$, we find that the pressure errors just before neighbor list updates follows the trend of missed interactions. We note further that the ΔP errors are positive for the Martini water, consistent with missed attractive Lennard-Jones interactions, as these lead to an underestimation of the cohesiveness. We thus conclude that missed interactions due to overly short outer cutoffs are the primary contributor to deviations ΔP in the pressure.

3.5 Anisotropic Errors in Pressure Tensor Deform Box Shape

The errors ΔP in the pressure shown in Figure 4 tend to be somewhat anisotropic even though the boxes had fixed cubic shape and volume. The lateral errors tend to be somewhat smaller than the normal errors, $\Delta P_{\parallel} < \Delta P_{\perp}$. Therefore, we hypothesized that in MD simulations with both semi-isotropic and anisotropic barostats, we should see a tendency for the boxes to grow in the z direction. We tested this hypothesis by running repeated MD simulations of Martini water systems with semi-isotropic and fully anisotropic barostats of PR and Berendsen type.

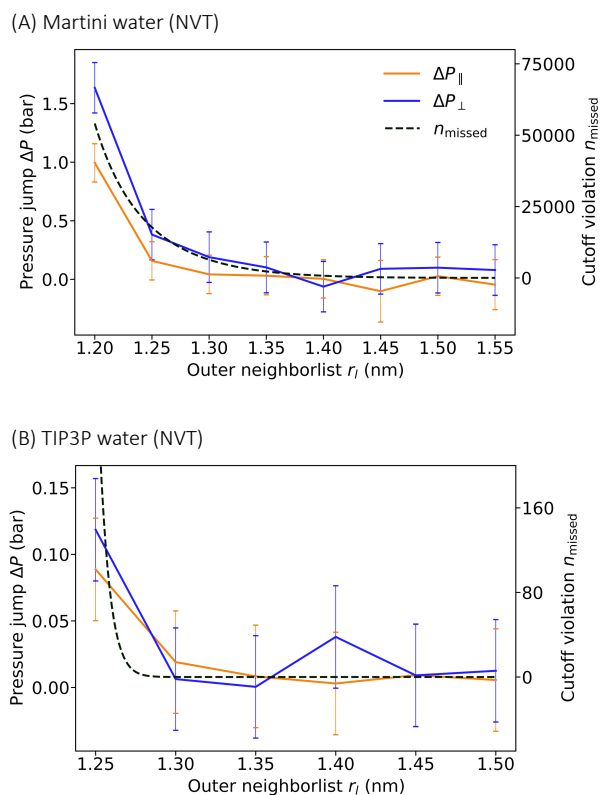


Figure 4: Deviations in the pressure ΔP between the time points just before and right after neighbor list updates follow the trend of missed interactions. ΔP (left scale) for lateral (orange) and normal pressures (blue) is shown as function of r_l (A) for Martini water and (B) for TIP3P water, both simulated in NVT conditions with `nstlist` = 20 and dual pair list disabled. Also shown (right scale) is the expected number of missed interactions (black dashed lines), evaluated using eq 1. For TIP3P water, we used a mass of 18 g/mol, hence ignoring the rotational motion.

The results of these simulations confirm a tendency for the simulation box to expand preferentially along z by the action of the barostat (Table 1). For fully anisotropic pressure coupling with the PR barostat, the two relatively poorer cutoff settings ($r_l = 1.9$ nm, `nstlist` = 20 and $r_l = 1.28$ nm, `nstlist` = 25) resulted in a significant tendency for the box to expand along z (p-value < 0.01 calculated from Pearson's chi-squared test). The observed tendency of the anisotropically coupled box to expand along z (Table 1) is consistent with the consistently larger error in the pressure along z seen in Figure 4. We note, however, that the direction of box expansion is biased towards z but not deterministic. The spatial isotropy was restored when $r_l = 1.9$ nm and `nstlist` = 1 was used (p-value ≈ 0.89). For this setting, all interactions should be counted at all time steps, as also the dual neighbor list is turned off.

Interestingly, we observed a statistically significant asymmetry in the box shape changes even for $r_l = 1.9$ nm and `nstlist` = 1 for MD simulations with a semi-isotropically coupled PR barostat (Table 1). Therefore, we checked alternative barostats with the same setting. For simulations with semi-isotropically coupled Berendsen and C-rescale barostats, the setting $r_l = 1.9$ nm and `nstlist` = 1 did not result in any significant bias in our tests (Table 1). Therefore, we suspect that the asymmetry in box shape changes with semi-isotropically coupled PR barostat may be a result of the specifics of the barostat implementation. For less stringent settings, $r_l = 1.9$ nm and `nstlist` = 20, we found that also the Berendsen barostat induced significant shape asymmetries, but not the C-rescale barostat (Table 1). This further hints at the possibility that subtle differences in barostat algorithm or implementation make the MD simulations susceptible to anisotropies in the box shape evolution. In practice, semi-isotropic coupling is typically used for systems with mechanical resistance, such as a lipid bilayer spanning the xy plane. For such systems, consequences of the asymmetries detected here should be negligible.

4 Discussion

Applying barostats (and thermostats) immediately after neighbor list updates, and thus with all interactions accounted for, should stabilize the simulations even in cases where the outer cutoff r_l is somewhat too short. To leave room for efficiency optimizations for specific MD runs, the actual combination of r_l and `nstlist` in GROMACS is determined by the Verlet-buffer-tolerance VBT. In addition, GROMACS does not ensure that `nsttcouple` and `nstpcouple` are identical to `nstlist` by default. All barostat types in GROMACS act with a fixed frequency (`nstpcouple`) according to the instantaneous pressure tensor. If `nstpcouple` is not an integer multiple of `nstlist` and if the outer cutoff r_l is too short so that pair interactions are occasionally missed, then the barostat will in regular intervals operate with instantaneous pressures calculated for systems with an incomplete list of pair interactions and thus underestimated cohesion. Consequently, the barostat then acts to reduce the artificially high apparent pressure by increasing the system volume. However,

Table 1: Poor cutoff treatment results in violation of spatial isotropy of NPT Martini water system. Four pressure coupling schemes were examined: semi-isotropically coupled PR, Berendsen and C-rescale barostats, and an anisotropically coupled PR barostat. Four different combinations of r_l and `nstlist` were tested with dual pair list disabled (column 1). Columns 2-5 list the number of times the system elongated along a specific principal axis. P-values were calculated under the null hypothesis that the probabilities are equal to 1/2 for contraction and elongation along z in the semi-isotropic case, and equal to 1/3 for expansions along x , y , and z in the fully anisotropic case.

r_l (nstlist)	Semi-isotropic			Anisotropic
	Parrinello-Rahman	Berendsen	C-rescale	Parrinello-Rahman
	Contraction : Elongation p-value			$X : Y : Z$ p-value
1.9 nm (1)	1992 : 2508 $\ll 0.001$	2274 : 2226 0.484	2232 : 2268 0.709	1486 : 1489 : 1525 0.890
1.9 nm (20)	1310 : 3190 $\ll 0.001$	1997 : 2503 $\ll 0.001$	2239 : 2261 0.634	1387 : 1536 : 1577 0.004
1.28 nm (25)	1650 : 2850 $\ll 0.001$	1977 : 2523 $\ll 0.001$	2154 : 2346 $\ll 0.001$	1540 : 1322 : 1638 $\ll 0.001$
1.28 nm (25) MxN disabled				1496 : 1531 : 1473 0.768

such an increase in volume is artificial, the effect of which is only partly restored at those time points where barostat action immediately follows a neighbor list update. It is thus advantageous to set `nsttcouple = nstpcouple = n × nstlist` with integer $n \geq 1$. Indeed, the unphysical distortion of the large Martini membrane system is greatly suppressed if `nstpcouple = nstlist` even with a relatively short outer cutoff of $r_l = 1.269$ nm (Figure S4). However, the observed membrane undulation amplitude in Figure S4 ($\approx 1.1 \pm 0.3 \text{ \AA}$) is still larger compared to the expected value (≈ 0.53 ; see SI Text). Hence, setting `nsttcouple = nstpcouple = nstlist` in combination with default r_l alone may be insufficient.

Setting the outer cutoff r_l to values somewhat larger than default further stabilizes the simulations. In numerical tests, we found it sufficient to obtain the default values of r_l and `nstlist` for the same system but with double the integration time step, and then to enforce these values manually together with `nstpcouple = nstlist = nsttcouple`. With VBT set to -1 , this procedure also automatically disables the dual pair list evaluation.

We also observed that the tendency to expand the simulation box along z (Table 1) is consistent with the distinct spatial anisotropy in the errors of the diagonal elements of the pressure tensor in Figure 4. A likely source of this anisotropy is the asymmetric implementation of the MxN algorithm.^{2,8} The algorithm constructs the neighbor lists by gridding the xy plane and binning the particles along the z axis. Clusters with insufficient numbers of particles are filled with dummy particles, which have zero contribution on the cluster volume.⁸ Hence, the average cluster dimensions along the z axis would be smaller than those of the lateral axes. This effectively results in an anisotropic buffer thickness. Consequently, we expect that a larger fraction of cohesive interactions is disregarded along the z axis, which would explain the observed spatial anisotropies. To test this hypothesis, we disabled the MxN algorithm by recompiling the GROMACS MD engine with `-DGMX_SIMD=None`. We then performed 4500 replicate MD simulations of the fully anisotropically coupled system containing pure Martini solvent with $r_l = 1.28$ nm and `nstlist = 25`, as before. Consistent with our hypothesis, we found that by disabling the MxN algorithm the asymmetry in the

box shape changes disappeared (last row in Table 1).

5 Recommendations for Practitioners

5.1 Diagnosis

Deviations between target and calculated pressures in NPT simulations serve as the primary indicator of possible issues with neighbor list construction and cutoff treatment. As shown in Figure 2, missed interactions as a result of an inadequate cutoff treatment tend to result in noticeable deviations between target and actual pressures, and in small but again noticeable anisotropy as manifested by differences among the diagonal elements of the pressure tensor.

As a complication, the pressure and energy are calculated only every `nstcalcenergy` time steps, whereas neighbor lists are updated every `nstlist` time steps. If `nstcalcenergy` is an integer multiple of `nstlist`, the pressure tensor is always evaluated immediately after a neighbor list update. This synchrony can mask deviations (Figure 2B right). More frequent or asynchronous pressure calculation (e.g., by setting `nstenergy = 1`) can then reveal cutoff issues, as seen by comparing Figure 2B center and right. However, variations of `nstlist` along atomistic MD trajectories can further complicate the analysis.

Trial trajectories with pressures calculated every time step (`nstenergy = 1`) provide the basis for a more detailed analysis. Running averages (Figure 2) and averages over blocks of `nstlist` time steps (Figure 3) help to pinpoint problems with missed interactions and the resulting pressure imbalances.

Oscillatory behavior in the scalar pressure as a result of neighbor list updates and barostat action can be revealed by a power-spectral analysis. For an inadequate outer cutoff, the PSD of the scalar pressure shows distinct spikes at frequencies that are integer multiples of $1/(\text{nstlist} \times \Delta t)$, as shown in Figure 5. The amplitude of these frequencies is modulated by the frequency of updates of the inner neighbor list. By enforcing a larger outer cutoff r_l with VBT set to -1 and without dual cutoff, these oscillations and the resulting features in the

power spectral density disappear (lower curve in Figure 5). For the pressure in MD simulations of TIP3P water, we similarly observed distinct spikes in the PSD at integer multiples of $1/(\text{nstlist} \times \Delta t)$, which disappeared for large values of r_l and with the MxN algorithm disabled (Figure S5). Note that we used the v-rescale thermostat to avoid the oscillatory contributions to the pressure PSD in simulations with the Nosé-Hoover thermostat.^{36,37}

5.2 Recommendations

For the existing software, pending updates, we recommend to manually define the outer cutoff r_l and the neighbor list update frequency `nstlist` of a system to be the default values obtained from the identical system with twice the time step ($\Delta t = 2\text{dt}$). This requires setting `verlet-buffer-tolerance = -1`. For Martini systems using the *new-rf* parameters with `nstlist = 20`, r_l should be at least 1.35 nm. For atomistic systems with `nstlist = 20`, r_l should be at least 1.3 nm. Finally, `nsttcouple` and `nstpcouple` should be exact multiples

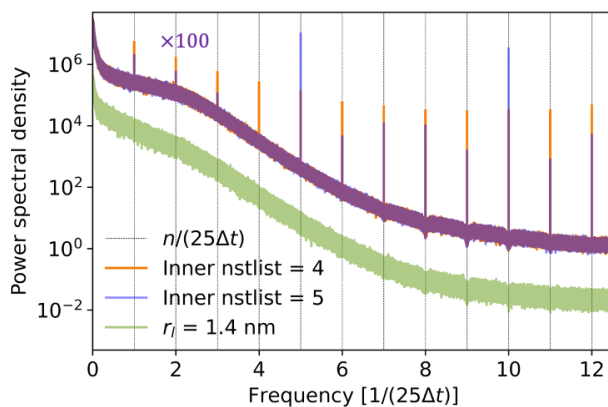


Figure 5: Power spectral density (PSD) of the scalar pressure calculated for Martini water in the NVT ensemble. Results are shown for $r_l = 1.27$ nm and `nstlist = 25` (top curves; scaled by factor 100) with dual pair list and inner neighbor list updates every four (orange) and five (blue line) integration steps, respectively. Except for the oscillations at frequency multiples of $1/(\text{nstlist} \times \Delta t)$, the two curves nearly superimpose. Also shown (fluorescent green at bottom) is the PSD for a simulation with $r_l = 1.422$ nm and `nstlist = 25` without dual cutoff, which does not show oscillations.

of `nstlist`. However, these measures may significantly reduce the computational performance, especially for large systems. Depending on the hardware specifications, a decrease in performance of up to $\sim 30\%$ may be observed.

In future, to minimize the impact on computational cost, a spatially isotropic neighbor search algorithm is desired also for the SIMD architecture. Spatial isotropy can for instance be restored to a significant degree with minimum computational overhead if the axis of the one-dimensional search is cycled through x , y , and z instead of keeping it fixed at z . More conservative default choices of r_l and `nstlist` will also help to ensure stable MD simulations. The choice of r_l and `nstlist` can be based on simple expressions for the probability of missed interactions such as eq 1. Finally, all `nstX` (where X is `energy`, `tcouple`, `pcouple` etc.) variables should be exact multiples of `nstlist`. In this way, energy output as well as thermostat and barostat action benefit from freshly updated neighbor lists. If the dual pair list is enabled, the interval between updates of inner and outer neighbor lists should be set similarly to ensure that barostats in particular act with all interactions considered.

6 Concluding Remarks

The efficient calculation of pair interactions is at the heart of modern MD simulation codes. The construction of neighbor lists is a critical factor to reduce computational cost and achieve near-linear scaling with system size as well as efficient parallelization. We found that even a small number of missed long-range interactions can impact the accuracy and qualitative behavior of MD simulations. For large membrane systems, we observed that a slight but systematic error in the pressure with default settings for cutoff distance and neighbor list updates resulted in artificial membrane buckling caused by the counter-action of the barostat. We suspect that similar behavior motivated the imposition of restraints on lipid motion normal to the membrane in earlier studies.^{15–20} Slight but systematic anisotropies in the errors of the pressure result in anisotropic deformations of the box shape. We also observed

distinct beating effects in the pressure time series. As underlying causes of these different but related problems, we identified (i) missed pair interactions as a result of too infrequent neighbor list updates, (ii) slight anisotropies in the neighbor list construction, and (iii) neighbor list and barostat updates at incommensurable time intervals. In most current MD simulations, in particular at all-atom resolution, we expect the slight pressure imbalances to have minimal impact. However, as MD simulations are used to study ever-larger systems, such as the coarse-grained membrane systems in Figure 1, the issues will have to be addressed. Whereas the immediate fixes of longer outer cutoff lengths r_l and disabled dual neighbor lists (VBT = -1) are associated with increased computational cost, we are confident that with the root causes identified, adjustments in algorithms and code can be made that will resolve the issues without major computational overhead.

Acknowledgement

The authors thank the International Max Planck Research School (IMPRS) on Cellular Biophysics (HK), the Max Planck Society (HK, BF, GH), and the Alexander von Humboldt-Foundation (BF) for support. The authors thank Sebastian Kehl (MPDCF) for discussions about the gmx source code.

References

- (1) Gelpi, J.; Hospital, A.; Goñi, R.; Orozco, M. Molecular dynamics simulations: Advances and applications. *Advances and Applications in Bioinformatics and Chemistry* **2015**, 37.
- (2) Abraham, M. J.; Murtola, T.; Schulz, R.; Páll, S.; Smith, J. C.; Hess, B.; Lindahl, E. GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. *SoftwareX* **2015**, 1-2, 19–25.

- (3) Lundborg, M.; Lindahl, E. Automatic GROMACS Topology Generation and Comparisons of Force Fields for Solvation Free Energy Calculations. *The Journal of Physical Chemistry B* **2015**, *119*, 810–823, PMID: 25343332.
- (4) Marrink, S. J.; Risselada, H. J.; Yefimov, S.; Tieleman, D. P.; de Vries, A. H. The MARTINI Force Field: Coarse Grained Model for Biomolecular Simulations. *The Journal of Physical Chemistry B* **2007**, *111*, 7812–7824.
- (5) Gecht, M.; Siggel, M.; Linke, M.; Hummer, G.; Köfinger, J. MDBenchmark: A toolkit to optimize the performance of molecular dynamics simulations. *The Journal of Chemical Physics* **2020**, *153*, 144105.
- (6) Kutzner, C.; Páll, S.; Fechner, M.; Esztermann, A.; de Groot, B. L.; Grubmüller, H. Best bang for your buck: GPU nodes for GROMACS biomolecular simulations. *Journal of Computational Chemistry* **2015**, *36*, 1990–2008.
- (7) Páll, S.; Abraham, M. J.; Kutzner, C.; Hess, B.; Lindahl, E. Tackling Exascale Software Challenges in Molecular Dynamics Simulations with GROMACS. Solving Software Challenges for Exascale. Cham, 2015; pp 3–27.
- (8) Páll, S.; Hess, B. A flexible algorithm for calculating pair interactions on SIMD architectures. *Computer Physics Communications* **2013**, *184*, 2641–2650.
- (9) Buyya, R.; Vecchiola, C.; Selvi, S. T. In *Mastering Cloud Computing*; Buyya, R., Vecchiola, C., Selvi, S. T., Eds.; Morgan Kaufmann: Boston, 2013; pp 29–70.
- (10) Hess, B.; Kutzner, C.; van der Spoel, D.; Lindahl, E. GROMACS 4: Algorithms for Highly Efficient, Load-Balanced, and Scalable Molecular Simulation. *Journal of Chemical Theory and Computation* **2008**, *4*, 435–447.
- (11) de Jong, D. H.; Baoukina, S.; Ingólfsson, H. I.; Marrink, S. J. Martini straight: Boost-

- ing performance using a shorter cutoff and GPUs. *Computer Physics Communications* **2016**, *199*, 1–7.
- (12) Abraham, M.; Alekseenko, A.; Bergh, C.; Blau, C.; Briand, E.; Doijade, M.; Fleischmann, S.; Gapsys, V.; Garg, G.; Gorelov, S.; Gouaillardet, G.; Gray, A.; Irngang, M. E.; Jalalypour, F.; Jordan, J.; Junghans, C.; Kanduri, P.; Keller, S.; Kutzner, C.; Lemkul, J. A.; Lundborg, M.; Merz, P.; Miletić, V.; Morozov, D.; Páll, S.; Schulz, R.; Shirts, M.; Shvetsov, A.; Soproni, B.; van der Spoel, D.; Turner, P.; Uphoff, C.; Villa, A.; Wingbermühle, S.; Zhmurov, A.; Bauer, P.; Hess, B.; Lindahl, E. GROMACS 2023.1 Manual. 2023; <https://doi.org/10.5281/zenodo.7852189>.
- (13) Berendsen, H. J. C.; Postma, J. P. M.; van Gunsteren, W. F.; DiNola, A.; Haak, J. R. Molecular dynamics with coupling to an external bath. *The Journal of Chemical Physics* **1984**, *81*, 3684–3690.
- (14) Parrinello, M.; Rahman, A. Polymorphic transitions in single crystals: A new molecular dynamics method. *Journal of Applied Physics* **1981**, *52*, 7182–7190.
- (15) Ingólfsson, H. I.; Melo, M. N.; van Eerden, F. J.; Arnarez, C.; Lopez, C. A.; Wassenaar, T. A.; Periole, X.; de Vries, A. H.; Tieleman, D. P.; Marrink, S. J. Lipid Organization of the Plasma Membrane. *Journal of the American Chemical Society* **2014**, *136*, 14554–14559, PMID: 25229711.
- (16) Larsen, A. H. Molecular Dynamics Simulations of Curved Lipid Membranes. *International Journal of Molecular Sciences* **2022**, *23*, 8098.
- (17) Vögele, M.; Köfinger, J.; Hummer, G. Hydrodynamics of Diffusion in Lipid Membrane Simulations. *Physical Review Letters* **2018**, *120*, 268104.
- (18) Duboué-Dijon, E.; Hémin, J. Building intuition for binding free energy calculations: Bound state definition, restraints, and symmetry. *The Journal of Chemical Physics* **2021**, *154*, 204101.

- (19) Mori, T.; Miyashita, N.; Im, W.; Feig, M.; Sugita, Y. Molecular dynamics simulations of biological membranes and membrane proteins using enhanced conformational sampling algorithms. *Biochimica et Biophysica Acta (BBA) - Biomembranes* **2016**, *1858*, 1635–1651, New approaches for bridging computation and experiment on membrane proteins.
- (20) Kolossváry, I.; Sherman, W. A Different Kind of Restraint Suitable for Molecular Dynamics Simulations. **2022**,
- (21) Tribello, G. A.; Bonomi, M.; Branduardi, D.; Camilloni, C.; Bussi, G. PLUMED 2: New feathers for an old bird. *Computer Physics Communications* **2014**, *185*, 604–613.
- (22) Frenkel, D.; Smit, B. *Understanding Molecular Simulation: From Algorithms to Applications*, 2nd ed.; Computational Science Series; Academic Press: San Diego, 2002; Vol. 1.
- (23) Allen, M. P.; Tildesley, D. J. *Computer Simulation of Liquids*; Oxford University Press, 2017.
- (24) Essmann, U.; Perera, L.; Berkowitz, M. L.; Darden, T.; Lee, H.; Pedersen, L. G. A smooth particle mesh Ewald method. *The Journal of Chemical Physics* **1995**, *103*, 8577–8593.
- (25) Sega, M.; Dellago, C. Long-range dispersion effects on the water/vapor interface simulated using the most common models. *The Journal of Physical Chemistry B* **2017**, *121*, 3798–3803.
- (26) Panigrahy, R. *Lecture Notes in Computer Science*; Springer Berlin Heidelberg, pp 387–398.
- (27) Páll, S.; Zhmurov, A.; Bauer, P.; Abraham, M.; Lundborg, M.; Gray, A.; Hess, B.; Lindahl, E. Heterogeneous parallelization and acceleration of molecular dynamics simulations in GROMACS. *The Journal of Chemical Physics* **2020**, *153*, 134110.

- (28) Thompson, A. P.; Aktulga, H. M.; Berger, R.; Bolintineanu, D. S.; Brown, W. M.; Crozier, P. S.; in 't Veld, P. J.; Kohlmeyer, A.; Moore, S. G.; Nguyen, T. D.; Shan, R.; Stevens, M. J.; Tranchida, J.; Trott, C.; Plimpton, S. J. LAMMPS - a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales. *Computer Physics Communications* **2022**, *271*, 108171.
- (29) Helfrich, W. Elastic Properties of Lipid Bilayers: Theory and Possible Experiments. *Zeitschrift für Naturforschung C* **1973**, *28*, 693–703.
- (30) Bhaskara, R. M.; Grumati, P.; Garcia-Pardo, J.; Kalayil, S.; Covarrubias-Pinto, A.; Chen, W.; Kudryashev, M.; Dikic, I.; Hummer, G. Curvature induction and membrane remodeling by FAM134B reticulon homology domain assist selective ER-phagy. *Nature Communications* **2019**, *10*.
- (31) Wassenaar, T. A.; Ingólfsson, H. I.; Böckmann, R. A.; Tieleman, D. P.; Marrink, S. J. Computational Lipidomics with insane: A Versatile Tool for Generating Custom Membranes for Molecular Simulations. *Journal of Chemical Theory and Computation* **2015**, *11*, 2144–2155, PMID: 26574417.
- (32) Bussi, G.; Donadio, D.; Parrinello, M. Canonical sampling through velocity rescaling. *The Journal of Chemical Physics* **2007**, *126*, 014101.
- (33) Bernetti, M.; Bussi, G. Pressure control using stochastic cell rescaling. *The Journal of Chemical Physics* **2020**, *153*, 114107.
- (34) Mark, P.; Nilsson, L. Structure and Dynamics of the TIP3P, SPC, and SPC/E Water Models at 298 K. *The Journal of Physical Chemistry A* **2001**, *105*, 9954–9960.
- (35) Jo, S.; Kim, T.; Iyer, V. G.; Im, W. CHARMM-GUI: A web-based graphical user interface for CHARMM. *Journal of Computational Chemistry* **2008**, *29*, 1859–1865.

- (36) Hoover, W. G. Canonical dynamics: Equilibrium phase-space distributions. *Physical Review A* **1985**, *31*, 1695–1697.
- (37) Nosé, S. A unified formulation of the constant temperature molecular dynamics methods. *The Journal of Chemical Physics* **1984**, *81*, 511–519.
- (38) Welch, P. The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Transactions on Audio and Electroacoustics* **1967**, *15*, 70–73.
- (39) Hernández-Muñoz, J.; Bresme, F.; Tarazona, P.; Chacón, E. Bending Modulus of Lipid Membranes from Density Correlation Functions. *Journal of Chemical Theory and Computation* **2022**, *18*, 3151–3163.