

TriMem: A Parallelized Hybrid Monte Carlo Software for Efficient Simulations of Lipid Membranes

Marc Siggel^{*, 1, 2, 3, a)} Sebastian Kehl^{*, 4} Klaus Reuter⁴ Jürgen Köfinger³ and Gerhard Hummer^{3, 5, b)}

¹⁾European Molecular Biology Laboratory Hamburg, Notkestraße 85, 22607 Hamburg, Germany

²⁾Centre for Structural Systems Biology (CSSB), Notkestraße 85, 22607 Hamburg, Germany

³⁾Department of Theoretical Biophysics, Max Planck Institute of Biophysics, Max-von-Laue-Straße 3, 60438 Frankfurt am Main, Germany

⁴⁾Max Planck Computing and Data Facility, Gießenbachstraße 2, 85748 Garching, Germany

⁵⁾Institute of Biophysics, Goethe University Frankfurt, 60438 Frankfurt am Main, Germany

(Dated: 25 May 2022)

Lipid membranes are integral building blocks of living cells and perform a multitude of biological functions. Currently, molecular simulations of cellular-scale membrane structures at atomic resolution are nearly impossible, due to their size, complexity, and the large times-scales required. Instead, elastic membrane models are used to simulate membrane topologies and transitions between them, and to infer their properties and functions. Unfortunately, efficiently parallelized open-source simulation code to do so has been lacking. Here, we present TriMem, a parallel hybrid Monte Carlo simulation engine for triangulated lipid membranes. The kernels are efficiently coded in C++ and wrapped with Python for ease-of-use. The parallel implementation of the energy and gradient calculations and of Monte Carlo flip moves of edges in the triangulated membrane enable us to simulate also large and highly curved sub-cellular structures. For validation, we reproduce phase diagrams of vesicles with varying surface-to-volume ratios and area difference. The software can tackle a range of membrane remodelling processes on sub-cellular and cellular scales. Additionally, extensive documentation make the software accessible to the broad biophysics and computational cell biology communities.

I. INTRODUCTION

Living cells are bounded by lipid membranes, and the interior of eukaryotic cells is filled with membranous organelles. Cellular membrane structures are highly dynamic, strongly curved, and branched^{1,2}. Membranes are flexible and behave like two-dimensional fluids. Super-resolution microscopy and cryo-electron tomography (cryo-ET) have given unprecedented insights into the spatial organization of cellular membranes and their associated protein structures³. However, a detailed and comprehensive biophysical understanding of the mechanisms underlying organellar membrane reshaping in the cellular context remains elusive.

Molecular dynamics simulations are increasingly becoming state of the art for understanding membrane biophysics in the cellular context⁴⁻⁷. Multi-scale approaches combining all-atom, coarse-grained and meso-scale methods are used to tackle various aspects of membrane remodeling^{8,9}. All-atom models are used to gain detailed insight into protein-lipid interactions and protein-protein interactions inside membranes but are limited in size and time scale⁹. Coarse-grained particle-based simulations approaches make it possible to study large membrane-associated complexes and complex membrane shape changes⁹⁻¹². However, simulations of large cell-scale membrane systems are currently impossible with molecular models at atomic or near-atomic resolution. Computational resources limit the time and length scales

that can be studied with particle-based approaches. Therefore, sub-cellular and cell-scale remodelling processes of organelles, vesicles, and cells are commonly studied by using meso-scale models and continuum approaches, primarily membrane-elastic theory^{13,14}.

Dynamic triangulated surfaces (DTSs) have emerged as a useful meso-scale model to solve the Helfrich Hamiltonian numerically and study large-scale membrane-shaping processes. Initially, DTSs were used to study shaping and properties of fluctuating giant unilamellar vesicle (GUV)¹⁵⁻¹⁸. The discretization of membranes makes it possible to sample non-axisymmetric shapes¹⁷ beyond symmetric arc length parameterizations^{19,20}. This method has been applied to a number of biological processes with increasing complexity including nano-particle wrapping^{21,22}, membrane tubulation²³, formation of autophagic vesicles²⁴, formation of Golgi stacks²⁵, and protein-induced membrane budding^{26,27}.

Here, we present TriMem, an open-source software package for efficient simulation and optimization of DTSs using a parallelized hybrid Monte Carlo (MC) approach. This approach overcomes limitations of existing serial implementations. In such non-parallel approaches, the number of vertices of the triangulated membrane representation is severely limited to keep computational times manageable. Yet, large numbers of vertices are needed to describe highly curved membranes, which are ubiquitous in cellular organelles such as mitochondria or the tubular ER. As a further challenge, the commonly used random single-particle MC moves are inefficient with respect to sampling. Importantly, only few code bases for these types of simulations have been made freely available to the broader community despite the large amount of scientific research in this area, and none of them are parallelized yet^{28,29}.

^{a)}Electronic mail: marc.siggel@embl-hamburg.de

^{b)}Electronic mail: gerhard.hummer@biophys.mpg.de

* These authors contributed equally

This paper is structured as follows. First, we briefly recapitulate the discretized version of Helfrich theory and the Helfrich Hamiltonian used throughout this work (section II). We then introduce the TriMem software, including its implementation strategy and algorithms (section III). We present benchmark results on timings for varying mesh sizes and compare the performance to a serial single-core implementation (section III D). We validate the software by reproducing well established phase diagrams of vesicle shapes (section IV). Finally, we provide a comprehensive outlook on the expected impact of our software and on future developments (section VI)

II. MEMBRANE MODEL

Membrane-elastic theory describes membranes as 2D surfaces with fluid properties embedded in 3D space. In its simplest form, the bending free energy of a membrane can be written in terms of the so-called Helfrich Hamiltonian^{13,14},

$$E_B = \frac{\kappa_B}{2} \oint dA (2H - C_0)^2 + \kappa_G \oint dA K_G. \quad (1)$$

with $H = 0.5(c_1 + c_2)$ the mean curvature and $K_G = c_1 c_2$ the Gaussian curvature, where c_1 and c_2 are the principal curvatures. C_0 signifies the intrinsic curvature of the membrane which is modulated by a variety of factors, including proteins, lipid composition and membrane asymmetry. κ_B and κ_G are the bending rigidity and Gaussian bending modulus, respectively, and describe the elastic properties of the studied membrane. The second term, the Gaussian curvature, can usually be neglected because it is constant in the absence of changes of the topology according to the Gauss-Bonnet theorem. Multiple extensions and variations of the Helfrich Hamiltonian have been introduced to include a broad spectrum of external factors such as area difference, protein inclusions, and osmotic pressure. Variational minimization of the Helfrich Hamiltonian with respect to the membrane shape leads to fourth-order differential equations, which have been solved analytically only for a limited number of cases of high symmetry^{30,31}. However, numerical solutions are achievable.

To make simulations of membrane shapes computationally tractable and amenable to integration and MC sampling, the surface is discretized as a triangular mesh, which can in principle represent arbitrary surface shapes. The Hamiltonian of this discretized system (H_{tot}) is more complex than the original Helfrich Hamiltonian. We decompose the energy as

$$H_{\text{tot}} = E_B + E_V + E_A + E_{\Delta A} + E_T + E_R \quad (2)$$

where E_B is the bending energy, E_V the volume energy to maintain the total internal volume, E_A the area energy to maintain the total surface area, $E_{\Delta A}$ the area-difference energy (ADE), E_T the tethering potential, and E_R the repulsive potential. In the following, we explain how we calculate these energies from triangulated surfaces.

To introduce the energy terms, we first have to define the DTS. A closed surface system consists of $N_T = 2(N_V - 2)$ triangles, with N_V vertices connected by $3(N_V - 2)$ tethers. To

this end we introduce a triangulation $\mathcal{T} := (\mathbf{x}, F)$ as a tuple of vertex positions $\mathbf{x} := \{x_i\}_{i=1}^{N_V}$ (with a single vertex $x_i \in \mathbb{R}^3$) and triangles $F := \{f_i\}_{i=1}^{N_T}$. Thereby a single oriented triangle $f_i := \{(i, j, k) | i, j, k \in [1, N_V], i \neq j \neq k\}$ is given by an ordered 3-tuple indexing into the vertices \mathbf{x} . For the convenience of notation, we define the vertex-vertex connectivity $v_i^v := \{j | j \in f_k \forall f_k \in F, i \neq j\}$, the vertex-face connectivity $v_i^f := \{j | j \in f_j \forall f_j \in F\}$, the set of edges $E := \{(i, j) | i, j \in f_k \forall f_k \in F, i \neq j\}$ and the edge-face connectivity $e_{ij}^f := \{k | i, j \in f_k\}$

We discretize the calculation of the Helfrich Hamiltonian H_{tot} on the DTS using a vertex-averaged formulation as previously described extensively in ref. 17. The total bending energy E_B is given by the sum over all vertices of the bending energy per vertex

$$E_B = 2\kappa_B \sum_{i=1}^{N_V} \frac{\hat{M}_i^2}{\hat{A}_i} \quad (3)$$

The area \hat{A}_i per vertex i is calculated as

$$\hat{A}_i = \frac{1}{3} \sum_{j \in v_i^f} A_j, \quad (4)$$

where A_j is the area of the single triangle f_j . The average mean curvature \hat{M}_i associated with vertex i can be calculated as sum over all adjacent edges

$$\hat{M}_i = \frac{1}{4} \sum_{j \in v_i^v} r_{ij} \phi_{ij}, \quad (5)$$

where $r_{ij} = \|x_j - x_i\|$ and ϕ_{ij} is the angle between the oriented normal vectors n_a, n_b of the triangles in e_{ij}^f , i.e., adjacent to the edge (i, j) , calculated as $\cos(\phi_{ij}) = n_a \cdot n_b$.

The total mean curvature of the system is then computed by

$$M = \sum_{i=1}^{N_V} \hat{M}_i. \quad (6)$$

The energies for the volume and area constraint are given by E_V and E_A , respectively, and can be written as simple harmonic potentials

$$E_V = \kappa_V \left(\frac{V}{V_0} - 1 \right)^2 \quad (7)$$

and

$$E_A = \kappa_A \left(\frac{A}{A_0} - 1 \right)^2 \quad (8)$$

with κ_V, κ_A being the coupling constants, respectively. The corresponding reference values are given by V_0 and A_0 . κ_V and κ_A are usually chosen several orders of magnitude larger than κ_B to avoid non-physical area and volume fluctuations. The total area A is calculated as

$$A = \sum_{j=1}^{N_T} A_j. \quad (9)$$

Also the volume V can be calculated as a discrete sum,

$$V = \sum_{j=1}^{N_T} V_j = \frac{1}{3} \sum_{i=1}^{N_T} (R_i \cdot n_i) A_j, \quad (10)$$

where V_i is the signed sub-volume of a single triangle, R_i the position vector, and n_i the unit normal vector.

The ADE, $E_{\Delta A}$, is given by

$$E_{\Delta A} = \kappa_{\Delta A} \left(\frac{\Delta A}{\Delta A_0} - 1 \right)^2 \quad (11)$$

The area difference ΔA of a bilayer with respect to its shape can be written similarly as previously defined for continuous representations^{32,33} as

$$\Delta A = 2h \sum_{i=1}^{N_V} \hat{M}_i, \quad (12)$$

where the sum runs over all vertices and \hat{M}_i is the mean curvature associated with each vertex, respectively, and h is the thickness of the neutral surfaces. When choosing $\kappa_{\Delta A} \rightarrow \infty$ we recover the bilayer couple model where ΔA acts as a constraint³³. Numerically we convert this by choosing $\kappa_{\Delta A}$ to be on the order of $10^5 - 10^6 k_B T$, where k_B is Boltzmann's constant and T is the temperature²³.

The overall tether-energy is given by a coupling constant κ_T times the sum over the tether pair-potential $T(r_{ij})$ over all edge lengths r_{ij} ,

$$E_T = \kappa_T \sum_{(i,j) \in E} T(r_{ij}). \quad (13)$$

This energy serves to constrain the edge length r_{ij} and to guarantee the efficient and accurate simulation of fluid triangulated surfaces. Previous work^{17,23} used discrete flat bottom potentials; however, these are not amenable to smooth time integration. In order to use the tether potential in hybrid MC simulations a continuous representation is required¹⁸. To be able to use large integration time steps Δt , it is crucial to avoid diverging branches that are present in previous formulations. Therefore, we use a continuous tether potential of the functional form

$$T(l) = \begin{cases} e^{\frac{l}{l_{c1}}} l^{-r} & \text{if } l \leq l_{c1} \\ r^{r+1} (l - l_{c0})^r & \text{if } l \geq l_{c0} \\ 0 & \text{otherwise} \end{cases} \quad (14)$$

with the l_{c1} and l_{c0} being the lower- and upper onset of penalization and the slope $r \in \mathbb{N}_1$ of the potential well.

The overall repulsion energy is given by

$$E_R = \kappa_R \sum_{i=1}^{N_V} \sum_{j=1, j \notin S_i}^{N_V} R(r_{ij}), \quad (15)$$

where S_i is a set of excluded vertices for vertex i , e.g., the set of directly connected neighbors. The $\mathcal{O}(N_V^2)$ complexity involved in the evaluation of this potential is in practice reduced

by the use of efficient neighbor tracking techniques³⁴. Mimicking the repulsive electrostatic membrane-membrane interactions, mesh intersection is prevented by the introduction of a penalty that applies a repulsive force on pairs of non-bonded vertices that are below a certain threshold distance l_{c1} . Such a penalty is implemented by the repulsive branch of the tether penalty, Eq. (14). It is given by

$$R(l) = \begin{cases} e^{\frac{l}{l_{c1}}} l^{-r} & \text{if } l \leq l_{c1} \\ 0 & \text{otherwise.} \end{cases} \quad (16)$$

The functional form of the membrane-membrane interaction could be adapted to model various properties of membrane adhesion and repulsion.

III. ALGORITHMS AND IMPLEMENTATION FOR SIMULATION AND ENERGY MINIMIZATION

In the following we discuss how to efficiently sample the configurational space of the DTS using the hybrid MC approach. First, we introduce the general sampling strategies and discuss how they are implemented. Next, we highlight the issues and solutions for efficient parallelization and evaluate their performance. Finally, we discuss important features to efficiently equilibrate and minimize structures in practice. The different algorithmic components introduced below are shown in the flow-chart in Fig. 1. We apply these tools in Sec. IV.

A. General Sampling Strategy

The statistical properties associated with the Hamiltonian (2) follow the canonical distribution function given by

$$p(\mathbf{x}, F) \propto \exp[-\beta H_{\text{tot}}(\mathbf{x}, F)]. \quad (17)$$

where we use $\beta = 1$ for the inverse temperature using reduced units. We use a Markov chain MC procedure to sample configurations (\mathbf{x}, F) . To account for the compound nature of the state of the triangulation, given by vertex position as well as vertex connectivity, we generate new samples using the Metropolis algorithm with two alternating, well established MC moves: (i) global vertex displacements and (ii) edge flips. After each step the resulting configuration is accepted or rejected according to the Metropolis criterion¹⁵ by evaluating the differences in the Hamiltonian H_{tot} as a function of \mathbf{x} and F . Edge flips ensure membrane fluidity.

Vertex displacements are generated by a Hybrid Monte Carlo scheme (HMC)³⁵. HMC draws on the lifting of the vertex coordinates \mathbf{x} onto an artificial phase-space $\mathbf{x} \mapsto (\mathbf{x}, \mathbf{p})$, where the components of the vector \mathbf{p} are the momenta for all vertices. On this phase-space, we impose a symplectic structure via the lifted Hamiltonian

$$H(\mathbf{x}, \mathbf{p}) = H_{\text{tot}}(\mathbf{x}) + \frac{1}{2} \mathbf{p}^\top \mathbf{M}^{-1} \mathbf{p}. \quad (18)$$

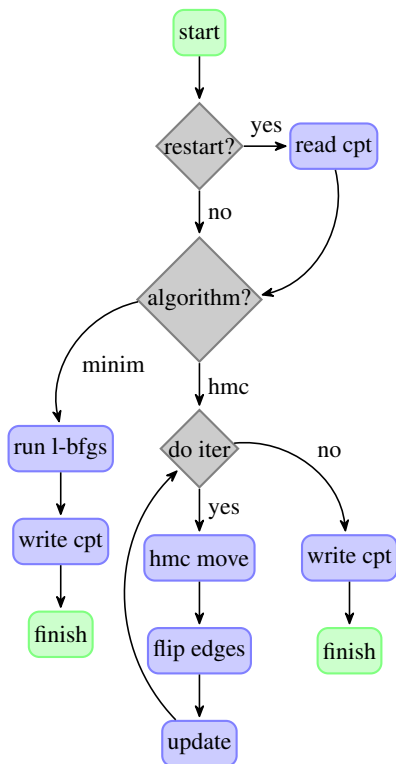


FIG. 1. Algorithmic flow-chart of the TriMem software. ‘run l-bfgs’ refers to section III F, ‘hmc move’ and ‘flip edges’ are introduced in Algorithm 1 and Algorithm 2, respectively. ‘update’ summarizes the update of the reference parameters η , see Sec. III E, and the temperature T , see Sec. III G.

The transpose of a vector is indicated by superscript T , and \mathbf{M}^{-1} is the inverse of the diagonal mass matrix \mathbf{M} . We exploit HMC to generate state transitions $(\mathbf{x}_n, \mathbf{p}_n) \rightarrow (\mathbf{x}_{n+1}, \mathbf{p}_{n+1})$ in the high-dimensional phase-space with high efficiency. Transitions generated in this way can then trivially be un-lifted $(\mathbf{x}_{n+1}, \mathbf{p}_{n+1}) \rightarrow \mathbf{x}_{n+1}$ to give a new sample \mathbf{x}_{n+1} . In practice, we use symplectic time-integration schemes such as the leapfrog/Verlet integration method³⁴. We show the general concept of a HMC step from $n \rightarrow n+1$ in Algorithm 1. Verlet-type integration accurately conserves $H(\mathbf{x}, \mathbf{p})$. As a result, the global vertex moves of HMC are accepted with a probability close to one.

Conceptually, a HMC step consists of a short molecular dynamics (MD) simulation of the vertices, with velocities drawn from a Maxwell-Boltzmann distribution and forces given by the negative gradient of H_{tot} with respect to the vertex positions (see appendix A). Subsequently, the lifted Hamiltonian is subject to the Metropolis criterion. The time integration parameter Δt and L can be used to tune the efficiency of the HMC algorithm. The time step Δt influences the acceptance probability within the Metropolis criterion. The number of steps L affects the mixing of the generated Markov chain. Tuning of these parameters is crucial for an efficient sampling scheme. The mass matrix \mathbf{M} is another free parameter that can be tuned to optimize the sampling performance. In this work we use a single mass m , $\mathbf{M} = m\mathbf{I}$, and set $m = 1$ as a default. This

Algorithm 1 One step of the HMC algorithm. We randomly draw new momenta from a Maxwell-Boltzmann distribution, perform a number of integration steps, and accept or reject the resulting configuration.

function HMC STEP($x_n, \Delta t, L, T = 1$)

$$p_0 \sim \mathcal{N}(0, m\mathbf{I})$$

$$s_0 \leftarrow x_n$$

$$i \leftarrow 0$$

while $i < L$ **do**

$$p_{i+1/2} = p_i - \frac{\Delta t}{2} \nabla H_{\text{tot}}(x)|_{x=s_i}$$

$$s_{i+1} = s_i + \Delta t \mathbf{M}^{-1} p_{i+1/2}$$

$$p_{i+1} = p_{i+1/2} - \frac{\Delta t}{2} \nabla H_{\text{tot}}(x)|_{x=s_{i+1}}$$

$$i \leftarrow i + 1$$

end while

$$\alpha \leftarrow \min(1, \exp(H(s_0, p_0) - H(s_L, p_L)) / T)$$

$$\sigma \sim \mathcal{U}(0, 1)$$

if $\sigma \leq \alpha$ **then**

return s_L

else

return x_n

end if

end function

setting has proven efficient in practice. For sampling according to Eq. (17), the application of the HMC framework leads to an enormous gain in efficiency compared to a sequential single-vertex-move dynamic. In particular in a high dimensional regime, i.e., large number of vertices N_V , HMC is essential for efficient sampling.

Edge flips ensure the 2D-fluidity of the membrane and alleviate a possible bias due to a fixed vertex connectivity. In an edge flip, the shared edge (i, j) of two adjacent triangles (i, j, k) and (i, l, j) is changed to (k, l) , resulting in modified triangles (i, l, k) and (j, k, l) (Fig. 2).

One step of edge flips consists of a sweep over a predefined fraction $\gamma \in [0, 1]$ (in the following also referred to a target flip rate) of the set of edges E . It is conceptually depicted in Algorithm 2.

B. Implementation

The vast part of the computational workload of the algorithms is the evaluation of the Hamiltonian and its gradient (see also Appendix A). This evaluation is needed from within the vertex moves as well as in the edge-flip move routine. At the core of the energy computation is the evaluation of the vertex-averaged quantities in Eqs. (4) and (5). Efficient evaluation of the vertex-connectivity is thus essential for overall performance. To this end, we utilize the concept of the half-edge data structure. This data structure for polygonal meshes efficiently tracks incidence information of edges, vertices and faces^{36,37}. In particular, we make use of the OpenMesh library, which provides a generic implementation of the half-edge data structure in C++³⁸ that can handle various mesh

Algorithm 2 One sweep of edge flips. The input is the set of triangles F , the set of edges E , and the fraction γ of flips to attempt. It makes use of the Hamiltonian $H(F)$ as a function of the mesh triangles, and the routines `shuffle_edges(E)` to shuffle the list of edges and `flip_egde(e)`, which takes an edge, flips it, and alters the mesh connectivity. The routine called `update_hamiltonian(h, F, e)` incrementally updates the total energy for the changed flip patch (see Fig. 2).

```

function FLIP_SWEEP( $F, E, \gamma$ )
   $h_n \leftarrow H(F)$ 
   $E \leftarrow \text{shuffle\_edges}(E)$ 
  for  $i$  in range( $\gamma|E|$ ) do
    flip_egde( $e_i$ )
     $h_{n+1} \leftarrow \text{update\_hamiltonian}(h_n, F, e_i)$ 
     $\alpha \leftarrow \min(1, \exp(h_n - h_{n+1}))$ 
     $\sigma \sim \mathcal{U}(0, 1)$ 
    if  $\sigma \leq \alpha$  then            $\triangleright$  evaluate Metropolis criterion
       $h_n \leftarrow h_{n+1}$ 
    else
      flip_egde( $e_i$ )            $\triangleright$  flip back in case of rejection
    end if
  end for
end function

```

sizes and geometries as input. To fully leverage the OpenMesh interface, we implement the evaluation of the Hamiltonian and its gradient in C++ and provide bindings to Python using the pybind11 Python package³⁹. This approach combines the efficiency of a compiled programming language with the convenience of the well-established Python ecosystem for clear and extensible algorithm development. Using C++ for the computationally intensive evaluations further allows us to exploit shared-memory parallelism and the speedup offered by modern multi-core architectures. To achieve this, we use OpenMP⁴⁰ to parallelize the loops that involve scans or reductions over all vertices in the triangulation, such as Eqs. (6)-(10) and (12) and the respective gradient evaluations (see Sec. III D). While parallelization of both the vertex moves and components of the integration of short trajectories is straightforward, flip moves are more complex and require a more detailed discussion.

C. A strategy for the parallel evaluation of edge flips

To increase the computational efficiency of flip moves, we pre-select a set of possible flips and evaluate in parallel independent changes of geometric membrane properties and of the energy. We then evaluate sequentially the corresponding Monte Carlo flip moves, including the associated energy changes from the aforementioned pre-calculated quantities. We designed this Algorithm 3 to alleviate two of the main efficiency bottle necks of Algorithm 2: (i) flip execution, i.e., the technical realization of the change in vertex connectivity that is provided by the underlying data structure; and (ii) energy evaluation, i.e., the evaluation of the change in the Hamilto-

nian due to the change in connectivity. Both components are crucially influenced by the evaluation of the vertex connectivity.

Due to the edge-based connectivity information implemented by the half-edge data structure, flip execution provided by OpenMesh is realized efficiently by simply swapping edge connectivity³⁸. Since the individual edge flip can technically be performed efficiently, the main computational workload results from the energy evaluation of the flip-patch associated to an edge flip. A straightforward exploitation of the speed-up provided by the evaluation of the change of the Hamiltonian for several edge flips in parallel is, however, still complicated by two intermingled issues.

First, individual edge flips tend to have very low acceptance probabilities due to the rather strict penalty on neighborhood distances that is imposed by the tether potential (Eq. (14)). The evaluation of the Metropolis criterion for a flip of several edges at once thus suffers from an exponential decay of the acceptance probability due the multiplicative nature of the acceptance probabilities of single flips. Consequently, the acceptance of edge flips must be evaluated sequentially to establish a reasonable flip rate ε , which is defined as

$$\varepsilon = \frac{\text{number of accepted flips}}{\gamma|E|}, \quad (19)$$

where $|E|$ indicates the cardinality of the set of edges E . This sequential evaluation therefore represents an inherent serial component of Algorithm 2 that limits the potential speed-up by construction.

Second, selected edges have to be independent such that we can evaluate the respective changes in the Hamiltonian in parallel. Afterwards, we sequentially evaluate the respective acceptance probabilities. The independence of edges means that they cannot be part of the same flip patch. A flip patch is the set of edges affected by a flip (see Fig. 2). This condition imposes a constraint on the set of edges for which the change in the Hamiltonian can be evaluated in parallel. In order to avoid the NP-hard problem of a deterministic pre-computation of sets of non-interfering edges, we propose a randomized approach that is outlined in the remainder of this section.

Inspired by the batch-parallel evaluation of mesh properties of Shang *et al.*⁴², we adopt a batch-parallel version of the flip-sweep Algorithm 3. Edges are first linearly assigned to participating threads by chunking the vector of edges as managed by the underlying data structure. During the sweep-iteration, edges are then selected at random by each thread from its respective chunk. To ensure independence of such a parallel batch of edges, each edge is attempted to be locked together with the associated flip-patch for the use by other threads. For technical reasons, we impose this strict condition of non-overlapping flip-patches within a batch of edges even though non-inclusion of a flipped edge in another patch would already suffice. The implementation draws upon a scoped data locking mechanism that allows for an efficient detection of patch clashes. Upon success of locking, a thread can independently evaluate the patch-local contributions to the geometric properties M (Eq. (6)), A (Eq. (9)) and V (Eq. (10)), as well as the contributions to the bending energy (Eq. (3)), the




FIG. 2. Extent of an exemplary flip patch. Shown is a patch of vertices involved in a flip which is locked during a the flip move. In a flip move the edge (i, j) (light red, solid) is flipped to (k, l) (light red, dashed). All edges involved in the computation of the averaged vertex properties (dark red) for the vertices i, j, k and l are included in the flip patch and must be locked during the flip operation. Note, for comparison, that this patch is significantly larger than the patch required for a flip subject to the Delaunay criterion⁴¹ where the first shell neighbors are not required.

tether-penalty (Eq. (14)) and repulsion-potential (Eq. (16)). The computationally intense evaluations are thus parallelized over a batch of edges with batch-size equal to the number of threads used in parallel. Subsequently, the change in the Hamiltonian and the acceptance probability for each edge in the batch are evaluated sequentially, thus preserving ergodicity of the resulting Markov chain. Although this inherent serial component reduces the theoretically achievable speedup as mentioned above, it has shown to be satisfactory in practice while maintaining a flip rate ε close to the serial version (see section III D for detailed results).

D. Parallel performance

We analyze the parallel performance of the algorithms and methods introduced in Sec. III on a unit sphere with different degrees of mesh discretization as a test geometry. We use the icosahedron recursion technique from the meshzoo Python package⁴³ to create high quality meshes with $N_V \in [624, 2562, 10242, 40962, 163842, 655362]$ vertices. The results referred to in this section are produced on a dual socket node with Intel® Xeon® Platinum 8280 CPUs with 56 cores in total.

Figure 3 shows results for the parallel scaling of the different algorithmic components with the number of threads for different numbers of vertices: energy evaluation (Eq. (2)), its gradient/force, flip-sweep according to Algorithm 3 and a full step of the MC-procedure outlined above, comprised of 1 HMC-step (Algorithm 1 with $L = 100$ and $\Delta t = 1.0 \times 10^{-4}$; note that the results in Fig. 3 and Fig. 4 are invariant with respect to the exact value of the time step Δt) plus 1 flip-sweep. To improve consistency, the measurements for energy, gradient and flip-sweep comprise 10 evaluations of the respective components due to their short run-time.

The high arithmetic intensity involved in the energy and gradient evaluations leads to efficient use of the available resources, which is reflected in good parallel speedup on the compute node. The flip-move sweep and MC-step evalua-

Algorithm 3 Parallel version of algorithm 2. The directives **# BARRIER** and **# CRITICAL** refer to OpenMP directives used in shared memory parallelism⁴⁰. **# BARRIER** indicates a barrier where threads have to wait for each other. **# CRITICAL** defines a region that can only be executed by one thread at a time.

```

function FLIP SWEEP( $F, E, \gamma$ )
     $h_n \leftarrow H(F)$ 
    # PARALLEL
     $E_j \leftarrow \text{assign\_edges\_to\_thread}(E, j, \gamma)$ 
     $E_j \leftarrow \text{shuffle\_edges}(E_j)$ 
    for  $i$  in range( $|E_j|$ ) do
        # BARRIER
        locks  $\leftarrow$  lock_patch( $e_i$ )
        # BARRIER
        if locks.empty() then
            continue
        end if
        release_locks(locks)
        flip_edge( $e_i$ )
         $h_{n+1} \leftarrow \text{update\_hamiltonian}(h_n, F, e_i)$ 
        # CRITICAL {
             $\alpha \leftarrow \min(1, \exp(h_n - h_{n+1}))$ 
             $\sigma \sim \mathcal{U}(0, 1)$   $\triangleright$  evaluate Metropolis criterion
            if  $\sigma \leq \alpha$  then
                 $h_n \leftarrow h_{n+1}$ 
            else
                flip_edge( $e_i$ )  $\triangleright$  flip back in case of rejection
            end if
        }  $\triangleright$  end critical section
    end for
end function

```

tions also benefit from an increase in parallel execution performance, but a saturation trend is visible. This trend is a result of both of these components containing a significant amount of low arithmetic intensity workload. Thus, they are more exposed to the bound of the memory bandwidth than the energy and gradient evaluation. In addition, the flip-sweep still has a serial portion that will inherently limit the achievable speedup. However, in the experiments presented here, this does not manifest itself as the critical component controlling the effective speedup. The scaling of the full MC-step is only minimally affected by the flip-sweep due to the small contribution of the flips-sweep to a full MC-step. Instead it is governed by the amount of low arithmetic intensity workload in the HMC-step (Algorithm 1).

Importantly, with our parallelized scheme, the computation of a mesh with 40962 vertices is comparable to a single-core run with a small system with 642 vertices. Without an explicit limit on mesh size in openmesh, it is possible in principle to simulate meshes at least up to a size of $10^6 - 10^7$ vertices (Fig. 3), as required for cellular scale membranes. The global vertex moves in HMC ensures efficient sampling of the high dimensional space even for large mesh sizes.

The effective time complexity of the algorithmic components is shown in Fig. 4. The distance calculations necessary

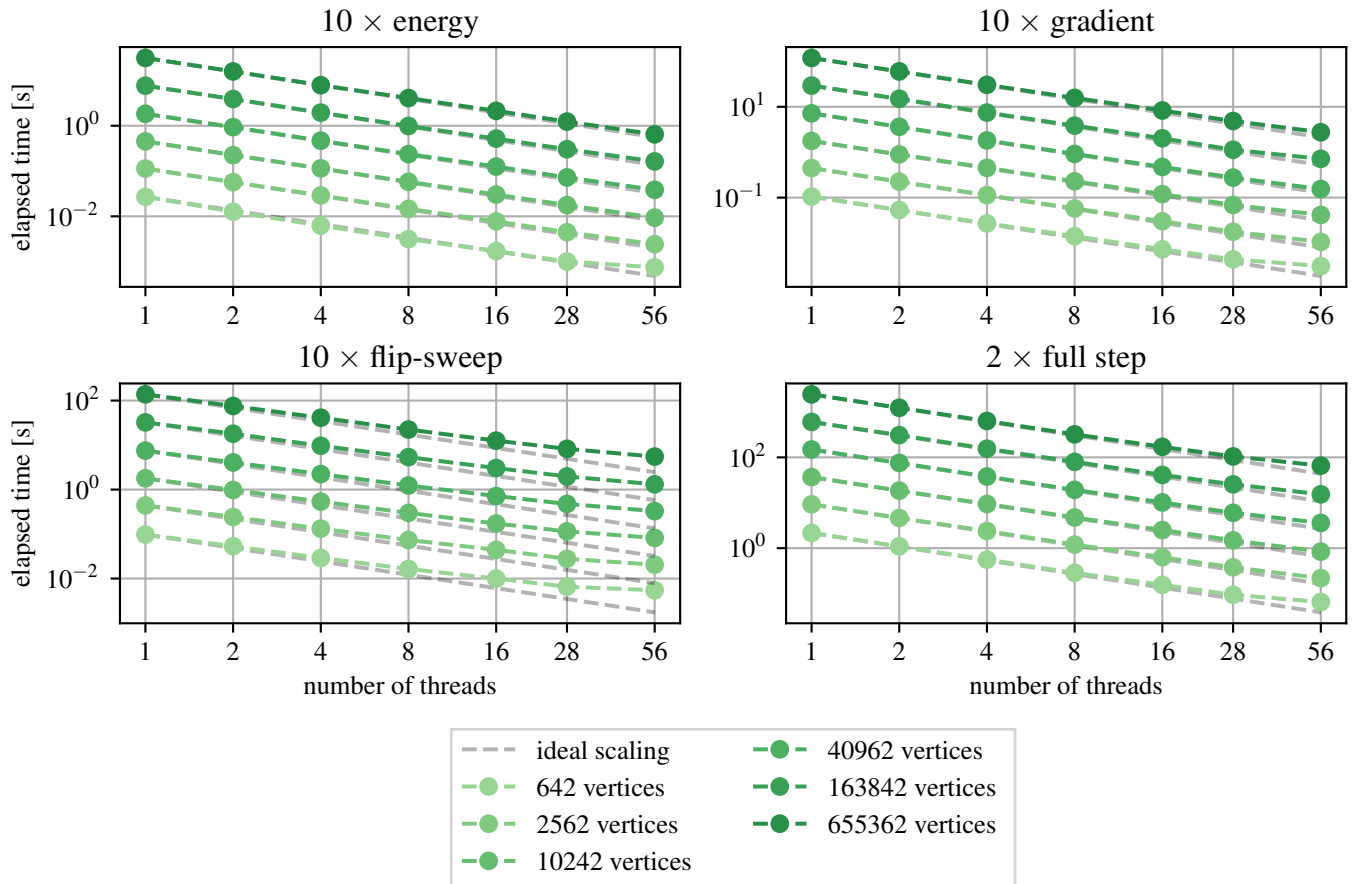


FIG. 3. Strong scaling of different algorithmic components with the number of threads for different problem sizes N_V (shades of green). The components are the evaluation of the energy, Eq. (2), its associated gradient/force, a flip-sweep according to Algorithm 3, and a full step of the MC-procedure outlined in Sec. III with 1 HMC-step (with $L = 100$) and 1 flip sweep. For the components energy, gradient, and flip-sweep, measurements consist of 10 evaluations each. The measurement for the component full-step consists of 2 steps.

for the computation of the repulsion penalty given by Eq. 16 are effectively reduced to $\mathcal{O}(N_V)$ by the neighbor list data structures. This complexity on the level of the energy and gradient evaluation is directly passed on to the evaluation of the full MC-step.

The influence of the parallel implementation (algorithm 3) on the flip rate ε of the flip-sweep is shown in Fig. 5. The measurements are carried out for a mesh with $N_V = 10242$ vertices and a target flip rate $\gamma = 10\%$. We found that the mean efficiency of $\varepsilon \approx 0.17\%$ observed for the serial implementation reduces to $\approx 0.15\%$ when using a full node with 56 cores. This slight decay is due to our randomized approach, in which the number of clashes in the locking of the necessary flip patches increases with the number of parallel threads used. These clashes lead to threads skipping a flip-batch and can consequently result in a reduced number of edge flips in total. Due to the linear distribution of edges to threads, the spatial order of edges as reflected in the data structure in memory can have an influence on the probability of clashes. As applied by⁴², a spatial pre-ordering of the edges (e.g., by space filling curves) might thus improve the efficiency for high thread counts. Since the spatial ordering

of vertices resulting from the icosahedron reconstruction used here is already rather good, preliminary tests have not shown to improve the presented results. Nevertheless, for general meshes the influence of spatial ordering is considered to have significant influence. Therefore, this topic might be followed up in future developments. In any case, the obtained effective flip rates have proven sufficient to provide the desired effect on the mesh mutability that is necessary to achieve membrane fluidity and accurate results, cf. section IV.

E. Parameter Continuation for Stiff Restraints on Membrane Shape

The accurate representation of the constraints regarding volume and surface area requires large values for the factors κ_V , κ_A and $\kappa_{\Delta A_0}$ in the respective penalty formulations given by Eqs. (7), (8), and (12). The conditioning of the Hamiltonian in Eq. (2) is determined by these penalty terms. That is, slight deviations of the initial configuration from the desired target configuration in terms of the area A and the volume V can have a large impact on the numerical stabil-

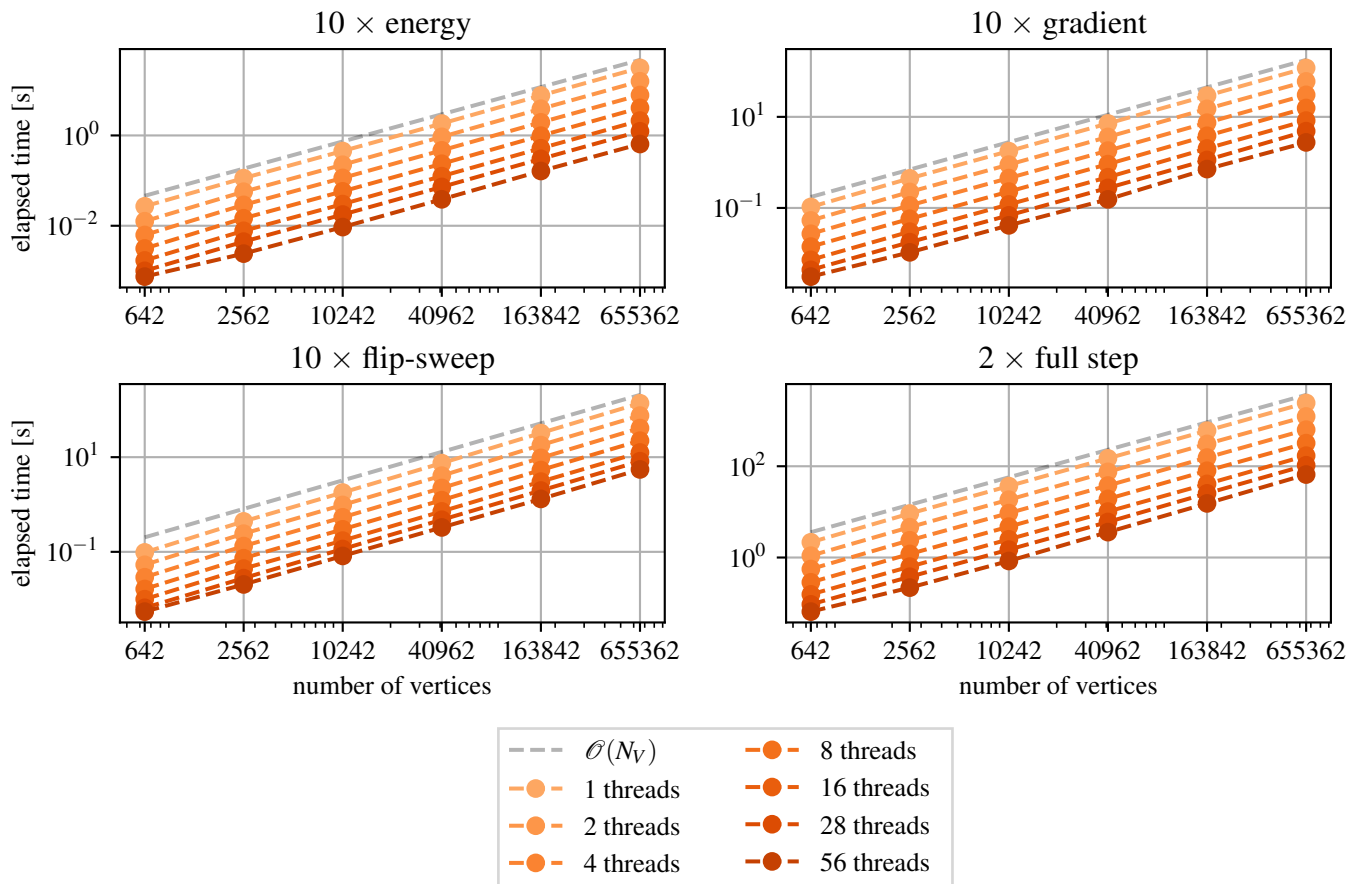


FIG. 4. Time complexity of different algorithmic components with the number of vertices N_V for different numbers of threads (shades of orange). The components are according to Fig. 3. For the components energy, gradient and flip-sweep, measurements consist of 10 evaluations each. The measurement for the component full-step consists of 2 steps. Color intensity indicates higher thread numbers. The slope of the grey dashed line indicates $\mathcal{O}(N_V)$ linear time.

ity of the Hamiltonian and the sampling performance. To mitigate such a performance degradation, TriMem offers the possibility to use a technique from the concept of parameter continuation⁴⁴. To this end, we introduce the parameterized Hamiltonian $H_p(\mathbf{x}, F; \eta)$, with the explicit definition of parameters $\eta := (V_0, A_0, \Delta A_0)$. This enables a smooth blending of the Hamiltonian $H_p(\mathbf{x}, F; \eta_0) \rightarrow H_p(\mathbf{x}, F; \eta_1)$ from some initial parameters η_0 to the target parameters η_1 via the linear interpolation

$$\eta = \lambda \eta_1 + (1 - \lambda) \eta_0 \quad (20)$$

by varying an interpolation parameter $\lambda \in [0, 1]$ smoothly from $0 \rightarrow 1$. By doing so, the sampling efficiency remains higher, even in situations in which the initial configuration is not consistent with the specific constraints in place. Parameter continuation can also help to overcome energy barriers that might appear in an immediate instantiation of the target Hamiltonian. More generally, the parameterized Hamiltonian allows us to integrate systematic approaches to branch tracking⁴⁴ in future work.

Currently, we interpolate the whole set of parameters η simultaneously by defining and implementing a single interpo-

lation parameter λ . The efficiency of this scheme could be improved by using a vector of interpolation parameters such that individual components of the Hamiltonian are transformed independently.

F. Gradient-based energy minimization

If the initial configuration is not consistent with the imposed constraints, we can improve the sampling efficiency by initial energy minimization. Such a minimization can be interpreted as the one-time application of a preconditioning and will bring the initial configuration closer to the equilibrium configuration, determined by Eq.(17). Energy minimization also avoids the necessity of running long simulations with parameter interpolation in the beginning. Since this is only an initial energy minimization prior to a HMC simulation, no particular requirements must be imposed on the convergence to the global optimum. By ignoring edge flips during minimization, we can use well established and efficient methods for function optimization such as the L-BFGS method⁴⁵. In TriMem, simulations can be run in minimization or HMC mode (see Fig.

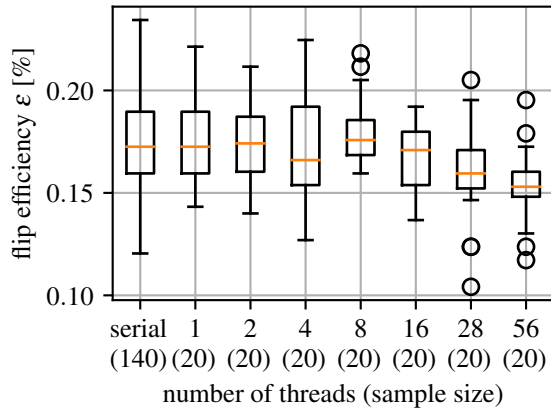


FIG. 5. Flip rate ε given by Eq. (19) of the parallel flip implementation given the number of threads. The target rate is $\gamma = 10\%$ and the number of vertices is $N_V = 10242$. For comparison, we show the rate for the serial version (leftmost data). The cardinality of the sample size is given in parentheses for every box. The flip efficiency decreases slightly with increased number of threads due to the locking of flip patches from different threads.

1).

G. Simulated Annealing

Gradient-based Energy minimization of the vertex positions must be followed by a global minimization strategy accounting for both vertex positions as well as mesh connectivity. We apply a simulated annealing procedure for the exploration of the domain of Eq. (17) that is capable of finding global minima/maxima. Following Ref.⁴⁶, we implement this method by simply modifying the temperature argument of the HMC step in Algorithm 1 according to a cooling schedule. Specifically, we apply an exponential cooling scheme $T_{n+1} = \max[\exp(-\lambda)T_n, T_{\min}]$, with the cooling factor $\lambda \in \mathbb{R}^+$ controlling the rate of cooling.

In TriMem, cooling can be initiated during HMC simulations directly in the beginning or after a longer equilibration period prior to the cooling. This enables sufficient sampling of the configurational space before settling into the global (or deep local) energy minimum.

IV. VALIDATION METHODS

We tested the robustness of the TriMem software and compared it to analytical and numerical calculations from the literature. We reproduced several well established aspects of the phase diagram for closed vesicles with $c_0 = 0$ and with respect to varying volume or area difference. In all simulations, we used a bending rigidity of $\kappa_B = 30k_B T$, which is a typical value for biological membranes. The coupling constants of the volume, κ_V , the area, κ_A , and the area difference, $\kappa_{\Delta A}$, were chosen several orders of magnitude larger

with $\kappa_V = \kappa_A = \kappa_{\Delta A} = 1 \times 10^6 k_B T$ to impose a strong restraint. All simulations were performed using a mesh size of $N_V = 1962$ or 642 vertices starting from a sphere. The initial shapes were generated with the meshzoo library⁴³. Our results are not affected by mesh size.

A practical instruction to setup and run such simulations with TriMem is available via Github (<https://github.com/bio-phys/trimem>) on the documentation webpage.

1. Volume Phase Diagram

The individual configurations of the branches of the volume phase diagram (Fig. 6) were generated as follows. The initial shapes for each branch were generated using the minimization procedure to initialize the system. Prolate simulations were started at the reduced volume $v = 3V/(4\pi R^3) = 1.0$ (Fig 6; blue triangles), oblate simulations at volume $v = 0.65$ (orange circles) and stomatocyte simulations at $v = 0.3$ (green squares). These initial shapes and volumes on the respective branches were achieved by using the preconditioning procedure (see Section III F). Then, using these initial shapes the reference reduced volume v was lowered or increased by 0.025 instantaneously in each step. In each step the previous final structure was used as initial configuration for the subsequent step. By doing so the respective branches could be mapped out by exploiting hysteresis. In all cases the simulations could, in principle, switch their respective branches on the phase diagram and otherwise no special restraint were applied. Each simulation was run for 5×10^5 steps using a temperature $T = 1k_B T$ and in an additional 6×10^5 steps the temperature was reduced from 1 to 0 following section III G. The rescaled bending energy of the lowest energy configuration was plotted for each value. In all HMC simulations an integration step size of 7^{-5} was used. The trajectories in each HMC step were 100 steps long. The flip ratio was set to $\gamma = 0.1$ of the vertex move steps. Previous tests showed this flipping ratio was sufficient for fast vertex diffusion.

2. Area-Difference Phase Diagram

The area-difference phase diagram is another commonly used benchmark to test the robustness of numerical membrane bending simulation codes^{23,28,47}. The simulations were performed as follows. The initial shape for the HMC run was generated in a two-step minimization procedure. First, the volume was reduced to the respective target reduced volumes of $v = \{0.5, 0.55, 0.6\}$, and second, the reduced area difference $\Delta a = \Delta A/(4\pi R)$ was adjusted to the respective target values between 1 and 1.8, by using the L-BFGS minimizer for gradient based energy minimization for both parameters successively. Then a HMC simulation was run for 1.5×10^7 steps. After 1.4×10^7 steps 1×10^6 steps of cooling were performed where the temperature was reduced from 1 to 0 according to section III G. The bending energy and shape at the final step of each simulation were extracted and normalized to the reduced energy by $8\pi\kappa_B$ to construct the phase diagram. Otherwise




FIG. 6. Phase diagram for a vesicle with $N_v = 1962$ vertices of spherical topology in the plane of the reduced volume ν and the bending energy E_B . Shown are minimum energies obtained for different membrane shapes as a function of ν . Filled symbols indicate the respective lowest-energy shapes. Open symbols indicate metastable shapes at a given reduced volume. All states with prolate shapes are shown as blue triangles. Oblate shapes are shown as orange circles. Stomatocytes are shown as green rectangles. The dot-dashed and dashed lines correspond to $\nu = 0.59$ and 0.56 , respectively, and mark transitions between stable branches. Exemplary energy-minimized shapes for all corresponding branches are shown in Fig. S1.

the same parameters were used for the HMC simulations as described above in section IV 1.

V. RESULTS AND DISCUSSION

We extensively validated our software for varying mesh sizes and exhaustively tested all features by reproducing known phase diagrams of vesicle shape^{47,48}. First, we ran a range of simulations to produce the volume phase diagram for a vesicle shown in Fig. 6. In the following, we consistently use reduced volumes ν as an order parameter. Starting with a sphere ($\nu = 1$), we explored the range between $\nu = 1$ and 0.2 to locate minimal shapes found as described in section IV. The y-axis shows the bending energy E_B and was re-scaled with $8\pi\kappa_B$, the bending energy of a sphere, as in previous work^{23,47}.

Here, we explored the stability of the three well established branches in the vesicle volume phase diagram; the prolate, the oblate and the stomatocyte^{17,28}. The transitions between the three branches are not smooth, i.e., the first derivative of the energies with respect to the reduced volume are discontinuous. Stable and metastable states representing the various shapes are separated by energy barriers.

The prolate branch (Fig. 6, blue triangles) gives the global energy minima for $\nu \gtrsim 0.65$. As the reduced volume is lowered, the vesicle changes its shape from the sphere at $\nu = 1$

to a prolate shapes and then to extended tubular structures. Low reduced volumes result in long metastable tubes with narrow diameter and a far higher bending energy than the corresponding oblate or stomatocyte shapes with the same reduced volume, which constitutes the global minimum for $\nu \lesssim 0.65$. Below this reduced volume the branch is metastable as previously outlined in^{23,28}. All prolate shapes and tubes below $\nu = 0.65$ are only found when starting from an initial prolate configuration and mapping out the branch by hysteresis, indicating that they are local minima and metastable.

The oblate branch is the global minimum between $0.59 \lesssim \nu \lesssim 0.65$ (Fig. 6 orange spheres). Simulations starting from the sphere instantaneously reducing the volume and using the L-BFGS minimizer to a target value $\nu < 0.65$ always converged to an oblate shape. This result is in agreement with observations of previous work^{17,28}. Above $\nu = 0.725$, the oblate is always unstable and the local energy minimum for the branch vanishes, and all initial shapes (sphere, prolate, oblate) converged to the prolate shape, which is consistent with previous theoretical work⁴⁹. Oblates in the range of $0.65 < \nu \lesssim 0.725$ can only be seen when starting with an oblate configuration in the globally stable range and tracing out the branch by hysteresis. By doing so the structures remaining in the metastable minimum. Similarly, the oblate shape is metastable at a reduced volume of $\nu \lesssim 0.59$.

For a reduced volume of $\nu \lesssim 0.59$ the stomatocyte is the global minimum, which is also achieved from the preconditioning step with the L-BFGS minimizer from a sphere using a stiff restraint on the target volume (Fig 6 green squares). For $\nu > 0.59$, we find that the stomatocyte shapes are metastable with respect to oblates. The energy of the stomatocytes are nearly independent of the reduced volume with $E_B \approx 16\pi\kappa_B$. Overall, all energies for the globally energy-minimized shapes are in good agreement with previous work following similar approaches^{17,28}.

To test if the area-difference is correctly evaluated, we also produced the phase diagram for varying area differences at multiple fixed reduced volumes. Fig. 7 shows the rescaled bending energy as a function of the area difference Δa_0 . Simulations were run for area-difference increments of 0.025 and the energy-minimized structures after cooling to $T = 0$ were plotted (see IV).

We verified that both the minimized energies and the corresponding shapes are in agreement with previous simulations and theoretical calculations^{23,28,47}. Small differences in the bending energies for some shapes in Fig. 7 are likely caused by hysteresis effects around boundaries in shape space, the use of edge flips here, and different resolutions of the triangulations. The oblate configuration and the tube shapes are metastable for $\nu = 0.55$ and are separated by an energy barrier. The top of the barrier corresponds to the non-axisymmetric paddle shape which separates the prolate and tube branch and converts to a tube with increasing Δa . For all values with $\Delta a < 1.4$ we observed structures with a D_{3h} symmetry. First triangular shaped oblates emerge which transition into three-armed starfish shapes (Fig. S2). We also find that the dumb-ell/tube structure⁴⁷ is a local minimum for $\Delta a_0 > 1.4$. Tubes are found for all higher values of Δa_0 . We note that extensive




FIG. 7. Energy as function of reduced area difference Δa at constant volume ν for a mesh with $N_V = 642$. The simulation results are shown as empty circles for a constant volume $\nu = 0.5$ (blue), 0.55 (green), and 0.6 (red), respectively. For reference, the light lines are adapted from earlier calculations⁴⁷ using numerical triangulation⁵⁰. Exemplary energy-minimized shapes for varying area difference are shown in Fig. S2

descriptions and discussions of these phase diagrams and their symmetries are provided in previous systematic studies^{23,28,47}.

Taken together, these simulations validate that TriMem faithfully reproduces data from previous studies and can be used to study a broad range of reshaping processes. All parameters are correctly assessed and our parallel dynamics protocol produces the expected results.

VI. CONCLUSIONS AND OUTLOOK

We presented TriMem, a parallelized open-source software package for hybrid MC simulations of triangulated meshes representing lipid membranes. TriMem offers a robust, efficient, and parallel implementation of the method that enables users to perform simulations efficiently with large mesh sizes of $10^6 - 10^7$ vertices, which was not previously possible. Many previously available software packages were limited to a few thousand vertices by the computational cost. We demonstrated how all relevant move types can be efficiently parallelized across multiple CPU cores up to a full two-socket shared memory compute-node (Fig. 3). While parallelization of energy and force routines to generate short MD trajectories for vertex moves is straightforward, we showed that flip parallelization requires a more complex algorithm to produce correct results. We introduced a novel method for flip moves that combines an effective generation of sets of edges with non-interfering flip-patches. The parallel computation of the independent energy changes and subsequent evaluation of the acceptance in series increases the performance. We bypass a major issue resulting from serial energy evaluation required for individual flip moves.

Additionally, the HMC approach allows for more efficient global vertex displacement moves with acceptance ratios of nearly one due to efficient proposals generated by short

molecular dynamics trajectories of the mesh and vertices¹⁸. This leads to fast convergence of sampling and will be particularly useful when simulating large meshes with millions of vertices.

We implemented an efficient neighbor-list algorithm to enable the use of repulsive, and possibly adhesive potentials, between non-neighbor vertices. The resulting code scales linearly, $\mathcal{O}(N_V)$, with mesh size, and makes calculations with large meshes possible. This feature proved particularly important to sample highly curved membrane shapes such as stomatocytes, where self-intersection of the mesh could result in faulty representations. It can also be further expanded to model the general physical properties of membrane repulsion due to charges, possible protein-mediated adhesion, and systems with multiple membrane meshes.

The user interface of the software was designed for versatility but with ease-of-use in mind. Some functionality, such as the configuration file, is kept loosely in the style of well established molecular simulation packages such as GROMACS⁵¹, which is widely used in the biophysics community. The ease-of-use and open-source nature of the code should empower a broad community of biophysicists to quickly set up, run, and evaluate large-scale membrane remodeling processes of interest. The Python front-end should also enable other users to quickly add novel features to the code.

We validated the software by reproducing several phase diagrams and results for vesicle shaping, which have been previously explored in various software packages and foot on theoretical calculations^{17,23,28,47,49}. We found in all cases that the validation was in good agreement with existing work, indicating the robustness of our algorithms.

The code was designed for performance with C++ in the backend and for ease-of-use with a Python interface in the frontend. The framework provided by the TriMem package can be extended in the future in a straightforward way to incorporate more complex biological assemblies of membranes. Protein membrane interactions, multiple membrane meshes, and cytoskeletal attachments will likely be added to the features of TriMem, enabling a broader range of cell-scale simulations. Systems containing multiple membranes in close vicinity are needed to build organelles and cell-scale membrane structures. Additionally, protein vertices and lipid domains can be included to study complex protein-induced membrane remodeling processes^{7,19}.

We anticipate that TriMem, with some modifications, will also enable simulations of complex membrane topologies. One possible direction is to include periodic boundary conditions (PBC) when defining the topology through the local connectivity of the graph. An implementation of PBC will enable the preparation of quasi-infinite membrane shapes from tubes (PBC in z direction) over flat membranes (PBC in the x - y plane) to lipidic mesophases (PBC in all directions x , y , z). Another possible direction is to simulate multiple disconnected membranes, such as one vesicle enclosing another, or stacks of membranes as they occur in the Golgi²⁵. Here, the vertex repulsion can ensure that the different membranes do not interpenetrate. It is also possible to fix the spatial position of certain vertices. In this way, the dynamic membrane can

be connected at fixed seams to other shapes, e.g., to form the half-toroidal pores connected with two flat background membranes as seen in nuclear pore complexes¹².

In the future, we envision that TriMem in combination with modern structural biology techniques such as cryo-ET will shed light on the structure and dynamics of membranes on an organelle and eventually cell-scale level. Cryo-ET will increasingly serve as data source to build accurate computational models for coarse-grained simulations. We envision a combination of multi-scale simulation techniques including TriMem's parallelized capability of performing large-scale simulations required to gain mechanistic understanding of cell-scale membrane remodeling processes.

VII. SOFTWARE AVAILABILITY

TriMem is available as free software under a GPL-3 license from (<https://github.com/bio-phys/trimem>) and can be installed as a Python package using *pip*⁵². As such, it can be used as a Python library with predefined building blocks for the methods and algorithms presented in the previous sections. It also provides a command line interface that encapsulates the algorithmic flow shown in Fig. 1, which is controlled via an input configuration file. For a more detailed description on the usage, we refer the reader to the documentation which can be found via the Github link above. Software contributions from third party developers are welcome as pull requests on Github.

SUPPLEMENTARY MATERIAL

See the supplementary material for Figs. S1 and S2 showing minimum-energy vesicle shapes.

ACKNOWLEDGMENTS

This work was supported by the Max Planck Society. G.H. and M.S. acknowledge funding by the Landes-Offensive zur Entwicklung Wissenschaftlich-Ökonomischer Exzellenz LOEWE of the State of Hesse (<https://wissenschaft.hessen.de/wissenschaft/landesprogramm-loewe>). M.S. acknowledges support by the EMBL Interdisciplinary Postdoc Programme under Marie Curie COFUND actions.

Appendix A: Gradient of the Hamiltonian

The gradient of the Hamiltonian (2) can be derived in terms of the derivatives of quantities defined on the n -simplexes that define the triangulation \mathcal{T} . For the derivatives of these quantities with respect to the vertex positions x_i , we apply efficient and well-known formulas from discrete geometry⁵³. In particular, we need the derivatives of the edge length, the face area, the face volume and the dihedral angle with respect to the involved vertex coordinates.

For the length $r_{ij} = \|u\| = \|x_j - x_i\|$ of the directed edge (x_i, x_j) , the gradient is given by

$$\frac{dr_{ij}}{dx_i} = -\frac{u}{\|u\|} \quad (\text{A1})$$

$$\frac{dr_{ij}}{dx_j} = \frac{u}{\|u\|}. \quad (\text{A2})$$

The gradient of the area A of an oriented triangle $(i, j, k) \in F$ with normal n is given by

$$\frac{dA}{dx_i} = \frac{1}{2}n \times (x_k - x_j), \quad (\text{A3})$$

$$\frac{dA}{dx_j} = \frac{1}{2}n \times (x_i - x_k), \quad (\text{A4})$$

$$\frac{dA}{dx_k} = \frac{1}{2}n \times (x_j - x_i). \quad (\text{A5})$$

The gradient of the volume V associated to an oriented triangle $(i, j, k) \in F$ (computed as the volume of the tetrahedron $(x_i, x_j, x_k, \mathcal{O})$, \mathcal{O} being the origin, is given by

$$\frac{dV}{dx_i} = \frac{1}{6}(x_j \times x_k), \quad (\text{A6})$$

$$\frac{dV}{dx_j} = \frac{1}{6}(x_k \times x_i), \quad (\text{A7})$$

$$\frac{dV}{dx_k} = \frac{1}{6}(x_i \times x_j). \quad (\text{A8})$$

And the gradient of the dihedral angle ϕ_{ij} between the normals of the two oriented triangles $(i, j, k) \in F$ and $(i, j, l) \in F$, with common edge (i, j) and normals n_1 and n_2 , is given by

$$\frac{d\phi_{ij}}{dx_i} = (\cot(\alpha_3)n_1 + \cot(\alpha_4)n_2)/r_{ij}, \quad (\text{A9})$$

$$\frac{d\phi_{ij}}{dx_j} = (\cot(\alpha_1)n_1 + \cot(\alpha_2)n_2)/r_{ij}, \quad (\text{A10})$$

$$\frac{d\phi_{ij}}{dx_k} = -(\cot(\alpha_1) + \cot(\alpha_3))n_1/r_{ij}, \quad (\text{A11})$$

$$\frac{d\phi_{ij}}{dx_l} = -(\cot(\alpha_2) + \cot(\alpha_4))n_2/r_{ij}, \quad (\text{A12})$$

whereby the α_i are the sector angles defined by

$$\alpha_1 := \angle x_k x_i x_j, \quad (\text{A13})$$

$$\alpha_2 := \angle x_j x_i x_l, \quad (\text{A14})$$

$$\alpha_3 := \angle x_i x_j x_k, \quad (\text{A15})$$

$$\alpha_4 := \angle x_l x_j x_i. \quad (\text{A16})$$

$$(\text{A17})$$

REFERENCES

- ¹H. T. McMahon and J. L. Gallop, *Nature* **438**, 590 (2005).
- ²H. T. McMahon and E. Boucrot, *J. Cell Sci.* **128**, 1065 (2015).

- ³J. Mahamid, S. Pfeffer, M. Schaffer, E. Villa, R. Danev, L. Kuhn Cuellar, F. Forster, A. A. Hyman, J. M. Plitzko, and W. Baumeister, *Science* (80-.), **351**, 969 (2016).
- ⁴S. J. Marrink, V. Corradi, P. C. Souza, H. I. Ingólfsson, D. P. Tieleman, and M. S. Sansom, *Chem. Rev.* **119**, 6184 (2019).
- ⁵H. I. Ingólfsson, M. N. Melo, F. J. van Eerden, C. Arnarez, C. A. Lopez, T. A. Wassenaar, X. Periole, A. H. de Vries, D. P. Tieleman, and S. J. Marrink, *J. Am. Chem. Soc.* **136**, 14554 (2014).
- ⁶M. Chavent, A. L. Duncan, and M. S. Sansom, *Curr. Opin. Struct. Biol.* **40**, 8 (2016).
- ⁷N. Ramakrishnan, P. Sunil Kumar, and R. Radhakrishnan, *Phys. Rep.* **543**, 1 (2014).
- ⁸W. Pezeshkian, M. König, S. J. Marrink, and J. H. Ipsen, *Front. Mol. Biosci.* **6**, 1 (2019).
- ⁹W. Pezeshkian and S. J. Marrink, *Curr. Opin. Cell Biol.* **71**, 103 (2021).
- ¹⁰B. J. Reynwar, G. Illya, V. A. Harmandaris, M. M. Müller, K. Kremer, and M. Deserno, *Nature* **447**, 461 (2007).
- ¹¹M. Siggel, R. M. Bhaskara, M. K. Moesser, I. Đikić, and G. Hummer, *J. Phys. Chem. Lett.* **12**, 1926 (2021).
- ¹²S. Mosalaganti, A. Obarska-Kosinska, M. Siggel, B. Turonova, C. E. Zimmerli, K. Buczak, F. H. Schmidt, E. Margiotta, M.-T. Mackmull, W. Hagen, G. Hummer, M. Beck, and J. Kosinski, *bioRxiv* (2021).
- ¹³W. Helfrich, *Zeitschrift für Naturforsch. C* **28**, 693 (1973).
- ¹⁴W. Helfrich, *Zeitschrift für Naturforsch. C* **29**, 510 (1974).
- ¹⁵G. Gompper and D. M. Kroll, *J. Phys. Condens. Matter* **9**, 8795 (1997).
- ¹⁶J. S. Ho and A. Baumgärtner, *Phys. Rev. Lett.* **63**, 1324 (1989).
- ¹⁷F. Jülicher, *J. Phys. II* **6**, 1797 (1996).
- ¹⁸H. Noguchi and G. Gompper, *Phys. Rev. E* **72**, 011901 (2005).
- ¹⁹B. Różycki, E. Boura, J. H. Hurley, and G. Hummer, *PLoS Comput. Biol.* **8**, e1002736 (2012).
- ²⁰F. Jülicher and R. Lipowsky, *Phys. Rev. Lett.* **70**, 2964 (1993).
- ²¹A. Šarić and A. Cacciuto, *Phys. Rev. Lett.* **108**, 1 (2012).
- ²²A. H. Bahrami, R. Lipowsky, and T. R. Weikl, *Phys. Rev. Lett.* **109**, 188102 (2012).
- ²³A. H. Bahrami and G. Hummer, *ACS Nano* **11**, 9558 (2017).
- ²⁴A. H. Bahrami, M. G. Lin, X. Ren, J. H. Hurley, and G. Hummer, *PLOS Comput. Biol.* **13**, e1005817 (2017).
- ²⁵M. Tachikawa and A. Mochizuki, *Proc. Natl. Acad. Sci.* **114**, 5177 (2017).
- ²⁶W. Pezeshkian, A. G. Hansen, L. Johannes, H. Khandelia, J. C. Shillcock, P. B. S. Kumar, and J. H. Ipsen, *Soft Matter* **12**, 5164 (2016).
- ²⁷W. Pezeshkian and J. H. Ipsen, *Soft Matter* **15**, 9974 (2019).
- ²⁸X. Bian, S. Litvinov, and P. Koumoutsakos, *Comput. Methods Appl. Mech. Eng.* **359**, 112758 (2020).
- ²⁹C. Zhu, C. T. Lee, and P. Rangamani, *bioRxiv* (2021).
- ³⁰U. Seifert, *Adv. Phys.* **46**, 13 (1997).
- ³¹K. Berndl, J. Käs, R. Lipowsky, E. Sackmann, U. Seifert, J. Kas, R. Lipowsky, E. Sackmann, and U. Seifert, *Europhys. Lett.* **13**, 659 (1990).
- ³²S. Svetina and B. Žekš, *Eur. Biophys. J.* **17**, 101 (1989).
- ³³L. Miao, U. Seifert, M. Wortis, and H.-G. Döbereiner, *Phys. Rev. E* **49**, 5389 (1994).
- ³⁴L. Verlet, *Phys. Rev.* **159**, 98 (1967).
- ³⁵S. Duane, A. Kennedy, B. J. Pendleton, and D. Roweth, *Phys. Lett. B* **195**, 216 (1987).
- ³⁶M. de Berg, M. van Kreveld, M. Overmars, and O. C. Schwarzkopf, “Polygon triangulation,” in *Computational Geometry: Algorithms and Applications* (Springer Berlin Heidelberg, Berlin, Heidelberg, 2000) pp. 45–61.
- ³⁷L. Kettner, *Computational Geometry* **13**, 65 (1999).
- ³⁸M. Botsch, S. Steinberg, S. Bischoff, and L. Kobbelt, *OpenSG Symp.* (2002).
- ³⁹W. Jakob, J. Rhineland, and D. Moldovan, “pybind11 – seamless operability between c++11 and python,” (2017), <https://github.com/pybind/pybind11>.
- ⁴⁰OpenMP Architecture Review Board, “OpenMP application program interface version 5.2,” (2021), <https://www.openmp.org/wp-content/uploads/OpenMP-API-Specification-5-2.pdf>.
- ⁴¹B. Delaunay, *Bulletin de l’Académie des Sciences de l’URSS. Classe des sciences mathématiques et na*, 793 (1934).
- ⁴²M. Shang, C. Zhu, J. Chen, Z. Xiao, and Y. Zheng, *Procedia Engineering* **163**, 289 (2016), 25th International Meshing Roundtable.
- ⁴³N. Schlömer, “meshzoo,” (2022), <https://github.com/meshpro/meshzoo>.
- ⁴⁴A. G. Salinger, N. M. Bou-Rabee, E. A. Burroughs, R. P. Palowski, R. B. Lehoucq, L. Romero, and E. D. Wilkes, *Sandia National Laboratories* (2002).
- ⁴⁵J. Nocedal and S. Wright, “Quasi-newton methods,” in *Numerical Optimization* (Springer New York, New York, NY, 2006) pp. 135–163.
- ⁴⁶R. Salazar and R. Toral, *Journal of Statistical Physics* **89**, 1047 (1997).
- ⁴⁷P. Zihlerl and S. Svetina, *Europhys. Lett.* **70**, 690 (2005).
- ⁴⁸U. Seifert, K. Berndl, and R. Lipowsky, *Phys. Rev. A* **44**, 1182 (1991).
- ⁴⁹M. Jarić, U. Seifert, W. Wintz, and M. Wortis, *Phys. Rev. E* **52**, 6623 (1995).
- ⁵⁰K. A. Brakke, *Exp. Math.* **1**, 141 (1992).
- ⁵¹M. J. Abraham, T. Murtola, R. Schulz, S. Páll, J. C. Smith, B. Hess, and E. Lindahl, *SoftwareX* **1-2**, 19 (2015).
- ⁵²<https://packaging.python.org/en/latest/guides/tool-recommendations/>.
- ⁵³M. Wardetzky, M. Bergou, D. Harmon, D. Zorin, and E. Grinspun, *Computer Aided Geometric Design* **24**, 499 (2007), discrete Differential Geometry.