

Newtonian Event-Chain Monte Carlo and Collision Prediction with Polyhedral Particles

Marco Klement,[†] Sangmin Lee,^{‡,¶} Joshua A. Anderson,[‡] and Michael Engel^{*,†}

[†]*Institute for Multiscale Simulation, IZNF, Friedrich-Alexander University
Erlangen-Nürnberg, 91058 Erlangen, Germany*

[‡]*Department of Chemical Engineering, University of Michigan, Ann Arbor, MI 48109, USA*

[¶]*Department of Biochemistry, University of Washington, Seattle, WA 98195, USA*

E-mail: michael.engel@fau.de

Abstract

Polyhedral nanocrystals are building blocks for nanostructured materials that find applications in catalysis and plasmonics. Synthesis efforts and self-assembly experiments have been assisted by computer simulations that predict phase equilibria. Most current simulations employ Monte Carlo methods, which generate stochastic dynamics. Collective and correlated configuration updates are alternatives that promise higher computational efficiency and generate trajectories with realistic dynamics. One such alternative involves event-chain updates and has recently been proposed for spherical particles. In this contribution, we develop and apply event-chain Monte Carlo for hard convex polyhedra. Our simulation makes use of an improved computational geometry algorithm XenoSweep, which predicts sweep collision in a particularly simple way. We implement Newtonian event chains in the open source general-purpose particle simulation toolkit HOOMD-blue for serial and parallel simulation. The speed-up over state-of-the-art Monte Carlo is between a factor of 10 for nearly spherical polyhedra and a factor of 2 for highly aspherical polyhedra. Finally, we validate the Newtonian event-chain algorithm by applying it to a current research problem, the multi-step nucleation of two classes of hard polyhedra.

1 Introduction

Nanocrystals can be synthesized in a variety of polyhedral Wulff shapes. They self-assemble by destabilization in solution or in evaporation and sedimentation experiments.^{1–4} The resulting superlattices have applications as functional materials.⁵ Some particles have several competing superlattice candidates. Octahedra, for example, can form four different superlattice structures. Three of these are extremal in packing density or contact area between particles.^{1,6} While the structures predicted can be reproduced in experiment reliably, the ordering phenomena and pathways to the superlattices are often not well understood. Varying the types and amounts of ligand molecules affects the structures found.^{7,8} Only for special cases, predictive rules have been found. One such case are hard particles, in which shape alone is the dominating factor. Hard-particle phase diagrams^{9–12} and equations of states^{13–15} help to understand experiments and to select promising polyhedral shapes and target structures. Other studies utilize simulations to predict photonic properties.¹⁶

Simulations of anisotropic particles are often performed stochastically with Monte Carlo methods. Molecular dynamics would in principle be favored because it equilibrates configurations more efficiently.¹⁷ The reason is that density fluctuations relax with the speed of

sound in molecular dynamics, whereas such fluctuations equilibrate by diffusion in local reversible Monte Carlo algorithms. Furthermore, molecular dynamics creates realistic trajectories, which are an advantage when comparing details of trajectories with experiment. However, integrating the equations of motion for anisotropic particles is not easily possible. This is in particular the case for hard anisotropic particles, which interact via discontinuous forces. Event-driven molecular dynamics is a well-established method for hard spheres to obtain Newtonian trajectories.¹⁸ But event-driven molecular dynamics for anisotropic particles requires solving collision equations that involve rotations and thus trigonometric functions. The solution of these non-linear equations is only possible numerically by iteration with approximations¹⁹ and is necessarily slow. In Monte Carlo, the problematic collision checks are replaced by overlap checks. Importantly, overlaps can be solved analytically in a finite number of steps for many shapes, including for polyhedra.^{20,21,21,22} It is thus desirable to combine the advantages of molecular dynamics (fast equilibration and realistic dynamics) with the advantages of Monte Carlo (simple algorithm and no approximation).

Here, we develop a simulation method that combines the advantages of Monte Carlo and molecular dynamics in an algorithm that avoids approximations and iterations to the extend possible. Our starting point is event-chain Monte Carlo,^{23,24} which offers efficient equilibration and prediction of structures, especially in the recently proposed variant of Newtonian event chains.¹⁷ We generalize Newtonian event-chain Monte Carlo to hard convex polyhedra. The algorithmic bottleneck for event-chain Monte Carlo is the prediction of collisions between non-rotating polyhedra. The detail that the polyhedra are non-rotating is crucial. It means the underlying equations do not contain trigonometric functions and can be solved analytically.

Fast collision detection is a classic problem in computer graphics.^{25–27} Solutions started with the evaluation of all vertex–face and edge–edge feature combinations of the two polyhedra in

question.²⁰ Because the number of such combinations increases rapidly with the number of polyhedron vertices, a brute-force calculation of all combinations is too time consumption. The algorithm by Gilbert, Johnson, and Kerthi²¹ (GJK) was a major breakthrough. It was not only faster because it avoided checking combinations of vertices, edges, and vertices, but supported all sorts of convex objects – as long as there is a support function that returns the furthest vector of an object for any given direction. Due to its conceptual complexity, the GJK algorithm remains cumbersome to implement. Snethen²² created with Xenocollide a derived algorithm for which the state can easily be visualized at any step. Xenocollide is the starting point of the collision prediction algorithm presented in this paper. Similar to the directional contact range²⁸ calculation, we determine the directional contact distance of convex polyhedra. This is the distance by which a polyhedron can be translated without rotation up to collision, a process called ‘sweeping’. The algorithm we present here, called XenoSweep, predicts sweep collisions and is particularly simple.

The central goal of the present manuscript is method development. We first implement Newtonian event-chain Monte Carlo of polyhedral particles utilizing XenoSweep in an open source software package and test it extensively. We then confirm that the algorithm indeed has the expected efficiency advantages. An advantage of Newtonian event-chain Monte Carlo is that the partially stochastic trajectories generated resemble Newtonian dynamics significantly better than the fully stochastic trajectories of Monte Carlo with local moves only.

2 Monte Carlo Simulation

We consider a system of anisotropic particles. The system is fully specified by the position vectors $\{\mathbf{x}_i\}$ and orientation quaternions $\{\mathbf{q}_i\}$ of all particles $i = 1, \dots, N$ and their shape. The most commonly employed strategy to simulate anisotropic particles is Monte Carlo with local updates of the positions and orientations.

2.1 Local Updates

Local-update Monte Carlo (LMC) applies a sequence of trial moves to update the configuration. A randomly selected particle is either translated by a random vector of length up to d_{trans} or rotated by addition of a random quaternion²⁹ of norm up to d_{rot} and subsequent re-normalization.

We denote the probability to execute a translation move as

$$p_{\text{trans}} = \frac{\# \text{ translation trials}}{\# \text{ translation trials} + \# \text{ rotation trials}}, \quad (1)$$

where the character ‘#’ means ‘number of’. The new configuration is accepted with a Boltzmann factor. The Boltzmann factor collapses to an overlap check when simulating hard particles.

2.2 Newtonian Event Chains

Event-chain Monte Carlo (ECMC)²³ applies a sequence of collective moves in the form of chains to update the configuration. Each chain consists of a random start followed by a series of deterministic collision events. A randomly selected particle is translated in a random chain direction up to its first collision. The collision partner then takes over and continues the translation in the chain direction up to its first collision. At this point the third particle takes over and so on. ECMC terminates the chain after a predefined chain length.

ECMC only translates particles and has been applied to disks and spheres.²³ It was recently extended to particles with three- and many-particle interaction.³⁰ Soft anisotropy can be reached by creating compound particles.³¹ A hard point-edge interaction allows simulation of nearly-hard needles.³² To treat hard anisotropic particles, we mix event-chain translations with LMC rotation trial moves as illustrated in Fig. 1. The probability to execute an event chain is

$$p_{\text{chain}} = \frac{\# \text{ event chains}}{\# \text{ event chains} + \# \text{ rotation trials}}. \quad (2)$$

We recently introduced Newtonian event-chain Monte Carlo (NEC) as a variant of ECMC




Figure 1: Flowchart of Monte Carlo simulation for anisotropic particles. The algorithm combines event chains for translation moves and random trial rotation moves as in local Monte Carlo. The type of move is chosen randomly. p_{chain} is the probability to execute an event chain.

that is more efficient but still follows the correct statistics, which is guaranteed by the fact that it obeys the balance condition.¹⁷ NEC associates each particle with a velocity vector. A randomly selected particle is translated in the direction of its velocity vector up to the first collision. During this translation, time advances by $\Delta t_i = \Delta x_i / v_i$, where Δx_i is the translation distance and v_i the velocity of the translated particle. The velocities of both particles involved in the collision are updated at the collision according to the rules of elastic collisions. The chain then continues with the collision partner as in ECMC. NEC terminates the chain after a predefined chain time $t_{\text{chain}} = \sum_i \Delta t_i$.

2.3 Parameters

LMC and NEC both have three dimensionless parameters, which must be tuned for maximal efficiency of the simulation algorithms:

(1) The NEC chain time is naturally expressed in units of the mean free time $t_{\text{mf}} = \langle \Delta t_i \rangle_{\text{NEC}}$. The average is conveniently computed in NEC after equilibration. In LMC, the translation trial move distance is naturally expressed in units of the mean free path d_{mf} . We approximate the mean free path using the

mean free time and the root mean square velocity measured in NEC, $d_{\text{mf}} \approx t_{\text{mf}} \sqrt{\langle v^2 \rangle}$. We call the joint parameter the translation parameter. It is given by

$$\tau = \begin{cases} \frac{d_{\text{trans}}}{d_{\text{mf}}} & \text{for LMC,} \\ \frac{t_{\text{chain}}}{t_{\text{mf}}} & \text{for NEC.} \end{cases} \quad (3)$$

The translation parameter τ behaves differently in LMC and NEC and will be treated separately below.

(2) The rotation trial move distance d_{rot} is identical in LMC and NEC. We call this parameter the rotation parameter,

$$\rho = \begin{cases} d_{\text{rot}} & \text{for LMC,} \\ d_{\text{rot}} & \text{for NEC.} \end{cases} \quad (4)$$

(3) The probability to execute a translation move p_{trans} in LMC can be mapped to an equivalent NEC parameter using p_{chain} by observing that the number of translation moves is given by $p_{\text{chain}}(N_{\text{chain}} + 1)$, where N_{chain} is the average number of collision events per chain. We call the fraction of translation moves the move ratio parameter. It has the expression

$$\begin{aligned} \mu &= \frac{\# \text{ translation moves}}{\# \text{ translation moves} + \# \text{ rotation moves}} \\ &= \begin{cases} p_{\text{trans}} & \text{for LMC,} \\ \frac{p_{\text{chain}}(N_{\text{chain}} + 1)}{p_{\text{chain}}(N_{\text{chain}} + 1) + (1 - p_{\text{chain}})} & \text{for NEC.} \end{cases} \end{aligned} \quad (5)$$

2.4 Parallel Event Chains

ECMC can be parallelized with a cell decomposition scheme.²⁴ As cells have a large inactive volume, we use domain decomposition³³ to parallelize NEC. Collisions with the domain wall and with particles outside of the cell are treated as elastic collisions with partners of infinite mass. Such collisions lead to ergodic dynamics if the domain walls are shifted every now and then.²⁴

We argued that NEC is efficient because the trajectories it generates follow Newtonian dynamics more closely.¹⁷ However, while collision

between particles conserve momentum, collisions with domain walls or outside particles break momentum conservation and break Newtonian dynamics near the walls. The result is a decrease of the advantage of NEC over LMC. This decrease can be critical especially for small domains.

Anisotropic particles have an additional issue. The number of chains per cell between two domain wall shifts is not fixed but Bernoulli-distributed with mean proportional to $p_{\text{chain}}N_{\text{domain}}$ and variance proportional to $p_{\text{chain}}(1 - p_{\text{chain}})N_{\text{domain}}$. Here, N_{domain} is the average number of particles in the domain. Whenever domain walls are to be shifted, the simulations for all domains must wait for the slowest domain (with most chains). The variance can be reduced by using shorter chains, increasing p_{chain} , or generally simulating longer between shifting domain walls. Too short chains destroy the advantage of event chains and are not efficient.¹⁷

3 Sweep Collisions of Convex Polyhedra

Event chains require the prediction of sweep collisions. A sweep collision is the collision of two particles when one of them is translated along a given direction, the other is not translated, and neither of them is rotated. While predicting sweep collisions of spheres is trivial, no simple expression exists for anisotropic particles.

We develop an iterative algorithm for sweep collision prediction of convex polyhedra. For this purpose, we simplify and extend Snethen's XenoCollide algorithm.²² XenoCollide utilizes Minkowski portal refinement. Portals are triangles that hit the origin when translated (or swept) along a normalized ray direction $\mathbf{r}$. This means the portal ray, which is the line $\{\lambda \mathbf{r} \mid \lambda \in \mathbb{R}\}$, intersects the triangle. Our modified XenoSweep algorithm not only returns the sweep distance ℓ at collision but also the normal vector to the collision plane $\mathbf{n}$. The normal vector is not calculated by XenoCollide, but we need it to perform elastic collisions in ECMC

or NEC.

Let A and B be two convex particles. Minkowski portal refinement utilizes the Minkowski difference $C = \{\mathbf{a} - \mathbf{b} \mid \mathbf{a} \in A, \mathbf{b} \in B\}$. C contains the origin, if and only if the intersection of A and B is non-empty. Sweep collision prediction is equivalent to checking where C hits the portal ray. A further simplification is that we need to know only the support vector $S_C(\mathbf{v}) = \mathbf{x}$, which is defined as the vector $\mathbf{x} \in C$ that maximizes $\mathbf{x} \cdot \mathbf{v}$. The support vector corresponds to the point in C that is extended furthest along $\mathbf{v}$.

The XenoSweep algorithm (Fig. 2) consists of three stages, an initialization stage and two iterative stages. Both iteration stages terminate for polyhedra after a finite number of iterations bounded from above by the number of vertices. In principle, the algorithm can be made to scale significantly faster than linearly with the number of polyhedron vertices by implementing a tree-like data structure for the calculation of the support vector function. However, we do not implement such a data structure here to keep the algorithm simple. The first iterative stage searches for the existence of a portal only in the plane perpendicular to the ray direction $\mathbf{r}$. The second iterative stage gradually improves the portal towards the collision plane. Differences to XenoCollide are a modified initialization and improved exit conditions.

3.1 Initialization

We initialize the algorithm with two vectors $\mathbf{v}_1, \mathbf{v}_2 \in C$. These vectors can be chosen arbitrarily, which is what we do for $\mathbf{v}_1$ (line 1). We find it efficient to set $\mathbf{v}_2$ as the support vector in direction $\mathbf{r} \times (\mathbf{r} \times \mathbf{v}_1) = \mathbf{r}(\mathbf{r} \cdot \mathbf{v}_1) - \mathbf{v}_1(\mathbf{r} \cdot \mathbf{r})$ on the opposite side of the particle as seen along $\mathbf{r}$ towards the origin (line 2). The normal vector $\mathbf{n} = \mathbf{r} \times (\mathbf{v}_2 - \mathbf{v}_1)$ is the search direction in the next step. We guarantee that $\mathbf{n}$ points towards the origin, by swapping the initial vectors if that is not yet the case by testing the condition $\mathbf{v}_1 \cdot \mathbf{n} < 0$ (lines 3-4).

Algorithm: XenoSweep

Inputs:	S_C	support vector function of C
	$\mathbf{r}$	ray direction
	$\mathbf{v}_{AB}$	point within C
Outputs:	ℓ	sweep distance
	$\mathbf{n}$	normal to collision plane
Variables:	$\mathbf{v}_i$	vertices of the portal

— Initialization

```

1   $\mathbf{v}_1 \leftarrow \mathbf{v}_{AB}$ 
2   $\mathbf{v}_2 \leftarrow S_C(\mathbf{r}(\mathbf{r} \cdot \mathbf{v}_1) - \mathbf{v}_1(\mathbf{r} \cdot \mathbf{r}))$ 
3  if  $\mathbf{v}_1 \cdot (\mathbf{v}_2 \times \mathbf{r}) > 0$ :
4     $\text{swap}(\mathbf{v}_1, \mathbf{v}_2)$ 

```

— Portal Discovery

```

5  while True:
6     $\mathbf{n} \leftarrow (\mathbf{v}_2 - \mathbf{v}_1) \times \mathbf{r}$ 
7     $\mathbf{v}_3 \leftarrow S_C(\mathbf{n})$ 
8    if  $\mathbf{n} \cdot \mathbf{v}_3 < 0$ :
9       $\text{return } \textit{NoCollision}$ 
10   if  $\mathbf{v}_1 \cdot (\mathbf{v}_3 \times \mathbf{r}) < 0$ :
11      $\mathbf{v}_2 \leftarrow \mathbf{v}_3$ 
12   elif  $\mathbf{v}_2 \cdot (\mathbf{v}_3 \times \mathbf{r}) > 0$ :
13      $\mathbf{v}_1 \leftarrow \mathbf{v}_3$ 
14   else:
15      $\text{break}$ 

```

— Portal Refinement

```

16 while True:
17    $\mathbf{n} \leftarrow (\mathbf{v}_3 - \mathbf{v}_1) \times (\mathbf{v}_2 - \mathbf{v}_1)$ 
18   if  $\mathbf{v}_1 \cdot \mathbf{n} \geq 0$ :
19      $\text{return } \textit{Overlapping}$ 
20    $\mathbf{v}_4 \leftarrow S_C(\mathbf{n})$ 
21   if  $(\mathbf{v}_4 - \mathbf{v}_1) \cdot \mathbf{n} \approx 0$ :
22      $\ell \leftarrow -(\mathbf{v}_1 \cdot \mathbf{n}) / (\mathbf{r} \cdot \mathbf{n})$ 
23      $\text{return } \ell, \mathbf{n}$ 
24    $\mathbf{x} \leftarrow \mathbf{v}_4 \times \mathbf{r}$ 
25   if  $\mathbf{v}_1 \cdot \mathbf{x} > 0$ :
26     if  $\mathbf{v}_2 \cdot \mathbf{x} \leq 0$ :
27        $\mathbf{v}_3 \leftarrow \mathbf{v}_4$ 
28     else:
29        $\mathbf{v}_1 \leftarrow \mathbf{v}_4$ 
30   else:
31     if  $\mathbf{v}_3 \cdot \mathbf{x} > 0$ :
32        $\mathbf{v}_2 \leftarrow \mathbf{v}_4$ 
33     else:
34        $\mathbf{v}_1 \leftarrow \mathbf{v}_4$ 

```

Figure 2: Python-like pseudocode of XenoSweep that extends XenoCollide.²² The algorithm consists of the three stages: Initialization, Portal Discovery, and Portal Refinement.

3.2 Portal Discovery

The portal discovery stage searches for the existence of a portal. A portal candidate $(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3)$ is created by including as $\mathbf{v}_3$ the support vector in a direction perpendicular to $\mathbf{r}$ and towards the origin (lines 6-7). By construction, we can exclude that the origin is behind the portal in negative normal direction (‘forbidden’ region in Fig. 3). If $\mathbf{n} \cdot \mathbf{v}_3 < 0$, we found a separating axis perpendicular to the ray direction $\mathbf{r}$ and can return ‘NoCollision’ (lines 8-9).

Otherwise, we test the intersection of the portal ray with the portal candidate. Depending on the relative location of the portal ray, portal discovery continues by replacing the vector furthest away (regions ‘replace $\mathbf{v}_1$ and ‘replace $\mathbf{v}_2$ ’ in Fig. 3), or breaks and continues with the next stage if the portal candidate is found to be a portal (lines 10-15).




Figure 3: Sketch of the portal discovery stage. With $\mathbf{n} = \mathbf{r} \times (\mathbf{v}_2 - \mathbf{v}_1)$ a third vector $\mathbf{v}_3 = S_C(\mathbf{n})$ is calculated. The portal candidate $(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3)$ is shown in projection along $\mathbf{r}$. The portal ray is a point in this projection. By construction, we know this point is located above the line through $\mathbf{v}_1$ and $\mathbf{v}_2$ and falls within the portal candidate if and only if the portal candidate is a portal. Portal discovery is iterated in the algorithm until either a portal is found or it is determined that no portal exists.

3.3 Portal Refinement

At this point in the algorithm, we successfully discovered the existence of a portal $(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3)$. But other portals might be closer to the origin.

The distance of the portal from the origin along $\mathbf{r}$ is $\ell = -(\mathbf{n} \cdot \mathbf{v}_1)/(\mathbf{n} \cdot \mathbf{r})$ with the portal normal vector $\mathbf{n} = (\mathbf{v}_3 - \mathbf{v}_1) \times (\mathbf{v}_2 - \mathbf{v}_1)$. Here, we utilized that $\mathbf{r}$ is a normalized vector. The portal refinement stage searches for the portal with the smallest ℓ . By construction $\mathbf{n} \cdot \mathbf{r} \geq 0$, and we know $\ell \leq 0$ if $\mathbf{n} \cdot \mathbf{v}_1 \geq 0$, which means the origin is on the back-side of the portal. At this point we cannot distinguish between particles moving apart and intersecting particles. Because we are not interested in negative sweeps, we return ‘Overlapping’ (lines 17-19). If $\mathbf{n} \cdot \mathbf{v}_1 \geq 0$ first occurs in a later iteration, then the origin is located between the previous and the current portal. Because all portals intersect C , the origin is then also inside C and we know the particles are intersecting.

Next, we compute a new support vector $\mathbf{v}_4$ in direction of the normal vector (line 20). We skip the check for a separating axis, $\mathbf{n} \cdot \mathbf{v}_4 < 0$, that would allow an early return in case the goal was only an overlap check instead of calculating the sweep distance. If $\mathbf{v}_4$ is coplanar to the other vectors (up to numerical precision), the portal is not moving towards the origin anymore, and we return ℓ as the sweep distance (lines 21-23). Finally, the portal is updated by replacing one of its vertices by $\mathbf{v}_4$ such that the new triangle remains a portal. We can replace $\mathbf{v}_i, i \in \{1, 2, 3\}$ by $\mathbf{v}_4$ when $\mathbf{v}_{(i+1) \bmod 3} \cdot \mathbf{x} > 0$ and $\mathbf{v}_{(i+2) \bmod 3} \cdot \mathbf{x} \leq 0$, where $\mathbf{x} = \mathbf{v}_4 \times \mathbf{r}$ (lines 24-34). One example is illustrated in Fig. 4. Based on the result of a first scalar product, we can choose the second scalar product either such that one of the replacement criterion is fulfilled directly or such that both scalar products have the same sign. Because the origin ray intersects the portal, the set of scalar products $\mathbf{v}_i \cdot \mathbf{x}, i = 1, 2, 3$ contain at least one positive and one negative result, and one replacement criterion is fulfilled implicitly.

3.4 Test and validation

To validate XenoSweep, we set up a well-defined test configuration. Two octahedra with edge length 1 are displaced by 1.8 in horizontal direction and 5 in vertical direction (Fig. 5(a)). The sweep distance is calculated in vertical direc-




Figure 4: Sketch of the portal refinement stage. The portal ($\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$) is projected along $\mathbf{r}$. We illustrate an example, where the origin $\mathbf{O}$ is located such that $\mathbf{v}_4$ replaces $\mathbf{v}_1$. $\mathbf{x} = \mathbf{v}_4 \times \mathbf{r}$ is a temporary vector used in the algorithm.




Figure 5: Comparison of collision prediction with XenoSweep and with the Bullet library.³⁴ (a) Geometry for collision prediction tests. Two octahedra (shown here in projection as squares) are placed with vertical and horizontal offset (denoted with numbers in the figure, drawn not to scale). (b) The sweep distance of the octahedra is predicted accurately with both algorithms (top). Time for the calculation of 10^4 collisions on an Intel Core i5-2400 CPU (bottom). In this test, XenoSweep is noticeably faster. The definition of the rotation angle θ (x-axis) is indicated in (a).

tion as the second octahedron is rotated. For a range of rotation angles around 45 degrees, the sweep distance is infinite because the octahedra do not collide. An advantage of this setup is that the sweep collision can be calculated analytically.

Fig. 5(b) compares the accuracy and performance of XenoSweep with the class `LinearConvexCastDemo` in the Bullet library,³⁴ which detects collisions by employing GJK Ray Casting.³⁵ The Bullet library is a physics engine and is used in many video games and for visual effects in movies. Our calculations show that the evaluated sweep distance is accurate in both algorithms and agrees with the analytic solution. We also evaluate the time for 10^4 collision. In this test, XenoSweep is significantly faster with a performance that only weakly depends on the specific geometry. We find an improvement of about a factor of 4 in cases with collision, and an improvement of about a factor of 2 if there is no collision. While this single test cannot replace a systematic benchmark, it already demonstrates that XenoSweep is not only simple but also sufficiently fast for our purposes, and that it compares well with established collision detection libraries. A further, more thorough validation will be presented below where we utilize NEC for the reproduction of multi-step nucleation in two thermodynamic systems of hard polyhedra.

4 Implementation and Parameterization

We implemented NECs in the open source general-purpose particle simulation toolkit HOOMD-blue.^{33,36} HOOME-blue has a well-tested and highly-efficient hard particle Monte Carlo (HPMC) package, which we extend. From the existing HOOMD-blue codebase, we use axis-aligned bounding boxes (AABBs), AABB-trees, the general management structure for memory, communication via MPI, and file input/output. As a result of this work, the XenoSweep algorithm and NEC for spheres¹⁷ and convex polyhedron will be included in HOOMD-blue in version 3.0.0.³⁷

It remains to tune the translation parameter, the rotation parameter, and the move ratio parameter in both LMC and NEC for maximum efficiency. Three polyhedra are tested, the tetrahedron with four vertices, the octahedron with six vertices, and the icosahedron with twelve vertices (Fig. 6). Every polyhedron is scaled to have volume $V_0 = 1$. The reference systems contains 8000 particles at 45% volume fraction. At this density all convex polyhedron remain in the fluid phase. Particle velocities in NEC are initialized randomly with mean velocity $\langle \mathbf{v} \rangle = 0$ and root mean square velocity $\sqrt{\langle \mathbf{v}^2 \rangle} = 1$. We quantify the efficiency of the algorithms by measuring the diffusion coefficient in units of CPU time (seconds), D_{CPU} . Simulations are performed in single-core mode on Intel Skylake compute nodes. Unless specified otherwise in each of the following sections, we use: $\tau_{\text{LMC}} = 2$ in LMC, $\tau_{\text{NEC}} = 30$ in NEC, $\mu = 0.5$, $\rho = 0.15$ for icosahedron and octahedron; $\rho = 0.07$ for tetrahedron. The unit of length x in Fig. 7, 8 and 9 is the length of a cube with the same volume for each shape.




Figure 6: The three polyhedra used for optimizing the parameters of LMC and NEC. The colors and symbols represent the key for Fig. 7, 8 and 9.

4.1 Translation parameter

We discuss the translation parameter τ separately for LMC and NEC because it affects the efficiency of both algorithms differently. In LMC, highest diffusion is obtained for a translation move distance of about double the mean free path, i.e. $\tau = 2$ (Fig. 7(a)). The acceptance

probability of translation moves is about 20% for this choice. Diffusion eventually decreases with increasing τ because random translations over large distances generate many overlaps.

In NEC, chains should be sufficiently long ($\tau \geq 10$) but increasing their length further has no significantly detrimental effect (Fig. 7(b)). This behavior agrees with NEC of hard spheres.¹⁷ In both cases, particles that are more spherical (icosahedron; yellow color) diffuse faster than particles that are less spherical (tetrahedron; purple) with the octahedron (green) located in the middle. We also observe a clear performance advantage (higher diffusion) of NEC over LMC that is greater for the icosahedron than the tetrahedron.

4.2 Rotation parameter

The rotation parameter ρ has similar behavior in LMC and NEC (Fig. 8). At density 45%, particles still have ample space to rotate, which means optimal performance is obtained for rather large orientation changes. Just like for the translation parameter in LMC, optimal performance generally occurs near values of the acceptance probability of around 20%. An overall vertical shift of the curves can be explained by the role of translation updates, which dominate diffusion. As before, more spherical particles diffuse faster, and NEC is generally more efficient than LMC.

4.3 Move ratio parameter

The move ratio μ is the ratio of translation moves among all moves in the simulation and ranges from zero (only rotation) to one (only translation). Simulations with parameters in either end of the parameter range are inefficient because particles get easily stuck in their neighbor shell (Fig. 9)). Shape plays an important role for the optimal choice of μ . Icosahedra prefer more translations ($\mu \gtrsim 0$), which is explained by the presence of a rotator phase prior to crystallization.³⁸ Tetrahedra prefer more rotations ($\mu \lesssim 1$) due to their strong preference for face-to-face contact.³⁹ In general, we expect nearly spherical polyhedra to rotate easily and





Figure 7: Effect of the translation parameter τ on the diffusion coefficient D_{CPU} for (a) LMC and (b) NEC. Highest diffusion occurs near $\tau = 2$ with a corresponding LMC acceptance probability of about 20%. As τ increases, diffusion decreases for LMC but plateaus for NEC.




Figure 8: Effect of the rotation parameter ρ on the diffusion coefficient D_{CPU} . Highest diffusion occurs for acceptance probabilities near 20%. We expect diffusion to decrease faster along the ρ -axis in more dense systems.




Figure 9: Effect of the move ratio parameter μ on the diffusion coefficient D_{CPU} . As expected, a mix of rotation and translation gives highest diffusion. The optimal mix changes with the polyhedron shape.

highly anisotropic polyhedra to preferentially align face-to-face.

Generally, the location of the diffusion maximum shifts towards lower μ in NEC compared to LMC. This confirms once more that translation moves are more efficient in NEC. We do not see a clear choice for μ that would allow us to select a particular optimal parameter value. For this reason, we set $\mu = 0.5$ in line with common practice.

5 Performance

After extending the NEC algorithm to convex polyhedra and implementing it in HOOMD-blue, we now quantify the speed-up of NEC over LMC. For this purpose, we analyze the ratio of the diffusion coefficients, $D_{\text{CPU}}^{\text{NEC}}/D_{\text{CPU}}^{\text{LMC}}$. The simulation parameters of the algorithms are chosen as $\tau_{\text{LMC}} = 2$, $\tau_{\text{NEC}} = 30$, $\mu = 0.5$. The rotation parameter is chosen depending on particle shape to account for variations in the importance of rotations as $\rho = 0.07$ for the tetrahedron, $\rho = 0.10$ for the triangular prism, $\rho = 0.15$ for the truncated tetrahedron, the cube, and the octahedron, $\rho = 0.3$ for the elongated dodecahedron, and $\rho = 1.0$ for all others polyhedra investigated.

We put a particular focus on the role of particle shape, which we describe by the sphericity (or isoperimetric quotient) of the polyhedron. Sphericity is defined for a convex particle as $Q = 36\pi V^2 A^{-3}$ with surface area A and volume V . Besides single-core performance, we also investigate the scaling behavior for parallel LMC and parallel NEC.

5.1 Single-Core Performance

We analyze the speed-up of NEC over LMC for different polyhedra in serial (single-core) mode of HOOMD-blue. As Fig. 10 shows, NEC is significantly more efficient than LMC for all tested convex polyhedra. The speed-up is always at least a factor of 2 and can be up to a factor of 10, with the minimum taken for tetrahedra and the maximum for icosahedra.

We observe the trend that higher sphericity

generally results in higher speed-up (Fig. 10). This makes sense because, as we have seen in Fig. 9, translations are less critical for equilibrating polyhedra with low sphericity. After all, translations benefit directly from NEC whereas rotations are not or only indirectly affected by the efficiency boost. The relationship between sphericity and speed-up is nearly linear.

The number of vertices also plays a role, but a minor one. For example the pentakis dodecahedron (sphericity 0.939, 32 vertices) is sped up by a factor of 10 and the truncated pentakis dodecahedron (sphericity 0.958, 120 vertices) is sped up by a factor of 8. The speed-up is smaller for the truncated pentakis dodecahedron, which has similar sphericity but more vertices than the pentakis dodecahedron. This behavior can be explained by details of the XenoSweep algorithm. While an overlap check with XenoCollide (implemented in LMC) only finds a separating axis or a shared point, the distance measurement with XenoSweep searches for the colliding surface elements. The latter task costs more compute time and, in comparison, slows down faster with the number of polyhedron vertices. Interestingly, the NEC speed-up for the sphere¹⁷ falls right on top of the polyhedron data despite the use of conceptually different algorithms for overlap and collision checks in both cases.

5.2 Parallel Performance

To test parallel performance of NEC, we simulate octahedra varying the system size and the number of CPU cores. Parallel performance is analyzed by the particle update frequency (i.e., how fast the algorithm generates and executed translations and rotations) and the diffusion coefficient (i.e., the efficiency of the algorithm to propagate through configuration space). We normalized the measurements to the values for a single-core simulation and divide them by the number of cores.

Results are shown in Fig. 11. NEC slows down as the number of cores increases, and scales better in the particle update frequency than the diffusion coefficient. This means domain decomposition has a negative influence




Figure 10: Speed-up of NEC over LMC as quantified by the ratio of diffusion coefficients. Twelve convex polyhedra and the sphere (labelled ‘Sph’) are tested. The polyhedra are, from bottom left (number of vertices): tetrahedron (4), triangular prism (6), truncated tetrahedron (12), cube (8), octahedron (6), elongated dodecahedron (18), cuboctahedron (12), dodecahedron (20), icosahedron (12), truncated icosahedron (60), pentakis dodecahedron; all vertices of same distance from the origin (32), and truncated pentakis dodecahedron (120).

on the parallel performance of chains, an effect that we find to be particularly strong for small systems. Our observation is in line with previous research.²⁴ NEC requires correlated translations in chains over significantly larger distances than translation trial moves in LMC. This can become problematic. For good parallel efficiency of NEC beyond a small number of cores, the system should be rather large, likely contain millions of particles. Massively parallel simulations on GPUs as performed with LMC⁴⁰ are not advisable with NEC, though chains with local times⁴¹ may be an alternative.

6 Pathways of multi-step nucleation

It remains to apply the polyhedron NEC code to a current research problem. We focus on systems of truncated tetrahedra (TTs) and triangular bipyramids (TBPs). Both of these systems show interesting phase behavior and un-

usual phase transformation pathways.⁴² TTs forms a high-density fluid consisting of shared dodecahedron motifs that nucleate a cF432 crystal. TBPs form a high-density fluid in form of an amorphous network and crystallize into a Clathrate I crystal. Importantly, when simulated at intermediate volume fraction, both systems exhibit two-step crystallization. This finding was unexpected and the systems are currently the only known examples of multi-step phase transitions in a hard particle system. Our aim here is to test whether NEC can reproduce the findings previously made with the LMC implementation in the HPMC package of HOOMD-blue. This is a good test, because such phase behavior is highly sensitive to mistakes in the collision prediction (i.e. problems in the XenoSweep algorithm) and to violations of the balance condition (i.e. incorrect statistics in the NEC implementation). Our analysis utilizes the same order parameters as in previous work.⁴²

Indeed, TTs (Fig. 12) and TBPs (Fig. 13)




Figure 11: Scaling behavior for parallel NEC for systems of different numbers of particles N_{particle} . The algorithm ran on the same CPUs but utilized different numbers of cores $N_{\text{core}} \in \{1, 2, 4, 8\}$. At all times, all CPU cores were completely filled with identical compute jobs to achieve comparable load. The particle update frequency (left) and the diffusion coefficient (right) are normalized to their single-core values and divided by the number of cores. A flat line corresponds to perfect scaling.

both reproduce all details of the reported two-step crystallization pathways. This finding demonstrates that our NEC algorithm follows the correct thermodynamics. Furthermore, NEC simulations of 4000 particles reach initial crystal formation after 430 core-hours (0.11 h/particle) for TTs and 292 core-hours (0.07 h/particle) for TBPs. In contrast, previous LMC simulations required 1725 core-hours (TT) for a system of 8000 TT particles (0.22 h/particle) and 1800 core-hours for a system of 20,000 TBP particles (0.09 h/particle) for crystallization to start.⁴² This means our observation is a first indication that it is possible to convert the speed-up derived from brief trajectories and diffusion measurements into an efficient equilibration of highly complex particle systems with a noticeable (factor 2.0) speed advantage of NEC over LMC for TTs a slight (factor 1.2) speed advantage for the highly aspherical TBPs. Still, given the difference in system sizes, these numbers are difficult to compare. Furthermore, because crystallization is a stochastic process that becomes more probable with larger system size, these speed advantages might in fact be underestimations. Clearly, more simulation work is necessary for a rigorous comparison.

As an added benefit, our simulations directly address a concern related to the way how Monte Carlo trajectories approach equilibrium. By design, Monte Carlo and molecular dynamics methods must always reach the same equilibrium given sufficient equilibration time. This equivalence of algorithms is, however, not necessarily the case for trajectories: Monte Carlo trajectories are distinct from trajectories produced by molecular dynamics at the microscopic level. Molecular dynamics trajectories are preferred, in principle, because they resemble the dynamics of experiments more closely. But molecular dynamics is not easily achievable for polyhedra, because the equations that need to be solved are highly non-linear. Event-driven hard particle simulations work only for spheres or with numerical approximations¹⁹ that might again affect the trajectories. In contrast, Monte Carlo is much simpler and therefore currently preferentially applied to anisotropic particles. A priori there is no guarantee that LMC reaches equilibrium in the same way as integrating Newton's equations of motions. Is the use of Monte Carlo as a replacement for molecular dynamics problematic when analyzing details of phase transformations trajectories?

NEC does not fully reproduce Newtonian




Figure 12: Analysis of a long NEC trajectory for tetrahedra with truncated tips and edges (TTs) at volume fraction 61%. The system is known to crystallize in two steps via a prenucleation motif with clusters.⁴² Views along the (111) direction depicting (a) the polyhedra and (b) spheres at the centers of the polyhedra. Particles are colored for better visibility of the crystal, with fluid particles translucent. A high-density fluid phase (orange, left) features a dodecahedron motif with particles shared across dodecahedron (SD). The crystal phase (red, right) features dodecahedron that do not share particles (NSD). (c) Characterization of the trajectory with order parameters. See Ref. 42 for a detailed description of the phases and order parameters.




Figure 13: Analysis of a long NEC trajectory for triangular bipyramids (TBPs) at volume fraction 50%. The system is known to crystallize in two steps via a prenucleation motif with a dense network.⁴² Views along the (100) direction depicting (a) the polyhedra and (b) spheres at the centers of the polyhedra tetramers. Particles are colored for better visibility of the crystal, with fluid particles translucent. A high-density fluid phase (blue) features tetrahedral coordination. The crystal phase (red) is of type clathrate I. (c) Characterization of the trajectory with two order parameters. See Ref. 42 for a detailed description of the phases and order parameters.

dynamics but resembles it significantly better than Monte Carlo.¹⁷ In fact, NEC can reasonably be considered intermediate to Monte Carlo and molecular dynamics. Our results in Fig. 12 and 13 demonstrate that the phenomenon of two-step crystallization remain present in NEC. This indicates that the choice of simulation algorithm is, after all, less crucial. We note that these results are only a preliminary evaluation of the phenomenon. More in-depth analysis will be required in future to settle the issue more conclusively.

7 Conclusion

We developed, implemented, and tested NEC for convex hard polyhedra. The new simulation method is an improvement over existing polyhedron simulation codes. It is multiple times more efficient. As an added benefit, NEC returns particle trajectories that are closer to molecular dynamics than conventional Monte Carlo simulation. Another advantage is that the virial expression for pressure, often a helpful observable to identify phase transformations, drops out as a side product of using event chains without additional effort.^{17,43–45} In LMC, pressure must be computed separately and requires additional effort. Our method is a first step towards more general and more versatile simulation algorithms for anisotropic particles.

At the core of NEC lies the sweep collision prediction algorithm XenoSweep. This algorithm heavily builds on prior work in the computer graphics community.²² XenoSweep solves the classic problem of polyhedron collision prediction and overlap detection^{46,47} in only 34 lines of pseudocode. The problem of collision prediction is of interest to researchers working on a broad range of problems in computer graphics, robotics, and granular dynamics.

Future extensions should advance our work in two directions: improved handling of rotations and generalization to anisotropic particles with extended interaction, i.e. particles that are not purely hard. Concerning the first point: While we sped up the effect of translation moves by about one order of magnitude, rotations remain

a bottleneck. This is apparent from Fig. 10. Efficiency reduces as the importance of rotations increases, which is the case towards more aspherical particles. Future work could develop collective rotations moves that resemble conceptually the idea of event chains for translations. Such an improved algorithm has the potential to reach a speed-up of up to one order of magnitude across all convex polyhedra. Second, the hard particle condition employed exclusively in this work is a simplification and limits comparison to experiments. Whereas nanoparticle shape is often of central importance, interactions such as van der Waals forces, electrostatic forces, or the effect of ligands can typically not be entirely ignored. Developing a general simulation method for anisotropic convex or concave polyhedra with arbitrary interactions remains a challenging open problem.

Acknowledgement This work has been funded by SFB1411 and the Cluster of Excellence Engineering of Advanced Materials of the German Research Foundation (DFG). Support by the Central Institute for Scientific Computing (ZISC), the Interdisciplinary Center for Functional Particle Systems (FPS), and computational resources and support provided by the Erlangen Regional Computing Center (RRZE) are gratefully acknowledged. Algorithm implementation and optimization in HOOMD-blue supported by the National Science Foundation, Division of Materials Research Award # DMR 1808342. Work implementing the test cases and model verification was supported as part of the Center for Bio-Inspired Energy Science, an Energy Frontier Research Center funded by the U.S. Department of Energy, Office of Science, Basic Energy Sciences under Award # DE-SC0000989. This work used the Extreme Science and Engineering Discovery Environment (XSEDE),⁴⁸ which is supported by National Science Foundation grant number ACI-1548562; XSEDE award DMR 140129.

Bibliography

References

- (1) Chang, C.-C.; Wu, H.-L.; Kuo, C.-H.; Huang, M. H. Hydrothermal Synthesis of Monodispersed Octahedral Gold Nanocrystals with Five Different Size Ranges and Their Self-assembled Structures. *Chemistry of Materials* **2008**, *20*, 7570–7574.
- (2) Eguchi, M.; Mitsui, D.; Wu, H.-L.; Sato, R.; Teranishi, T. Simple Reductant Concentration-dependent Shape Control of Polyhedral Gold Nanoparticles and Their Plasmonic Properties. *Langmuir* **2012**, *28*, 9021–9026.
- (3) Henzie, J.; Grünwald, M.; Widmer-Cooper, A.; Geissler, P. L.; Yang, P. Self-assembly of Uniform Polyhedral Silver Nanocrystals into Densest Packings and Exotic Superlattices. *Nature Materials* **2012**, *11*, 131.
- (4) Avci, C.; Imaz, I.; Carné-Sánchez, A.; Pariente, J. A.; Tasios, N.; Pérez-Carvajal, J.; Alonso, M. I.; Blanco, A.; Dijkstra, M.; López, C.; Maspocho, D. Self-assembly of Polyhedral Metal-organic Framework Particles into Three-dimensional Ordered Superstructures. *Nature Chemistry* **2017**, *10*, 78–84.
- (5) Boles, M. A.; Engel, M.; Talapin, D. V. Self-assembly of Colloidal Nanocrystals: From Intricate Structures to Functional Materials. *Chemical Reviews* **2016**, *116*, 11220–11289.
- (6) Gong, J.; Newman, R. S.; Engel, M.; Zhao, M.; Bian, F.; Glotzer, S. C.; Tang, Z. Shape-dependent Ordering of Gold Nanocrystals into Large-scale Superlattices. *Nature Communications* **2017**, *8*, 14038.
- (7) Wang, S.-S.; Xing, Z.-H.; Chen, G.-Y.; Cölfen, H.; Xu, A.-W. Cuboctahedral Sb_2O_3 mesocrystals organized from octahedral building blocks: more than self-similarity. *Crystal Growth & Design* **2016**, *16*, 3613–3617.
- (8) Lu, F.; Vo, T.; Zhang, Y.; Frenkel, A.; Yager, K. G.; Kumar, S.; Gang, O. Unusual packing of soft-shelled nanocubes. *Science Advances* **2019**, *5*, eaaw2399.
- (9) Ni, R.; Gantapara, A. P.; Graaf, J. d.; Roij, R. v.; Dijkstra, M. Phase Diagram of Colloidal Hard Superballs: From Cubes Via Spheres to Octahedra. *Soft Matter* **2012**, *8*, 8826.
- (10) Damasceno, P. F.; Engel, M.; Glotzer, S. C. Crystalline Assemblies and Densest Packings of a Family of Truncated Tetrahedra and the Role of Directional Entropic Forces. *ACS Nano* **2012**, *6*, 609–614.
- (11) Damasceno, P. F.; Engel, M.; Glotzer, S. C. Predictive Self-assembly of Polyhedra into Complex Structures. *Science* **2012**, *337*, 453–457.
- (12) Gantapara, A. P.; Graaf, J. d.; Roij, R. v.; Dijkstra, M. Phase Diagram and Structural Diversity of a Family of Truncated Cubes: Degenerate Close-packed Structures and Vacancy-rich States. *Physical Review Letters* **2013**, *111*, 015501.
- (13) Marechal, M.; Löwen, H. Density Functional Theory for Hard Polyhedra. *Physical Review Letters* **2013**, *110*, 137801.
- (14) Irrgang, M. E.; Engel, M.; Schultz, A. J.; Kofke, D. A.; Glotzer, S. C. Virial Coefficients and Equations of State for Hard Polyhedron Fluids. *Langmuir* **2017**, *33*, 11788–11796.
- (15) Tian, J.; Jiang, H.; Mulero, A. Equations of the state of hard sphere fluids based on recent accurate virial coefficients B5–B12. *Phys. Chem. Chem. Phys.* **2019**, *21*, 13070–13077.

- (16) Cersonsky, R. K.; Dshemuchadse, J.; Antonaglia, J.; Anders, G. v.; Glotzer, S. C. Pressure-tunable photonic band gaps in an entropic colloidal crystal. *Physical Review Materials* **2018**, *2*, 125201.
- (17) Klement, M.; Engel, M. Efficient equilibration of hard spheres with Newtonian event chains. *The Journal of Chemical Physics* **2019**, *150*, 174108.
- (18) Alder, B.; Wainwright, T. Studies in Molecular Dynamics. II. Behavior of a Small Number of Elastic Spheres. *Journal of Chemical Physics* **1960**, *33*, 1439–1451.
- (19) Donev, A.; Torquato, S.; Stillinger, F. H. Neighbor list collision-driven molecular dynamics simulation for nonspherical hard particles. I. Algorithmic details. *Journal of Computational Physics* **2005**, *202*, 737–764.
- (20) Canny, J. Collision detection for moving polyhedra. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **1986**, 200–209.
- (21) Gilbert, E. G.; Johnson, D. W.; Keerthi, S. S. A fast procedure for computing the distance between complex objects in three-dimensional space. *IEEE Journal on Robotics and Automation* **1988**, *4*, 193–203.
- (22) Snethen, G. In *Game Programming Gems 7*; Jacobs, S., Ed.; Charles River Media, 2008; pp 165–178.
- (23) Bernard, E. P.; Krauth, W.; Wilson, D. B. Event-chain Monte Carlo Algorithms for Hard-sphere Systems. *Physical Review E* **2009**, *80*, 056704.
- (24) Kampmann, T. A.; Boltz, H.-H.; Kierfeld, J. Parallelized Event Chain Algorithm for Dense Hard Sphere and Polymer Systems. *Journal of Computational Physics* **2015**, *281*, 864–875.
- (25) Jiménez, P.; Thomas, F.; Torras, C. 3d Collision Detection: A Survey. *Computers & Graphics* **2001**, *25*, 269–285.
- (26) Brochu, T.; Edwards, E.; Bridson, R. Efficient geometrically exact continuous collision detection. *ACM Transactions on Graphics* **2012**, *31*, 1–7.
- (27) Huang, J.; Chen, J.; Xu, W.; Bao, H. A survey on fast simulation of elastic objects. *Frontiers of Computer Science* **2019**, *13*, 443–459.
- (28) Choi, Y.-K.; Li, X.; Rong, F.; Wang, W.; Cameron, S. Determining the directional contact range of two convex polyhedra. *Computer-Aided Design* **2010**, *42*, 27–35.
- (29) Frenkel, D.; Smit, B. *Understanding molecular simulation: from algorithms to applications*; Elsevier, 2001; Vol. 1.
- (30) Harland, J.; Michel, M.; Kampmann, T. A.; Kierfeld, J. Event-chain Monte Carlo algorithms for three- and many-particle interactions. *Europhysics Letters* **2017**, *117*, 30001.
- (31) Müller, D.; Kampmann, T. A.; Kierfeld, J. Chaining of hard disks in nematic needles: particle-based simulation of colloidal interactions in liquid crystals. *Scientific Reports* **2020**, *10*, 12718.
- (32) Kampmann, T. A.; Müller, D.; Weise, L. P.; Vorsmann, C. F.; Kierfeld, J. Event-Chain Monte-Carlo Simulations of Dense Soft Matter Systems. *Frontiers in Physics* **2021**, *9*.
- (33) Anderson, J. A.; Irrgang, M. E.; Glotzer, S. C. Scalable Metropolis Monte Carlo for Simulation of Hard Shapes. *Computer Physics Communications* **2016**, *204*, 21–30.
- (34) *Bullet Physics Library - Version 2.81*; <http://bulletphysics.org>.
- (35) van den Bergen, G. *Ray Casting against General Convex Objects with Application to Continuous Collision Detection*; 2004.
- (36) Anderson, J. A.; Lorenz, C. D.; Traveset, A. General Purpose Molecular Dynamics Simulations Fully Implemented on

- Graphics Processing Units. *Journal of Computational Physics* **2008**, *227*, 5342–5359.
- (37) HOOMD-blue; <http://hoomd-blue.readthedocs.io/>.
- (38) Mravlak, M.; Schilling, T. On the phase diagram of Mackay icosahedra. *The Journal of Chemical Physics* **2018**, *149*, 134502.
- (39) Haji-Akbari, A.; Engel, M.; Keys, A. S.; Zheng, X.; Petschek, R. G.; Palfy-Muhoray, P.; Glotzer, S. C. Disordered, quasicrystalline and crystalline phases of densely packed tetrahedra. *Nature* **2009**, *462*, 773–777.
- (40) Anderson, J. A.; Jankowski, E.; Grubb, T. L.; Engel, M.; Glotzer, S. C. Massively parallel Monte Carlo for many-particle simulations on GPUs. *Journal of Computational Physics* **2013**, *254*, 27–38.
- (41) Li, B.; Todo, S.; Maggs, A. C.; Krauth, W. Multithreaded event-chain Monte Carlo with local times. *Computer Physics Communications* **2021**, *261*, 107702.
- (42) Lee, S.; Teich, E. G.; Engel, M.; Glotzer, S. C. Entropic colloidal crystallization pathways via fluid–fluid transitions and multidimensional prenucleation motifs. *Proceedings of the National Academy of Sciences* **2019**, *116*, 14843–14851.
- (43) Michel, M.; Kapfer, S. C.; Krauth, W. Generalized Event-chain Monte Carlo: Constructing Rejection-free Global-balance Algorithms from Infinitesimal Steps. *Journal of Chemical Physics* **2014**, *140*, 054116.
- (44) Engel, M.; Anderson, J. A.; Glotzer, S. C.; Isobe, M.; Bernard, E. P.; Krauth, W. Hard-disk Equation of State: First-order Liquid-hexatic Transition in Two Dimensions with Three Simulation Methods. *Physical Review E* **2013**, *87*, 042134.
- (45) Isobe, M.; Krauth, W. Hard-sphere Melting and Crystallization with Event-chain Monte Carlo. *Journal of Chemical Physics* **2015**, *143*, 084509.
- (46) Williams, J. R.; O’Connor, R. Discrete element simulation and the contact problem. *Archives of Computational Methods in Engineering* **1999**, *6*, 279–304.
- (47) de Berg, M.; Cheong, O.; van Kreveld, M.; Overmars, M. *Computational Geometry: Algorithms and Applications*, 3rd ed.; Springer, 2008.
- (48) Towns, J.; Cockerill, T.; Dahan, M.; Foster, I.; Gaither, K.; Grimshaw, A.; Hazelwood, V.; Lathrop, S.; Lifka, D.; Peterson, G. D.; Roskies, R.; Scott, J. R.; Wilkins-Diehr, N. XSEDE: Accelerating Scientific Discovery. *Computing in Science & Engineering* **2014**, *16*, 62–74.

Graphical TOC Entry

