



HAL
open science

Fast event-driven simulations for soft spheres: from dynamics to Laves phase nucleation

Antoine Castagnède, Laura Fillion, Frank Smallenburg

► To cite this version:

Antoine Castagnède, Laura Fillion, Frank Smallenburg. Fast event-driven simulations for soft spheres: from dynamics to Laves phase nucleation. *The Journal of Chemical Physics*, 2024, 161 (2), <10.1063/5.0209178>. <hal-04779880>

HAL Id: hal-04779880

<https://hal.science/hal-04779880v1>

Submitted on 13 Nov 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Fast event-driven simulations for soft spheres: from dynamics to Laves phase nucleation

Antoine Castagnède^{1,*} and Frank Smallenburg¹

¹*Université Paris-Saclay, CNRS, Laboratoire de Physique des Solides, 91405 Orsay, France*

Laura Filion²

²*Soft Condensed Matter and Biophysics, Debye Institute for Nanomaterials Science, Utrecht University, Utrecht, Netherlands*

Conventional molecular dynamics (MD) simulations struggle when simulating particles with steeply varying interaction potentials, due to the need to use a very short time step. Here, we demonstrate that an event-driven Monte Carlo (EDMC) approach first introduced by Peters and de With [Phys. Rev. E **85**, 026703 (2012)] represents an excellent substitute for MD in the canonical ensemble. In addition to correctly reproducing the static thermodynamic properties of the system, the EDMC method closely mimics the dynamics of systems of particles interacting via the steeply repulsive Weeks-Chandler-Andersen (WCA) potential. In comparison to time-driven MD simulations, EDMC runs faster by over an order of magnitude at sufficiently low temperatures. Moreover, the lack of a finite time step in EDMC circumvents the need to trade accuracy against simulation speed associated with the choice of time step in MD. We showcase the usefulness of this model to explore the phase behavior of the WCA model at extremely low temperatures, and to demonstrate that spontaneous nucleation and growth of the Laves phases is possible at temperatures significantly lower than previously reported.

I. INTRODUCTION

Conventional molecular dynamics (MD) simulations function by numerically integrating Newton's equations of motion for all particles in the system. To do this, MD simulations typically discretize time using a fixed time step. Ideally, this time step is chosen as large as possible to optimize the overall simulation speed. However, to obtain accurate results, the time step must be chosen small enough to avoid systematic errors due to the approximations involved in the discretization of time. Generally, the time step should be small enough that the forces on any given particle do not change significantly during a single step. This poses a significant challenge for many-body systems where the interaction potential and its derivatives are sharply varying functions of the interparticle distance.

The extreme version of an isotropic steeply varying interaction potential is the hard-sphere model. In this fundamental model system [1], the interaction potential between two particles jumps from zero to infinity at contact, rendering conventional MD algorithms useless. Intriguingly, this problem can be overcome via a different strategy: event-driven molecular dynamics (EDMD) simulations [2–6]. In the event-driven approach, collisions between particles are resolved as instantaneous events, which can be efficiently predicted based on the observation that between collisions, hard particles move in simple straight lines at constant velocity. However, this approach usually limits event-driven methods to systems of particles where the interaction potential only changes via

discontinuous jumps, such as hard spheres and square-well or square-shoulder models (see e.g. [3, 4, 7–11]).

In principle, it would be interesting to have simulation methods available that benefit from the potential efficiency of event-driven approaches but are still capable of handling particles with continuously varying interaction potentials. One way of accomplishing this is the discretization of continuous potentials [12, 13], although this necessarily results in approximated interactions. A novel alternative route was introduced in 2012 by Peters and de With [14], who developed an event-driven approach to simulate particles with continuous interaction potentials in the canonical ensemble. In this approach, which can be regarded as a rejection-free Monte Carlo method, the particles move in straight lines between discrete collisions with their neighbors. The collision distance between two particles is stochastically chosen based on the Boltzmann distribution, ensuring that the method correctly samples the canonical ensemble. By handling the interactions between pairs of particles as discrete events, this method neatly sidesteps the issues associated with finite integration time steps encountered in conventional MD of steep potentials. The approach of Peters and de With has proven useful in the extension of event-chain Monte Carlo (ECMC) methods [15] to continuous potentials [14, 16, 17]. ECMC simulations are an extremely efficient method to sample the static properties of many-body systems (especially those involving long-range interactions [18]), made possible by the use of simplified, unphysical dynamics involving chains of colliding particles [15]. However, this approach makes them less suitable for studying the dynamics of these systems. In contrast, the approach originally proposed in Ref. [14] is capable of modeling systems with near-physical dynamics, where all particles move collectively. This raises an interesting question: can this event-driven Monte Carlo

*Author to whom correspondence should be addressed:
antoine.castagnede@universite-paris-saclay.fr

(EDMC) approach be efficiently used as a substitute for molecular dynamics when exploring the dynamics of particles with almost-hard interactions?

Here, we combine the approach of Peters and de With with an efficient implementation of event-driven molecular dynamics, constructing an efficient simulation tool for exploring the behavior of short-ranged, steeply varying potentials. We test this method on systems of steeply repulsive particles (modeled via the Weeks-Chandler-Andersen (WCA) potential [19]), in order to assess its similarity to Newtonian dynamics and its efficiency. As expected, EDMC correctly reproduces the static equilibrium properties of the simulated systems, such as the equation of state and its phase behavior. More importantly, over a large temperature regime, the EDMC *dynamics* also closely resemble MD trajectories in terms of diffusion and structural relaxation. Moreover, EDMC can be simulated significantly more efficiently, especially close to the hard-sphere limit where conventional MD struggles. To highlight the usefulness of this method, we show two applications based on recent literature studies of low-temperature WCA models: one investigating the low-temperature phase behavior of the WCA model [20], and one examining the effect of softness on the nucleation of the Laves phases in almost-hard spheres [21].

In the remainder of this manuscript, we first describe the EDMC method and its application to the WCA potential. Next, we apply it to monodisperse systems of WCA particles and examine its dynamics and phase behavior. Finally, we use it to examine nucleation in binary mixtures of WCA particles of two different sizes. We end with a brief discussion and conclusions.

II. METHODS

A. Event-driven Monte Carlo simulations

In Ref. 14, Peters and de With proposed an innovative event-driven approach to simulate many-body systems with continuous pair potentials. Their approach is a rejection-free Monte Carlo method in the canonical ensemble (i.e. at fixed number of particles N , volume V , and temperature T). Similar to an EDMD simulation, the particles move in straight lines between collisions, and collide as if they are hard spheres. However, the distance of collision is stochastically determined based on the interaction potential and temperature of the system. As a result, the traditional acceptance rule of the Metropolis scheme [22] is replaced with a stochastic choice of a maximum energy gain between two particles before a collision occurs. In particular, the collision distance r_{col} between two particles is chosen such that the total energy gain of the particles as they approach is equal to a random energy ΔU taken from the Boltzmann distribution at the chosen temperature T . Any decrease in energy during this interval is ignored. At the time of collision, the two particles exchange momentum as if they are two hard




FIG. 1: Schematic view of the collision process between two particles as processed by the conventional time-driven MD approach (dashed line), and the present EDMC approach (solid lines). The stochastic nature of the EDMC simulation method can yield a range of different trajectories. Note that for simplicity, the blue particle is considered to be fixed. The background shading and dotted contour lines illustrate the interaction potential $U(r)$, with darker shading indicating higher energy.

spheres undergoing an elastic collision. As shown in Ref. 14, it can be proven that this correctly reproduces the equilibrium distribution in the canonical ensemble.

If we compare the dynamics generated by this EDMC scheme to those expected in a standard MD simulation (in the micro-canonical ensemble), then it is clear that the dynamics of this rejection-free method are not identical for any pair potential. Firstly, EDMC natively simulates the canonical ensemble, while MD simulations would require a thermostat to do so. Moreover, as illustrated in Fig. 1, in a normal MD trajectory of repulsive spheres, the velocities of a pair of “colliding” particles change continuously as the particles approach and separate again. In contrast, the EDMC trajectories consist of straight lines, meeting at a sharp corner. Although these two trajectories converge in the limit of infinitely steep repulsive (i.e. hard-sphere) potentials, this implies that the microscopic dynamics of the two simulation methods can in principle be expected to differ significantly. The effect of this change on the macroscopic dynamics remains to be investigated.

Event-driven simulations can be implemented extremely efficiently for isotropic hard particles [6, 23], suggesting that an event-driven approach to simulate soft potentials could potentially be more efficient than time-driven schemes. This is particularly the case in the limit of hard-sphere-like potentials, where MD typically requires small integration time steps. This limit is also the regime where we would expect the dynamics resulting from EDMC and MD to be most similar, and hence

where this approach may be suitable for studying dynamical features of the system under consideration.

To explore how EDMC simulations compare to MD methods both in terms of resulting dynamics and efficiency, here we apply this method to a commonly used interaction potential for sharply repulsive soft spheres: the WCA potential [19]. In this model, the pair potential between particles i and j is given by

$$U_{ij}(r) = \begin{cases} 4\epsilon \left[\left(\frac{\sigma_{ij}}{r} \right)^{12} - \left(\frac{\sigma_{ij}}{r} \right)^6 + \frac{1}{4} \right], & \frac{r}{\sigma_{ij}} < 2^{1/6} \\ 0 & \text{otherwise} \end{cases}, \quad (1)$$

where r is the distance between two particles, ϵ controls the strength of the interaction, and $\sigma_{ij} = (\sigma_i + \sigma_j)/2$ with σ_i the diameter of particle i . The cutoff for this WCA potential can then be written as $\sigma_{\text{cut}} = 2^{1/6}\sigma_{ij}$. Note that in the low-temperature limit, this system reduces to a hard-particle model, with the contact distance given by $\sigma_{\text{HS}} = \sigma_{\text{cut}}$. Note that Eq. 1 can be analytically inverted to obtain the distance corresponding to a certain pair energy U :

$$r_{\text{WCA}}(U) = \left(\frac{2}{1 + \sqrt{U/\epsilon}} \right)^{1/6} \sigma_{ij}. \quad (2)$$

To simulate the WCA potential, we modify an efficient implementation of EDMD for hard spheres [6]. The main change is in the prediction of collisions. For two spheres i and j that move in straight lines and collide at a separation distance r_{col} , the time of collision can be analytically predicted based on the instantaneous positions and velocities of said particles:

$$t_{\text{col}} = \frac{-\mathbf{r}_{ij} \cdot \mathbf{v}_{ij} - \sqrt{(\mathbf{r}_{ij} \cdot \mathbf{v}_{ij})^2 - v_{ij}^2 (r_{ij}^2 - r_{\text{col}}^2)}}{v_{ij}^2}, \quad (3)$$

where $\mathbf{r}_{ij} = \mathbf{r}_j - \mathbf{r}_i$ is the vector indicating the relative position of the two particles, and $\mathbf{v}_{ij} = \mathbf{v}_j - \mathbf{v}_i$ is their relative velocity. For hard spheres, r_{col} is simply equal to σ_{ij} . For the WCA potential, we instead find the collision distance stochastically, using [14]

$$r_{\text{col}} = r_{\text{WCA}}(U_{ij}(t) + \Delta U), \quad (4)$$

where $U_{ij}(t)$ is the pair energy between the two particles at the current simulation time, and ΔU is a random energy drawn from the (exponential) Boltzmann distribution at the chosen temperature T (see SM for implementation details [24–26]). Equation 3 then provides the moment of collision. If the argument of the square root is negative, no solution exists, and the two particles will not collide. Note that since the WCA potential is purely repulsive, we only have to consider pairs of particles that are currently approaching each other ($\mathbf{r}_{ij} \cdot \mathbf{v}_{ij} < 0$).

Predicted collisions are resolved in the event-driven fashion as if they are normal hard-sphere collisions, conserving kinetic energy and momentum. Note that as a

result, this scheme inherently conserves the total kinetic energy of the system, while the total potential energy fluctuates. Hence, it is important to ensure that at the start of the simulation, the average kinetic energy of the particles matches the chosen temperature T .

B. Comparing EDMC and MD

For comparison, we also simulate our systems using the simulation package LAMMPS [27], using a Nosé-Hoover thermostat to keep the temperature fixed.

Simulation times are compared in terms of the “thermal” time unit

$$\tau_{kT} = \sqrt{\beta m \sigma^2}, \quad (5)$$

where $\beta = 1/k_B T$ with k_B Boltzmann’s constant, m is the mass of a particle (which we always choose equal for all particles), and our length unit σ is chosen to be the diameter of the largest sphere size in our system. This time unit is a measure for the time it takes a free particle at temperature T to move a distance on the order of σ . Additionally, it is convenient to introduce an “energetic” time unit

$$\tau_\epsilon = \sqrt{\frac{m\sigma^2}{\epsilon}} = \sqrt{\frac{k_B T}{\epsilon}} \tau_{kT}, \quad (6)$$

as this is more conventionally used in time-driven simulations.

To initialize our systems in the fluid state, we make use of the hard-sphere EDMD code of Ref. 6, where the particles are grown as purely hard spheres from a random dilute initial configuration to reach the desired packing fraction. These configurations are then equilibrated with the WCA potential at the temperature of interest, using either EDMC or LAMMPS. The equilibrated configurations are used as initial configurations for the simulations where we perform our measurements of e.g. the potential energy, pressure, diffusion coefficient, intermediate scattering function, and simulation speed. For the specific case of benchmarking, we simulate multiple runs over a time of $100\tau_{kT}$, each sequentially starting from the last snapshot of the previous one. Benchmarking results were obtained on a dedicated machine running Ubuntu 20.04 and powered by two 8-core 3.3Ghz Intel Xeon Gold 6234 processors. All simulations used only a single core.

III. MODEL

In this paper, we consider both monodisperse and bidisperse systems of WCA particles.

A. Monodisperse WCA

The phase behavior of the monodisperse WCA model has been studied extensively (see e.g. Refs. 20, 28), and

consists of a fluid at low densities and a face-centered cubic crystal at high densities, separated by a coexistence region. In the limit of low temperatures ($k_B T/\epsilon \rightarrow 0$), the WCA model reduces to hard spheres with an effective diameter equal to $2^{1/6}\sigma$. As a result, the freezing and melting densities in this limit can be determined by simply rescaling the hard-sphere binodals [29], and are given by $\rho_F \sigma^3 \simeq 0.664$ and $\rho_M \sigma^3 = 0.774$, respectively. At higher temperatures, both the freezing and melting densities shift to higher values, as the particles effectively shrink.

As such, it is common to map the freezing transition of the monodisperse WCA system onto that of a pure hard sphere (HS) system. This way, one can define a temperature-dependent effective diameter σ_{eff} for the WCA particles and thus a temperature-dependent effective reduced density $\rho^* = \rho \sigma_{\text{eff}}^3$ for the system. To this end, we use the fit of the WCA freezing density ρ_F introduced by Filion *et al.* [30] based on the data from Ahmed and Sadus [28]:

$$\rho_F \sigma^3 \simeq 0.635 + 0.473 \left(\frac{k_B T}{\epsilon} \right)^{1/2} - 0.236 \frac{k_B T}{\epsilon}. \quad (7)$$

B. Binary WCA

As our second model, we consider a binary WCA mixture, with size ratio $\gamma = \sigma_s/\sigma_L = 0.78$ and fixed composition $x_L = N_L/(N_L + N_S) = 1/3$. Here, the subscripts S, L denote the small and large particles, respectively. In the low-temperature limit, this model reduces to a binary hard-sphere mixture of the same size ratio. Hard-sphere mixtures of approximately this size ratio are known to be capable of forming a stable binary crystal phase with the same structure as MgZn_2 , one of the Laves phases [31]. In fact, the other two Laves phase structures (MgCu_2 and MgNi_2) are only slightly less stable than MgZn_2 for hard-sphere mixtures as well, with only small free-energy differences separating them [31].

The low-temperature behavior of this WCA mixture, and in particular the interplay between glassy dynamics and the nucleation of the Laves crystal phases, has been studied previously by Dasgupta, Coli, and Dijkstra in Ref. 21. They observed the spontaneous nucleation of the Laves phases at relatively high temperatures $k_B T/\epsilon \gtrsim 0.025$, but concluded that at low temperatures, where the model approaches the hard-sphere limit, spontaneous nucleation is pre-empted by dynamical arrest associated with the glass transition. In particular, they determine both an instability line, marking the state points where the fluid phase becomes unstable with respect to crystallization, and a glass transition line, where the fluid is no longer able to dynamically relax, and argued that these cross at a low temperature on the order of $k_B T/\epsilon \simeq 0.01$. However, we note that this observation conflicts with recent observations of spontaneous nucleation of Laves phases in simulation studies of

binary mixtures of true hard spheres [32, 33], making this an interesting question to revisit.

To detect nucleation in the binary WCA model, we perform local structure identification in each snapshot using a supervised machine learning algorithm. In particular, we use a single-layer neural network classifier, following the methodology from Ref. 34. The network is trained to recognize particles whose local environment matches one of the different considered phases, specifically fluid, FCC, and each of the three Laves phases (MgCu_2 , MgZn_2 , and MgNi_2). The performed analysis is based on characterization of the local environment of particles in the system using the locally averaged bond order parameters introduced by Lechner and Dellago [35]. The weights and biases of the machine learning model we use in our work have been obtained by training this model on a set of snapshots of fluid, FCC, and all three Laves phases generated from simulations of a binary hard-sphere mixture with the same size ratio and composition as our chosen binary WCA mixture. Note that this system corresponds to the zero-temperature limit of our binary WCA model. We report that the accuracy of the model on the training set exceeds 98% for all considered phases, similar to the performances reported in Ref. 34.

IV. TEST CASE 1: MONODISPERSE WCA MODEL

We first focus on the monodisperse WCA model. In order to explore the dynamics produced by the EDMC method, we initially focus on a dense fluid phase at density $\rho \sigma^3 = 0.65$, which is just below freezing in the zero-temperature limit, and remains in the fluid phase at higher temperatures. At temperatures above $k_B T/\epsilon \simeq 0.01$, this density is in the stable fluid regime instead [20, 28].

In Fig. 2a, we show the mean squared displacement of the WCA model as measured using both EDMC and LAMMPS as a function of time. We see that for all investigated temperatures, there is very little difference between the two methods. A very close examination shows mild deviations at high temperatures, with the LAMMPS simulations resulting in slightly faster dynamics than EDMC.

In Fig. 2b, we examine the intermediate scattering function (ISF) $F(q, t)$ of the same systems, defined as

$$F(q, t) = \frac{1}{N} \langle \rho(q, t) \rho(-q, 0) \rangle, \quad (8)$$

where $\rho(q, t) = \sum_j \exp(-i\mathbf{q} \cdot \mathbf{r}_j)$ is the Fourier transform of the density and $F(q, t)$ describes the correlation of collective density fluctuations in the system. The ISF is measured at the wave vector $q = 2\pi/\sigma$. We observe much the same results as for the mean squared displacement: EDMC simulations result in slightly slower dynamics at high temperatures, but the differences are very subtle. Finally, to confirm that the switch to EDMC does




FIG. 2: Comparison between dynamics of EDMC (marked by solid dots (\bullet) and lines) and LAMMPS (marked by empty squares (\square) and dashed lines) for monodisperse WCA at density $\rho\sigma^3 = 0.65$, for a range of temperatures. a) Mean squared displacement as a function of time. b) Intermediate scattering function $F(q, t)$ ($q = 2\pi/\sigma$) as a function of time. c) Relaxation time τ_α as a function of the wave-vector q . Note that error bars (determined as twice the standard error) are shown for all points but are typically smaller than the point size. All lines are guides to the eye.

not introduce deviations in the dynamics at other length scales, we show in Fig. 2c the variation of the structural relaxation time τ_α (determined as the point where $F(q, t) = 1/e$) with the wave vector q . We see the same qualitative behavior at all examined temperatures over the range of wave vectors considered, meaning that similar relaxation dynamics are observed at all investigated length scales.

To demonstrate the effectiveness of EDMC for establishing equilibrium properties of systems near the hard-sphere limit, we turn our attention to determining the phase boundaries of the low-temperature WCA system, using the direct coexistence approach of Ref. 29. In particular, we perform direct coexistence simulations in the canonical ensemble using an elongated simulation box, such that the fluid and FCC crystal phases coexist in a slab geometry, separated by two interfaces normal to the long z -axis of the box (see Fig. 3a). As the lattice spacing is fixed by the initial configuration and the box size in the xy -plane, we perform multiple runs at each temperature with several different lattice spacings chosen roughly near the melting density, and identify the lattice spacing which corresponds to the equilibrium phase coexistence (see Ref. 29 for more details). Figure 3 shows the coexistence pressure as a function of temperature for systems of $N = 6144$ particles. We compare this result to the data from Attia, Dyre, and Pedersen, who used an interface pinning approach combined with MD simulations [20] for approximately the same system size, finding excellent agreement. A key advantage of the EDMC method is that it can readily be extended to arbitrarily low temperatures, as it does not rely on discrete time steps. Hence, we are able to extend our results to much lower temperatures than were easily accessible in the MD simulations of Ref. 20. As expected, at low tempera-

tures the coexistence pressure closely approaches that of a monodisperse hard-sphere system with effective particle diameter $\sigma_{HS} = 2^{1/6}\sigma$.

The above results clearly show that the EDMC approach accurately reproduces the thermodynamics and dynamics of the studied monodisperse WCA system, even at temperatures where the interaction potential is significantly removed from hard spheres. The next question that we would like to answer is how efficient this method is in comparison to a conventional MD simulation. To do this, however, we first have to establish a reasonable choice of integration time step for the MD simulations. Specifically, we want the time step δt to be small enough to reproduce the thermodynamics of the system associated with the limit of infinitely small time steps. On the other hand, from an efficiency point of view, we want the time step to be large, such that fewer integration steps are needed to simulate the desired time interval.

In Fig. 4, we plot the behavior of the potential energy of the WCA system as measured in LAMMPS for different choices of the time step and for several different temperatures. Clearly, when the time step is too large, (e.g. $\delta t = 0.01\tau_\epsilon$), we observe a shift to noticeably higher energies in comparison to the true value reached in the limit of small time steps. Based on these results, and on the weak effect of temperature on this behavior, we choose $\delta t = 0.001\tau_\epsilon$ for all of our MD simulations in this paper. Note that the time required to perform an MD simulation over a fixed total time interval trivially scales as $(\delta t)^{-1}$. The following benchmarking results should be regarded in light of this performance dependence: by raising the time step by e.g. a factor of 10, the performance of the MD simulations could be boosted by a similar factor, at the cost of a small but noticeable devi-




FIG. 3: a) Typical snapshot of a direct coexistence simulation. b) WCA phase diagram at very low temperatures. Empty circles represent recent estimates also obtained from interface pinning simulations[20]. The estimate for $\beta P_{\text{coex}}^{\text{HS}} \sigma_{\text{HS}}^3 = 11.55668$ was obtained using the same methodology for a pure hard-sphere system of the same size [29]. Error bars are determined as twice the standard error on five independent measurements, and are mostly smaller than the points. The dashed line shows the expected low-temperature scaling (see Ref. 20).

ation of the thermodynamic properties of the simulated system (see Fig. 4).

In Fig. 5, we show benchmarking results comparing EDMC to the the LAMMPS simulations. In particular, we plot the ratio of number of time units simulated per second in both an EDMC and a LAMMPS simulation at the same state point. In Fig. 5a we plot this quantity for a fixed density $\rho\sigma^3 = 0.68$ as a function of $k_B T / \epsilon$, reported for a series of system sizes. We have confirmed that for all these simulations, the systems remained in the fluid state. The EDMC approach significantly outperforms the MD approach over most of the investigated temperature range, and especially in the low-temperature regime, where it is faster than MD by a factor of approximately 20. This speed-up can be understood by considering the number of simulation “steps” (i.e. collisions or integration steps) required to simulate $1 \tau_\epsilon$ in time. For MD, this is fixed, as it is purely determined by the integration time step. For EDMC, this is simply the number of particle collisions during the chosen time interval,




FIG. 4: MD measurements of the average energy of the WCA model at density $\rho\sigma^3 = 0.68$ as a function of the MD time step, for different temperatures as indicated. The solid orange lines are cubic fits to the data at $\delta t / \tau_\epsilon < 0.03$, and the dashed green lines indicate the extrapolation of that fit to $\delta t = 0$. For these results, we used $N = 8000$.

which scales $\propto \sqrt{k_B T / \epsilon}$ due to the thermal slowdown of the particles, and hence vanishes in the low-temperature limit. Note that this observation only holds because the pressure P (in units of ϵ / σ^3) in the low-temperature limit vanishes for this system: for a system in which the particles still feel interactions with their neighbors in the ground state, we would not expect the same scaling. We also observe that LAMMPS scales slightly better with system size, as can be seen from the decrease in relative performance for larger system sizes. The exact effect of system size on the relative performance is not monotonic, likely due to details in e.g. the way the number of cells in the cell list is chosen in each simulation, which can lead to discrete ‘steps’ in performance. Although it is possible to further optimize both simulation methods by fine-tuning the neighbor list parameters for each density, we have not done so here. However, we have checked for specific state points that changing the neighbor list shell size does not significantly enhance the speed of either simulation method: for both methods this optimization leads to only small performance differences that do not exceed $\sim 15\%$.

Figure 5b shows the relative performance of the two methods as a function of density $\rho\sigma^3$ for a monodisperse WCA system equilibrated in the fluid phase at

temperature $k_B T/\epsilon = 0.001$. We observe that the relative efficiency of our simulation method slightly decreases as a function of density but consistently outperforms LAMMPS by a significant factor. The high efficiency of the EDMC simulation at low densities is understandable, since the computational effort required by this method is concentrated in handling collisions. Overall, both methods are more efficient at low densities, due to the fact that during either collision predictions or force calculations, fewer neighboring particles have to be considered. In addition to this, the EDMC simulation benefits from a lower number of collisions per unit time, while the MD simulation still needs to consider the same number of integration steps per time unit.

In short, for the case of monodisperse WCA systems, we find that the EDMC approach significantly outperforms the MD approach for a large range of (low) temperatures and system sizes. The EDMC method is competitive with MD even at temperatures around $k_B T/\epsilon = 1$, and – as expected – performs best in the limit of low temperatures. To examine how this method performs in a more complex system, we explore in the next section the application of the EDMC approach to a binary WCA system, and in particular examine glassy behavior and crystal nucleation.

V. TEST CASE 2: LAVES PHASE NUCLEATION IN A BINARY WCA MODEL

For the binary WCA mixture, we first compare the dynamics obtained from EDMC and MD in Fig. 6. In Fig. 6a we show the ISF $F(q, t)$ at low temperature for different densities, which in contrast to the monodisperse case shows the clear formation of a plateau at intermediate times. This plateau, which lengthens with increasing system density, can be attributed to the caging of particles by their neighbors, and is typical of glassy systems. We notice a much larger range of relaxation times compared to the monodisperse case (figure 6b) which implies much longer simulation times are required, even for the sole purpose of equilibration. Again, EDMC accurately reproduces the dynamics of conventional MD, with the LAMMPS simulations possibly having a slightly shorter relaxation time, although this difference is hard to resolve accurately. Similarly to the monodisperse case, we also recover the expected relaxation dynamics over a large span of temperatures and length scales, as shown in figure 6b.

In Fig. 7, we compare the simulation speed of the EDMC method to the LAMMPS package as a function of temperature (Fig. 7a) and density (Fig. 7b). In particular at low temperatures, we again find a significant speedup (up to a factor of approximately 40).

The efficiency of the EDMC method at low temperatures allows us to use this simulation approach to look at the interplay between glassy dynamics and Laves phase nucleation at temperatures lower than those that were

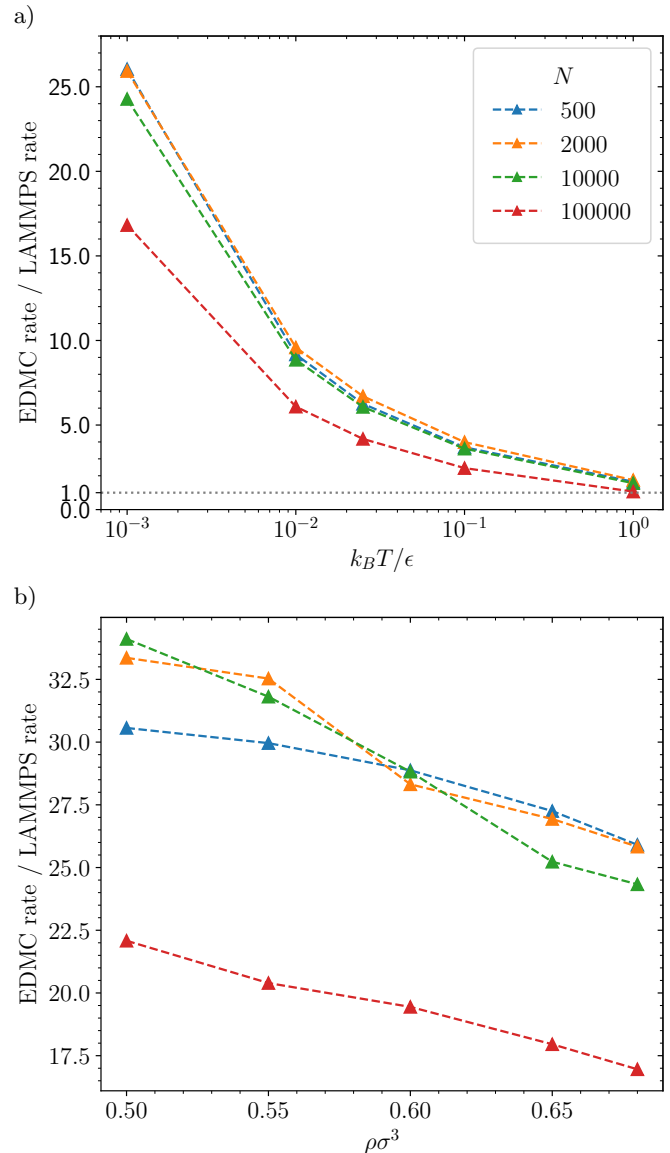


FIG. 5: Relative performance of the EDMC and MD methods for monodisperse WCA systems of several sizes, expressed as the ratio of the number of time units simulated per second in each simulation method, a) at fixed density $\rho\sigma^3 = 0.68$ and b) at fixed temperature $k_B T/\epsilon = 0.001$. Performance evaluation was obtained from successive short simulations of duration $100\tau_{kT}$ started from an equilibrated fluid system. Standard deviations are typically found to be much less than 1%.

easily accessible to the time-driven molecular dynamics methods previously used to study nucleation in this system (Ref. 21). To this end, we run brute-force nucleation simulations for a range of densities $\rho\sigma^3$ and temperatures $k_B T/\epsilon = 0.005$ and $k_B T/\epsilon = 10^{-4}$, using system sizes of $N = 2000$ particles. We run our simulations for at least $6 \times 10^6 \tau_{kT}$. For comparison, the lowest temperature examined in Ref. 21 was $k_B T/\epsilon = 0.005$.

In contrast to Ref. 21, we do observe spontaneous nucleation at a temperature $k_B T/\epsilon = 0.005$. As an ex-

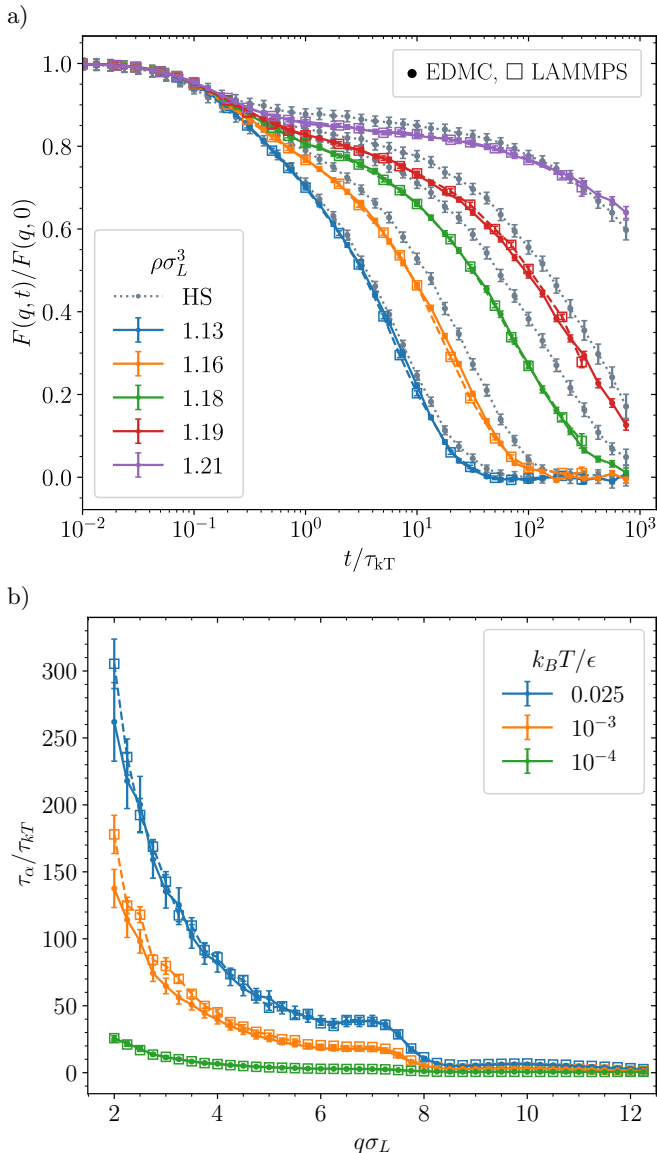


FIG. 6: Comparison between dynamics of EDMC and LAMMPS for a binary WCA mixture with composition $x_L = 1/3$ and size ratio $\gamma = 0.78$. Symbols and lines are the same as in figure 2. a) Intermediate scattering function $F(q, t)$ ($q = 2\pi/\sigma_L$) as a function of time at low temperature $k_B T/\epsilon = 10^{-4}$, and at several densities as indicated. HS systems at the equivalent effective density (mapped using Eq. 7) are shown as dotted lines. b) Relaxation time τ_α as a function of the wave-vector $q\sigma_L$ for multiple temperatures, at density $\rho\sigma_L^3 = 1.175$. Error bars are determined as twice the standard error.

ample of this, we show in Fig. 8 the growth of the largest connected cluster of particles in a Laves-phase-like environment (see Methods). In particular, we show four nucleation trajectories taken at the same state point inside the coexistence region ($\rho\sigma^3 = 1.213$, $k_B T/\epsilon = 0.005$). We observe a clear (random) waiting time followed by a jump to a finite-sized cluster consisting of approximately half of the total number of particles in the system ($N = 2000$).

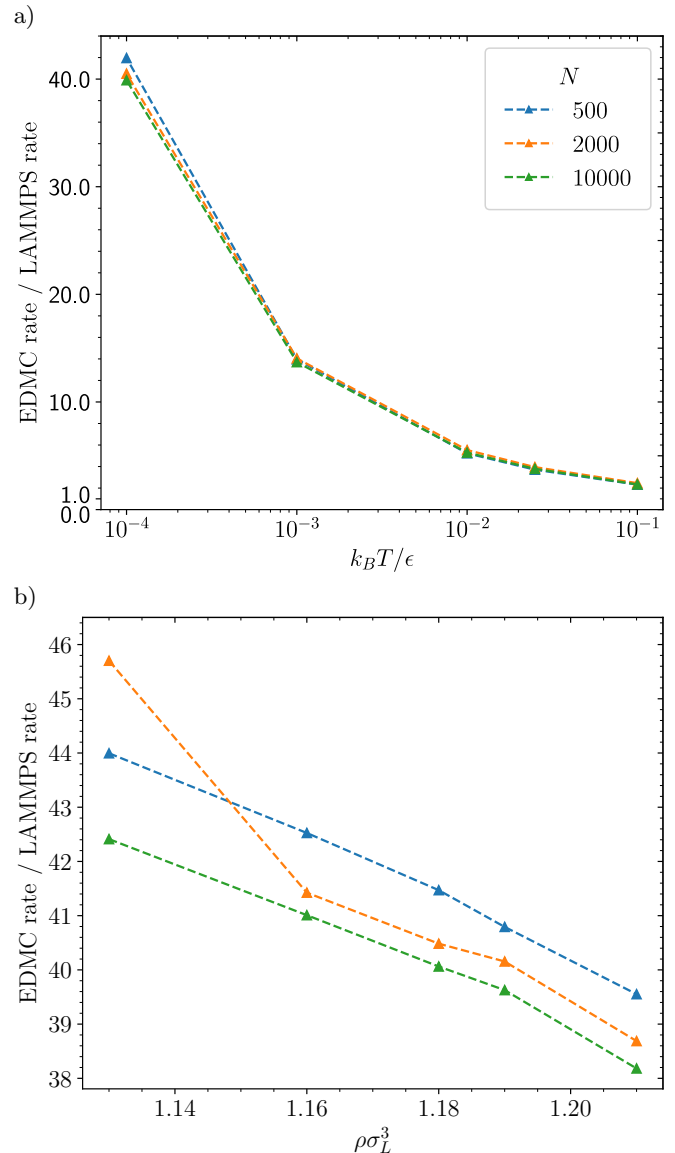


FIG. 7: Relative performance of the EDMC and MD methods for several system sizes, expressed as the ratio of the number of time units simulated per second in each simulation method, a) at fixed density $\rho\sigma_L^3 = 1.175$ and b) at fixed temperature $k_B T/\epsilon = 10^{-4}$. Performance evaluation was obtained from successive short simulations of duration $100\tau_{kT}$ started from an equilibrated fluid system. Standard deviations are typically found to be much less than 1%.

Analysis of the system using our order parameter shows that this state corresponds to a roughly spherical cluster (second snapshot in Fig. 8). Running the simulation longer, we also sometimes observe a second transition to a system-spanning cluster (third snapshot in Fig. 8). This can sometimes be seen in the evolution of the cluster size over time as a secondary jump to a larger value (as visible in the red line in Fig. 8). In terms of structure composition, we typically observe a mixture of the competing MgZn_2 and MgCu_2 Laves phases, with MgNi_2 appearing

only in small amounts, and typically at later stages of the simulation. Similar to what is reported in Ref. 21, it appears that MgZn_2 -like environments emerge rapidly from the metastable fluid and that they later transform into MgCu_2 -like environments as the crystal nucleus extends across the simulation box. These observations are consistent with the previously reported small ($\sim 10^{-3}k_B T$ per particle) free-energy difference that exists between these phases [31].

We note that the state point investigated in Fig. 8 lies below the ‘‘instability line’’ determined by Dasgupta *et al.*[21], which they defined as the temperature below which the supersaturated fluid becomes unstable to crystallization and hence where they expected no discrete nucleation events to be observable. We also report in Fig. 9 a selection of nucleation trajectories recorded at an even lower temperature $k_B T/\epsilon = 10^{-4}$, at different densities. As one might expect, we observe a random waiting time at low densities ($\rho\sigma_L^3 \lesssim 1.188$), and an immediate onset of crystallization at higher densities. On further increasing the density, crystallization still sets in immediately, but the crystal growth becomes slower, due to the overall slowdown of the dynamics. In these glassy systems, we also sometimes see jumps in the cluster size that rapidly revert (see e.g. the curve associated with $\rho\sigma_L^3 = 1.211$). We attribute this to the crystal nucleus consisting of multiple domains that have joined, but are not fully aligned yet, such that thermal fluctuations may temporarily ‘break’ the connection between parts of the cluster. In our simulations at this temperature, we no longer observe crystallization for densities higher than $\rho\sigma_L^3 \gtrsim 1.22$, as the system becomes too glassy. A broader selection of nucleation trajectories can be found in the Supplementary Material (SM).

In contrast to the findings of Ref. 21, we do not find a low-temperature regime where spontaneous crystal nucleation and growth is completely pre-empted by glassy behavior. This may be attributable to the ability of the EDMC approach to reach longer time scales at low temperatures than the conventional MD simulations of Ref. 21. Additionally, Dasgupta *et al.* used Monte Carlo (MC) simulations to determine the glass transition for their system. As MC simulations have their own microscopic dynamics (which can depend on the chosen step size), this could influence the predicted critical packing fraction associated with dynamical arrest. Note that our observation of nucleation and growth of the Laves phases at low temperatures is consistent with past observations of Laves phase nucleation in pure hard spheres at similar size ratios [32, 33], which can be seen as the low-temperature limit of the WCA model considered here.

VI. CONCLUSIONS

We have demonstrated that the rejection-free event-driven Monte Carlo approach introduced by Peters and de With [14], combined with an efficient implementation

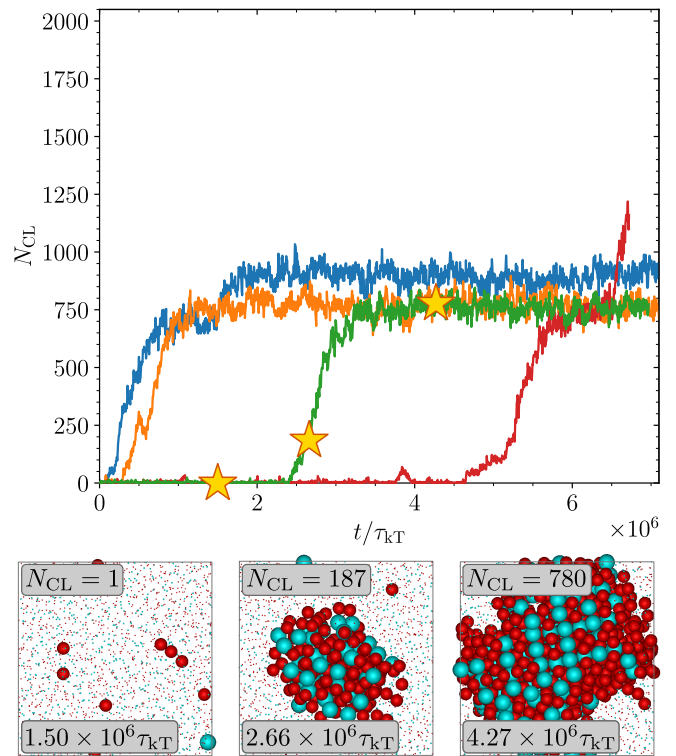


FIG. 8: Nucleation events at low temperature for the binary WCA system ($x_L = 1/3$, $\gamma = 0.78$), showing multiple nucleation events with waiting time at the same state point ($k_B T/\epsilon = 0.005$, $\rho\sigma_L^3 = 1.213$). The total system consisted of $N = 2000$ particles. The snapshots show the evolution of one system (green line) undergoing crystallization. Large particles are depicted in blue and small particles in red. The yellow stars in the plot report where each snapshot of the system was taken.




FIG. 9: Nucleation events recorded at very low temperature ($k_B T/\epsilon = 10^{-4}$) for a binary WCA system with composition $x_L = 1/3$ and size ratio $\gamma = 0.78$ for a range of densities.

of event-driven molecular dynamics [6], provides an accurate and efficient method for simulating the dynamics of model systems with sharply varying interactions like the WCA model. At sufficiently low temperatures, this approach circumvents the time-step issues encountered by molecular dynamics simulations, resulting in significant boosts in speed.

We should acknowledge here that we have considered an ideal case for the EDMC method: a short-ranged, purely repulsive interaction potential, which can be easily analytically inverted to obtain an expression for the distance as a function of interaction energy. When considering other systems, the suitability of EDMC as a substitute for MD simulations comes with a number of caveats.

First, every collision prediction in an EDMC simulation requires the evaluation of both $U_{ij}(r)$ and its inverse $r(U_{ij})$. This means that the function should ideally be analytically invertible. In turn, this would incur a significant performance penalty if one were to solve numerically for the inverse during every collision. However, this issue could in principle be addressed with e.g. lookup tables, or by avoiding most exact evaluations of the potential using an approximate inverse of the potential (thinning) [36].

Second, for interaction potentials with both attractive and repulsive components (e.g. Lennard-Jones), collision prediction needs to check for both attractive and repulsive collisions [14], similar to EDMD simulations of square-well systems. This will slow down the simulation speed. Note that oscillatory interaction potentials (with multiple switches between attractive and repulsive regimes) would exacerbate this slowdown. Preliminary tests on the Lennard-Jones model show that the EDMC model can still be competitive with conventional MD.

Third, as potentials become more long-ranged, the size of the neighbor lists increases. In event-driven simulations of large systems, this will significantly increase the cost of collision calculations due to scattered memory access.

Fourth, we note that similar to EDMD, EDMC could in principle be adapted to non-spherical interactions. However, for the vast majority of anisotropic models, an analytical prediction of a collision time is not feasible. Numerically solving for collision times (see e.g. [37]) would be significantly slower, and unlikely to be competitive with conventional MD.

Fifth, the EDMC approach inherently simulates a constant-temperature ensemble. While this may be advantageous in some cases, there are scenarios where energy conservation is crucial, e.g. when studying heat transport through a system.

Finally, despite a few existing routes in this direction [38, 39], we point out that EDMC, similar to EDMD, is difficult to parallelize efficiently, which may limit its suitability to extremely large-scale simulations.

The EDMC approach also offers several advantages (apart from its efficiency under suitable conditions).

Most important is the ability of EDMC to simulate arbitrarily steep interaction potentials, without impacting the required time step. Second, unlike either MD or EDMD, EDMC can handle interactions which combine discontinuous and continuous elements. Examples of these would be e.g. hard-spheres combined with a screened-Coulomb repulsion, or an Asakura-Oosawa depletion attraction. Third, the avoidance of a finite integration step size also implies that EDMC circumvents any systematic inaccuracies this discretization might introduce: the EDMC approach simulates the implemented interaction potential exactly, without requiring the approximations necessary for MD. Finally, in systems that are (at least partially) in a gas phase, EDMC will be highly efficient in comparison to MD, as the free flight of particles is extremely computationally cheap in event-driven simulations.

In conclusion, for studying the dynamics of suitable model systems, the EDMC approach is an appealing alternative to molecular dynamics simulations in the canonical ensemble, with potential speedups of over an order of magnitude.

VII. SUPPLEMENTARY MATERIAL

Additional details on random number generation as well as additional low temperature nucleation trajectories may be found in the Supplementary Material available online.

VIII. ACKNOWLEDGMENTS

AC and FS gratefully acknowledge funding from the Agence Nationale de la Recherche (ANR), grant ANR-21-CE30-0051. LF acknowledges funding from the Vidi research program with project number VI.VIDI.192.102 which is financed by the Dutch Research Council (NWO).

IX. AUTHOR DECLARATIONS

A. Conflict of Interest

The authors have no conflicts to expose.

X. DATA AVAILABILITY

A version of the EDMC simulation code can be found at <https://github.com/FSmallenburg/EDMC/>. The data that support the findings of this study is available in a data package on Zenodo [40].

XI. REFERENCES

- [1] C. P. Royall, P. Charbonneau, M. Dijkstra, J. Russo, F. Smallenburg, T. Speck, and C. Valeriani, arXiv preprint arXiv:2305.02452 (2023).
- [2] B. J. Alder and T. E. Wainwright, *The Journal of Chemical Physics* **31**, 459 (1959).
- [3] D. C. Rapaport, *Progress of Theoretical Physics Supplement* **178**, 5 (2009).
- [4] M. N. Bannerman, R. Sargant, and L. Lue, *Journal of computational chemistry* **32**, 3329 (2011).
- [5] C. De Michele, *Computer Physics Communications* **182**, 1846 (2011).
- [6] F. Smallenburg, *The European Physical Journal E* **45**, 22 (2022), ISSN 1292-8941, 1292-895X.
- [7] M. Bannerman, J. Magee, and L. Lue, *Physical Review E* **80**, 021801 (2009).
- [8] T. Dotera, T. Oshiro, and P. Ziherl, *Nature* **506**, 208 (2014).
- [9] V. R. Akkaya, I. Kandemir, et al., *Mathematical Problems in Engineering* **2015** (2015).
- [10] I. E. Paganini, C. Pastorino, and I. Urrutia, *The Journal of Chemical Physics* **142**, 244707 (2015).
- [11] A. Plati, R. Maire, E. Fayen, F. Boulogne, F. Restagno, F. Smallenburg, and G. Foffi, *Nature Physics* **20**, 465 (2024).
- [12] C. Thomson, L. Lue, and M. N. Bannerman, *The Journal of Chemical Physics* **140**, 034105 (2014).
- [13] S. Ucyigitler, M. C. Camurdan, and J. R. Elliott, *Industrial & engineering chemistry research* **51**, 6219 (2012).
- [14] E. A. J. F. Peters and G. de With, *Physical Review E* **85**, 026703 (2012), ISSN 1539-3755, 1550-2376.
- [15] E. P. Bernard, W. Krauth, and D. B. Wilson, *Physical Review E* **80**, 056704 (2009).
- [16] E. P. Bernard and W. Krauth, *Physical Review E* **86**, 017701 (2012).
- [17] M. Michel, S. C. Kapfer, and W. Krauth, *The Journal of chemical physics* **140**, 054116 (2014).
- [18] S. C. Kapfer and W. Krauth, *Physical Review E* **94**, 031302 (2016).
- [19] J. D. Weeks, D. Chandler, and H. C. Andersen, *The Journal of chemical physics* **54**, 5237 (1971).
- [20] E. Attia, J. C. Dyre, and U. R. Pedersen, *The Journal of Chemical Physics* **157**, 034502 (2022), ISSN 0021-9606, 1089-7690, 2205.00063.
- [21] T. Dasgupta, G. M. Coli, and M. Dijkstra, *ACS Nano* **14**, 3957 (2020), ISSN 1936-0851, 1936-086X.
- [22] D. Frenkel and B. Smit, *Understanding molecular simulation: from algorithms to applications* (Elsevier, 2023).
- [23] M. Klement and M. Engel, *The Journal of Chemical Physics* **150**, 174108 (2019), ISSN 0021-9606, 1089-7690, 1903.11392.
- [24] D. Blackman and S. Vigna, *ACM Transactions on Mathematical Software (TOMS)* **47**, 1 (2021).
- [25] J. H. Ahrens and U. Dieter, *Communications of the ACM* **15**, 873 (1972).
- [26] L. Devroye, *Continuous Univariate Densities* (Springer New York, New York, NY, 1986), pp. 379–484, ISBN 978-1-4613-8643-8.
- [27] A. P. Thompson, H. M. Aktulga, R. Berger, D. S. Bolinteanu, W. M. Brown, P. S. Crozier, P. J. In 'T Veld, A. Kohlmeyer, S. G. Moore, T. D. Nguyen, et al., *Computer Physics Communications* **271**, 108171 (2022), ISSN 00104655.
- [28] A. Ahmed and R. J. Sadus, *Physical Review E* **80**, 061101 (2009), ISSN 1539-3755, 1550-2376.
- [29] F. Smallenburg, G. Del Monte, M. de Jager, and L. Fillion, *The Journal of Chemical Physics* **160**, 224109 (2024).
- [30] L. Fillion, R. Ni, D. Frenkel, and M. Dijkstra, *The Journal of Chemical Physics* **134**, 134901 (2011), ISSN 0021-9606, 1089-7690.
- [31] A.-P. Hynninen, L. Fillion, and M. Dijkstra, *The Journal of chemical physics* **131**, 064902 (2009).
- [32] P. K. Bommineni, M. Klement, and M. Engel, *Physical Review Letters* **124**, 218003 (2020).
- [33] S. Marín-Aguilar, H. H. Wensink, G. Foffi, and F. Smallenburg, *Physical Review Letters* **124**, 208005 (2020).
- [34] E. Boattini, M. Ram, F. Smallenburg, and L. Fillion, *Molecular Physics* **116**, 3066 (2018), ISSN 0026-8976, 1362-3028.
- [35] W. Lechner and C. Dellago, *The Journal of Chemical Physics* **129**, 114707 (2008), ISSN 0021-9606, 1089-7690, 0806.3345.
- [36] G. Tartero and W. Krauth, *American Journal of Physics* **92**, 65 (2024).
- [37] L. Hernández de la Peña, R. van Zon, J. Schofield, and S. B. Opps, *The Journal of chemical physics* **126**, 074106 (2007).
- [38] S. Miller and S. Luding, *Journal of Computational Physics* **193**, 306 (2004).
- [39] M. A. Khan and M. C. Herbordt, *Journal of computational physics* **230**, 6563 (2011).
- [40] A. Castagnède, L. Fillion, and F. Smallenburg, [Data set] Zenodo (2024), URL <https://doi.org/10.5281/zenodo.12608286>.