Massively parallel Monte Carlo for many-particle simulations on GPUs

Joshua A. Anderson [a], Eric Jankowski [a], Thomas L. Grubb [b], Michael Engel [a],
Sharon C. Glotzer [a,b,*]

[a] Department of Chemical Engineering, University of Michigan, Ann Arbor, MI 48109, USA
[b] Department of Materials Science and Engineering, University of Michigan, Ann Arbor, MI 48109, USA

A R T I C L E   I N F O

A B S T R A C T

Current trends in parallel processors call for the design of efficient massively parallel algorithms for scientific computing. Parallel algorithms for Monte Carlo simulations of thermodynamic ensembles of particles have received little attention because of the inherent serial nature of the statistical sampling. In this paper, we present a massively parallel method that obeys detailed balance and implement it for a system of hard disks on the GPU. We reproduce results of serial high-precision Monte Carlo runs to verify the method. This is a good test case because the hard disk equation of state over the range where the liquid transforms into the solid is particularly sensitive to small deviations away from the balance conditions. On a Tesla K20, our GPU implementation executes over one billion trial moves per second, which is 148 times faster than on a single Intel Xeon E5540 CPU core, enables 27 times better performance per dollar, and cuts energy usage by a factor of 13. With this improved performance we are able to calculate the equation of state for systems of up to one million hard disks. These large system sizes are required in order to probe the nature of the melting transition, which has been debated for the last forty years. In this paper we present the details of our computational method, and discuss the thermodynamics of hard disks separately in a companion paper.

© 2013 Elsevier Inc. All rights reserved.

1. Introduction

During the last decades computational scientists have enjoyed a doubling of performance for single-threaded applications every two years solely from improvements in computer architecture. This is no longer the case. Current processor designs run into the Power Wall, which limits attainable clock speeds in a given power budget, and the Instruction Level Parallelism (ILP) Wall, which exists because there is only so much ILP that can be extracted from a typical program [1]. Moore's law still holds, for now, and the additional transistors go into increasing the core counts on each new chip. Consequently, researchers must utilize parallelism to execute larger, longer, or more demanding calculations and simulations.

As of this publication, a cluster of networked multi-core CPUs is the most common system architecture. In an alternative approach, a single graphics processing unit (GPU) can execute thousands of instructions at the same time and provides the performance of a small cluster at a fraction of the cost [2]. GPUs are becoming popular as desktop 'personal supercomputers' and as coprocessors in heterogeneous clusters. A successful GPU algorithm divides a given computation into a maximal number of identical, fully independent, and simple tasks, called *threads*. To fully utilize their potential it is necessary to design not just parallel, but *massively parallel* algorithms that scale to thousands of threads. Over the last few years, many

---

* Corresponding author.
 *E-mail address:* sglotzer@umich.edu (S.C. Glotzer).

Fig. 1. In a serial hard disk simulation, each trial move typically consists of the following: (i) Select a single disk at random, (ii) apply a random displacement to it, and (iii) accept the move if it generates no overlaps. The cell list enables $O(1)$ overlap checks by limiting the search space to only nine local cells. A sweep is defined as N consecutive trial moves, where N is the number of disks in the simulation.

problems have successfully been adapted to GPUs. One example is the molecular dynamics (MD) method for simulating thermodynamic ensembles of particles, which is well suited for massive parallelization. Numerous MD software packages support GPUs, including HOOMD-blue [3,4], LAMMPS [5], AMBER [6,7], NAMD [8], OpenMM [9], FEN ZI [10], HALMD [11], the work by Rapaport [12], GROMACS [13] and ACEMD [14].

A case where implementation to GPUs has so far not been achieved is Monte Carlo (MC) applied to off-lattice many-particle systems. MC is a statistical, rather than deterministic, sampling method that, appropriately implemented, samples the microstates of desired thermodynamic ensembles. It is the method of choice in many situations because it only requires an interaction potential, not a force field, allowing e.g. the use of non-differentiable pair potentials. Such potentials are useful in the simulation of hard particles, which interact solely via excluded volume. MC is also flexible in the sense that a wide variety of update moves can be applied [15–18]. MC is easy to implement on serial machines where each step selects a new microstate at random and accepts or rejects the new microstate based on the Boltzmann factor. An example trial move involves translating a single particle in a random direction. Since the acceptance of a trial move depends on results of prior moves, subsequent moves usually cannot be performed independently. A massively parallel algorithm must not only update a large number of particles at the same time, but also do so in a statistically correct way. This can be achieved by obeying detailed balance, which is not a trivial task. For these reasons, parallel MC codes for particle systems have received much less attention in the literature than parallel MD.

Available parallel MC algorithms fall into several categories. Most of them rely on domain decomposition schemes to update portions of the problem in parallel. Lattice MC has been employed for Ising models [19,20] including GPU implementations [21,22]. Related schemes for particle systems exist [23–27], but none scale up to thousands of threads. Moreover, the updating of domains employed in some of those methods can introduce a sampling bias that precludes the balance conditions. Asynchronous parallel algorithms [28,29] are poorly suited for GPUs because of their extensive inter-thread communication. Much work has been done on parallelizing Kinetic Monte Carlo simulations [30,31], but these do not directly apply to the systems we are interested in. Hybrid approaches that employ MD trajectories to create trial configurations [32,33] can be an effective way to exploit GPU parallelism, but require substantial additional programming and are not guaranteed to evolve any faster than MD alone.

In this paper, we develop an algorithm for massively parallel Monte Carlo (MPMC) simulations of many-particle systems that obeys detailed balance. As a test case, we implement it on the GPU for a system of hard disks in two dimensions (Fig. 1) and validate it with comparisons to recent large-scale serial event-chain MC simulations [34]. Our algorithm is not specific to this system, and it is valid for any MC simulation with local interactions between particles or lattice sites on any massively parallel computer architecture.

2. Algorithm

In developing any parallel application, the programmer must identify the computations that can be executed simultaneously. As the total work is broken into smaller tasks, opportunities for scaling to more processors increase. GPUs execute especially fine grained work loads. In GPU MD applications, one thread typically acts on a single particle [3,5,8,11,12]. Such
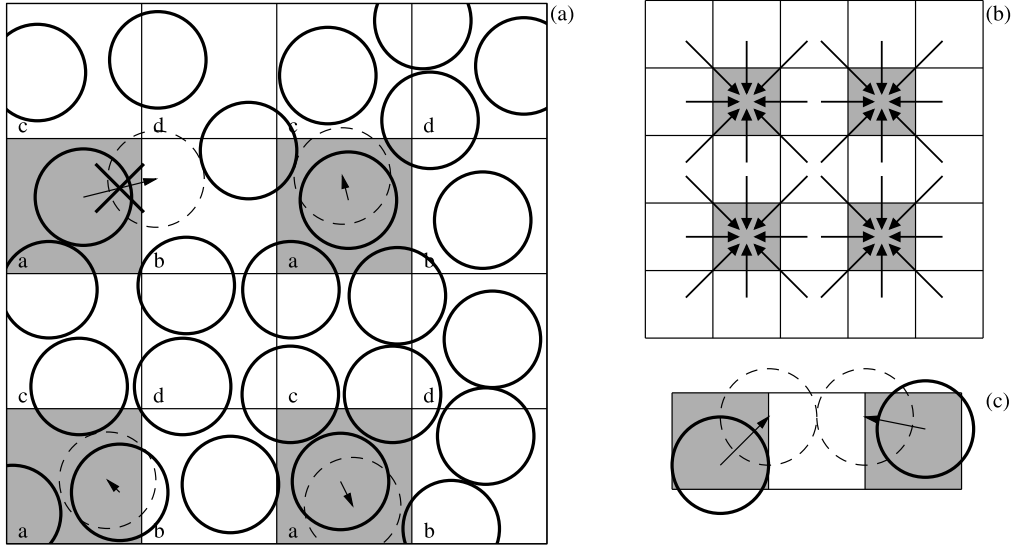
Fig. 2. (a) In massively parallel Monte Carlo (MPMC), trial moves are concurrently applied to particles in a subset of the cells. Moves that leave the cell are rejected. (b) Selected cells are separated by one row or one column of inactive cells. During the evaluation of the acceptance criterion, each active cell reads the particles in the eight neighboring inactive cells. (c) Simultaneous trial moves do not interact when the cell width is greater than the interaction range σ.

a decomposition is not directly applicable to traditional MC because each trial move depends on the state of the neighboring particles.

In MPMC, we utilize the *cell list* data structure for parallel decomposition, as well as overlap checks in the case of hard particles. A checkerboard decomposition permits many cells to be updated independently [19]. Although similar to applying trial moves to particles in a particular sequential order, checkerboard decomposition differs from the serial algorithm in one key way. Particle positions, not labels (or indices), determine the order of updates, so the order will change as particles migrate. Consequently, careless choices can lead to erroneous simulations. We prove that our implementation of MPMC obeys detailed balance to ensure that no incorrect choices are made in its design.

2.1. Checkerboard decomposition

The checkerboard domain decomposition scheme [19,23] divides the simulation volume into sets of square (cubic) cells (see Fig. 2). Checkerboarding maps well to MC simulations because it allows parallel updates of each set, comprising one quarter (one eighth in three dimensions) of the simulation volume. The x and y coordinates of the cell (and z in three dimensions) determine its checkerboard set $Q \in \{a, b, c, d, \ldots\}$:

$$
Q = \begin{cases}
a & \text{if } (x \in \text{Even}) \text{ and } (y \in \text{Even}), \\
b & \text{if } (x \in \text{Odd}) \text{ and } (y \in \text{Even}), \\
c & \text{if } (x \in \text{Even}) \text{ and } (y \in \text{Odd}), \\
d & \text{if } (x \in \text{Odd}) \text{ and } (y \in \text{Odd}), \\
\ldots & \ldots,
\end{cases}
\tag{1}
$$

where $a, b, \ldots$ indicate labels of checkerboard sets.

The width of the cell w must be chosen greater than the diameter of the disk σ (generally, the pair interaction cutoff). At the minimum $w = \sigma$, two particles separated by one cell can move without interacting (see Fig. 2(c)). Thus, the moves available to particles in a cell are independent from those in other cells of the same checkerboard set.

Most previous parallel MC simulations with mobile particles use stripe domain decomposition [24–26], a one-dimensional version of the checkerboard decomposition, which minimizes the interface (and therefore communication) between domains. However, the number of stripes and therefore the number of trial moves that can be conducted in parallel is low. This means stripe decomposition is not efficient for parallelization on more than a few cores.

2.2. Sweep structure

Algorithm 1 outlines the structure of MPMC. It splits each sweep over cells into sub-sweeps (four in two dimensions, eight in three dimensions), one handling each checkerboard set. Line 2 shuffles the order of checkerboard sets using Fisher-

Algorithm 1 Monte Carlo sweep

```
 1: C ← {a, b, c, d, . . .}
 2: rng.shuffle(C)
 3: for Q ∈ C  do                                                        ▷ Loop over sub-sweeps
 4:     for c ∈ cells(Q), in parallel do                                 ▷ Loop over cells
 5:         rng.shuffle(c.particles)
 6:         for s ∈ [0 . . . n_M) do
 7:             p ← c.particles[mod(s, len(c.particles))]
 8:             Generate trial move
 9:             if p remains in cell and move accepted then
10:                 Move p
11:             end if
12:         end for
13:     end for
14: end for
15: d ← rng.uniform(0, w/2)
16: f⃗ ← rng.choose(−x, +x, −y, +y, . . .)
17: shift_cells(f⃗, d)
```

Yates [35] to guarantee a random permutation. During a sub-sweep, the algorithm concurrently processes all of the cells in the active set (line 4).

Each concurrent cell update shuffles and then loops over n_M trial moves (lines 5, 6). Line 7 selects the particle from the cell, repeating from the start of the list when $n_M > n$, where n is the number of particles in the cell. Fixing the number of moves in all cells to the same number distributes computational effort most evenly among GPU cores [36]. Line 8 generates a trial move for each particle. Line 9 accepts the move if it passes the normal Metropolis acceptance criterion [37] and the particle center remains in the cell [24]. Lines 15–17 maintain ergodicity by performing a cell shift, which redraws the cell boundaries in a randomly chosen location.

2.3. Detailed balance

An MC simulation obeys detailed balance if, for every internal process evolving the system there exists a reverse process occurring at the same rate. This ensures that a sequence of configurations converges to the correct equilibrium distribution, regardless of the initial condition. Markov-chain MC generates a sequence of configurations where the probability $x_j(t+1)$ of observing the system in state j at step $t+1$ is determined only by the previous state i at step t. This can be expressed by

$$x_j(t+1) = x_i(t)P_{ij}, \tag{2}$$

where $\boldsymbol{x}(t) = \{x_1(t), x_2(t), \ldots, x_n(t)\}$ is the probability distribution at step t. The elements P_{ij} of the transition matrix represent the probabilities that the system will transition from state i to state j. If there exists an equilibrium distribution of states $\boldsymbol{x}^*$ for which $\boldsymbol{x}^* = \boldsymbol{x}^* P$, then $\boldsymbol{x}(t)$ is guaranteed to converge to $\boldsymbol{x}^*$ as $t \to \infty$ when P satisfies detailed balance:

$$x_i^* P_{ij} = x_j^* P_{ji}. \tag{3}$$

Although detailed balance of a Markov chain is a sufficient condition to ensure convergence, it is not necessary. Manousiouthakis and Deem show that an irreducible transition matrix that enforces regular sampling ($\exists m : (P^m)_{ij} > 0 \ \forall i, j$) and obeys balance ($\boldsymbol{x}^* = \boldsymbol{x}^* P$) is both necessary and sufficient for convergence to the correct equilibrium distribution [38]. In this work we choose to enforce detailed balance.

The MPMC algorithm constructs a Markov chain and obeys detailed balance on the level of an MC sweep. This follows directly from the observation that for each sweep there is exactly one inverse sweep, which can be seen as follows. Take a particular sequence of sub-sweeps and a sequence of n_M moves (either accepted or rejected) within each cell. The reverse sweep consists of the reverse sequence of sub-sweeps and the reverse sequence of moves within each cell, with each move following the negative of the original vector. For example, with $n_M = 6$ trial moves per cell and $n = 4$ particles in the cell, the original sequence would be $[0, 1, 2, 3, 0, 1]$. There is exactly one particle shuffling, $[1, 0, 3, 2]$, that generates the reverse sequence $[1, 0, 3, 2, 1, 0]$. Since each sequence is chosen randomly from all possible permutations, the forward and reverse sequences occur with equal probability, and thus detailed balance holds.

2.4. Pitfalls leading to incorrect statistical sampling

Detailed balance is ensured when the following three steps are in place.

1. *The particle center must not leave the cell* [24]. If particles are allowed to leave their cells during a sub-sweep, the reverse sequence of moves cannot be generated and detailed balance is not ensured. When we skip this restriction in the hard disk system, it always develops order in the same orientation and at a lower than expected density.

Table 1
Simulations of $N = 256^2$ hard disks with particle (P.) and checkerboard (CB.) shuffling enabled or disabled. The comparison of equilibrium pressures P^* for runs at three different packing fractions ϕ demonstrates the necessity to shuffle particles. Checkerboard shuffling is not required to obtain correct results.

Shuffling		P^* in the hard disk system at		
P.	CB.	$\phi = 0.698$	$\phi = 0.708$	$\phi = 0.716$
Yes	Yes	9.17079(5)	9.18214(6)	9.1774(2)
Yes	No	9.1707(1)	9.1821(2)	9.1775(3)
No	Yes	9.1716(1)	9.1876(2)	9.1831(3)
No	No	9.1715(1)	9.1873(1)	9.1823(2)

2. *Shuffling the particles in each cell.* Particles entering a cell are added at the end of the cell list and cell lists are partially maintained during a cell shift. When we skip particle shuffling, a temporal memory of previous states builds up over many sweeps, violating the Markov property.
3. *Shuffling the checkerboard set.* Without shuffling of the checkerboard set, the reverse sweep cannot be generated, which violates the condition of detailed balance.

To increase the number of accepted moves per sweep one might be tempted to allow particles to leave the cell, compensating the violation of step (1) by ensuring that each particle moves exactly once per sweep. However, this procedure does not guarantee balance. Cell updates with moves that leave a cell in the 'middle' of the cell update (i.e. not the first or last successful move of the sub-sweep into a given neighboring cell) are not reversible and generate an incorrect probability distribution. When we apply this scheme in the hard disk system, the pressure is shifted slightly away from the correct value and the magnitude of shift depends on the maximum trial move distance.

The particle shuffling step (2) is often explicitly omitted in favor of sequential updating [20,22,25,26]. As a justification, these authors refer to the analysis of Manousiouthakis and Deem [38], who showed that shuffling is not necessary for the Ising lattice model away from infinite temperature. However, their analysis cannot be transferred to systems of mobile particles if the sequence of particles is determined dynamically. Our simulations for hard disks confirm (Table 1) that skipping the particle shuffling step alters the pressure close to the melting transition. In contrast, the checkerboard set shuffling step (3) is not necessary for correct sampling.

3. Implementation

We implement MPMC for hard disks using the NVIDIA CUDA programming model and execute benchmarks on a Tesla K20 graphics processor. CUDA is an established parallel programming language; details may be found in the CUDA programming guide [39], text books [40–42] or in other publications, Refs. [2,3] for example. The pseudocode presented in this paper is general enough that it could be adapted to any data-parallel language (e.g. OpenCL or OpenMP).

3.1. Data structures

The proper choice of data structures can make or break an implementation's performance. We keep all of the architecture details of NVIDIA GPUs in mind when designing our implementation. MD codes store particles in a flat array with N elements, and auxiliary data structures indirectly reference this list by index [3,5,11,12]. Such a data structure is not appropriate for MC as it would be expensive to rebuild the cell list every sweep.

Instead, we store the particle positions directly in the cell list. That data is a sparse flat array, $\texttt{disk}[x, y, i]$, with storage for $m \cdot m \cdot n_{max}$ particle positions, where m is the number of cells on the side of the simulation box and n_{max} is the maximum number of particles allowed in a cell. The auxiliary array, $\texttt{n}[x, y]$ stores the number of particles in each cell, where the particles are placed in elements $i \in [0 \ldots n)$.

We minimize the number of overlap checks and maximize parallelism by setting the cell width w small, but not so small that a large fraction of moves will cross the cell boundaries. The size must also be chosen so that n_{max} is known. Fig. 3 shows that the largest cell that can fit no more than four particles has a width $w < \sqrt{2}\sigma$. Furthermore, we avoid expensive boundary condition checks by choosing m as a multiple of 2 times the block size, because each thread handles every other cell. For the 32-thread blocks used here, we set $m = \lfloor L/(\sqrt{2}\sigma) \rfloor$ and then round up to the nearest multiple of 64.

Each cell has a local coordinate system to mitigate floating point cancellation errors that would otherwise occur for large absolute coordinate values. Whenever difference vectors are computed a coordinate system transformation is needed. The components of the translation vector are either $+w$, $-w$, or 0 depending on the relative location of the neighboring cell.

3.2. Kernel

Algorithm 2 implements the MPMC sub-sweep (lines 4–13 of Algorithm 1) update in a CUDA kernel. One thread is launched for each cell in the active checkerboard set. Lines 1 and 2 compute the (x, y) index of the cell to which the thread is assigned. Rows in a thread block handle rows in the cell data. For example, threads with x ids 0, 1, 2, 3 are assigned to

Fig. 3. Four disks of diameter σ are placed on the corners of the cell and a fifth in the center. The smallest cell that can contain five disk centers has a diagonal of 2σ and an edge length of $w_5 = \sqrt{2}\sigma$. Thus, the largest cell that can contain a maximum of four disk centers has a width $w < \sqrt{2}\sigma$.

Algorithm 2 Sub-sweep GPU kernel

Require: gdim is $(m/\text{bdim.x}/2, m/\text{bdim.y}/2)$
Require: (off.x, off.y) is the offset to the lower-leftmost active cell in the sub sweep
Require: seed is a random number seed chosen by the user and fixed for the duration of a run
Require: sweep is the index of the current sweep
Require: $\vec{D}_{\text{sh}}[p]$ is stored in shared memory such that each thread indexes unique elements in memory for $p \in [0 \ldots n_{\max})$
1: $x \leftarrow 2(\text{bidx.x} \cdot \text{bdim.x} + \text{tidx.x}) + \text{off.x}$
2: $y \leftarrow 2(\text{bidx.y} \cdot \text{bdim.y} + \text{tidx.y}) + \text{off.y}$
3: $\text{rng} \leftarrow \textbf{Saru}(mx + y, \text{step}, \text{seed})$
4: $n \Leftarrow \text{n}[x, y]$
5: **if** $n == 0$ **then**
6:     **return**
7: **end if**
8: $\vec{D}_{\text{sh}}[i] \Leftarrow \text{disk}[x, y, i] \ \forall i \in [0 \ldots n_{\max})$
9: $\textbf{rng.shuffle}(\vec{D}_{\text{sh}}[0 \ldots n])$
10: $i \leftarrow 0$
11: **for** $s \in [0 \ldots n_M)$ **do**
12:     $\vec{D}_{\text{move}} \leftarrow \vec{D}_{\text{sh}}[i] + \textbf{rng.inCircle}(d)$
13:     $\text{overlap} \leftarrow \textbf{False}$
14:     **for** $(x_{\text{neigh}}, y_{\text{neigh}}) \in$ neighborhood of cell (x, y) **do**
15:         $\vec{s} \leftarrow$ vector pointing to current neighbor
16:         **for** $j \in [0 \ldots n_{\max})$ **do**
17:             **continue** when $(x_{\text{neigh}}, y_{\text{neigh}}) = (x, y) \wedge i = j$
18:             $\vec{D} \Leftarrow \text{disk}[x_{\text{neigh}}, y_{\text{neigh}}, j]$
19:             **if** $|\vec{D}_{\text{move}} - (\vec{D} + \vec{s})| < \sigma$ **then**
20:                 $\text{overlap} \leftarrow \textbf{True}$
21:             **end if**
22:         **end for**
23:     **end for**
24:     **if** $\neg \text{overlap}$ **then**
25:         **if** $\vec{D}_{\text{move}} \in$ cell (x, y) **then**
26:             $\vec{D}_{\text{sh}}[i] \leftarrow \vec{D}_{\text{move}}$
27:         **end if**
28:     **end if**
29:     $i \leftarrow i + 1$
30:     **if** $i \geqslant n$ **then**
31:         $i \leftarrow 0$
32:     **end if**
33: **end for**
34: $\vec{D}_{\text{sh}}[i] \Rightarrow \text{disk}[x, y, i] \ \forall i \in [0 \ldots n_{\max})$

Algorithm 3 Index cell data

1: **procedure** CELL_INDEX(x, y, i)
2:     **if** $x \in$ Odd **then**
3:         $q \leftarrow (x + m)/2$
4:     **else**
5:         $q \leftarrow x/2$
6:     **end if**
7:     **return** $(i \cdot m + y) \cdot m + q$
8: **end procedure**

---

cells with x coordinates 0, 2, 4, 6 relative to some offset. Similarly, each row in a thread block is assigned to alternating y rows in the cell data.

Each thread initializes its own random number stream (line 3). We use the Saru PRNG, developed by Steve Worley [43], to create uncorrelated random number streams from a hash of the thread index, current step index and a user chosen

---

Algorithm 4 Cell shift GPU kernel

Require: gdim is $(m/\text{bdim.x}, m/\text{bdim.y})$
Require: $\vec{D}_{\text{sh}}[p]$ is stored in shared memory such that each thread indexes unique elements in memory for $p \in [0 \ldots n_{\max}]$
1: **procedure** SHIFT_CELLS($\vec{f}, d$)
2:     $x \leftarrow (\text{bidx.x} \cdot \text{bdim.x})$
3:     $y \leftarrow (\text{bidx.y} \cdot \text{bdim.y})$
4:     $n_{\text{current}} \Leftarrow \text{n}[x, y]$
5:     $\vec{D}_{\text{sh}}[i] \leftarrow (-10, -10) \; \forall i \in [0 \ldots n_{\max}]$
6:     $n_{\text{new}} \leftarrow 0$
7:     **for** $i \in [0 \ldots n_{\text{current}})$ **do**
8:         $\vec{D} \Leftarrow \text{disk}[x, y, i]$
9:         $\vec{D} \leftarrow \vec{D} - \vec{f} \cdot d$
10:        **if** $D.x > 0 \wedge D.y > 0 \wedge D.x \leqslant w \wedge D.y \leqslant w$ **then**
11:            $\vec{D}_{\text{sh}}[n_{\text{new}}] \leftarrow \vec{D}$
12:            $n_{\text{new}} \leftarrow n_{\text{new}} + 1$
13:        **end if**
14:    **end for**
15:    $(x_{\text{neigh}}, y_{\text{neigh}}) \leftarrow$ cell in direction of $\vec{f}$
16:    $\vec{s} \leftarrow$ vector pointing to neighbor
17:    $n_{\text{neigh}} \Leftarrow \text{n}[x_{\text{neigh}}, y_{\text{neigh}}]$
18:    **for** $i \in [0 \ldots n_{\text{neigh}})$ **do**
19:        $\vec{D} \Leftarrow \text{disk}[x_{\text{neigh}}, y_{\text{neigh}}, i]$
20:        $\vec{D} \leftarrow \vec{D} - \vec{f} \cdot d$
21:        **if** $D.x > 0 \wedge D.y > 0 \wedge D.x \leqslant w \wedge D.y \leqslant w$ **then**
22:   ▷ Particle stays in neighbor cell, do nothing
23:        **else**
24:            $\vec{D} \leftarrow \vec{D} + \vec{s}$
25:            $\vec{D}_{\text{sh}}[n_{\text{new}}] \leftarrow \vec{D}$
26:            $n_{\text{new}} \leftarrow n_{\text{new}} + 1$
27:        **end if**
28:    **end for**
29:    $\vec{D}_{\text{sh}}[i] \Rightarrow \text{disk\_dbl}[x, y, i] \; \forall i \in [0 \ldots n_{\max}]$
30:    $n_{\text{new}} \Rightarrow \text{n\_dbl}[x, y]$
31: **end procedure**

---

seed. NVIDIA's CURAND library is an alternative, but requires reading and writing a large state in each thread, which slows performance by 30%. See Ref. [44] for more details on the tradeoffs of various parallel PRNG schemes.

Lines 4–9 read the assigned cell into shared memory and shuffle the particles. Line 11 starts a loop over n_M trial moves. For each selected particle i in sequence, line 12 generates the trial move and lines 13–23 check for any overlaps with particles in the neighboring cells. Lines 24–28 update the particle to its new position if the move generates no overlaps and remains in the cell. Lines 39–32 wrap i back to the start of the cell when the end is reached.

When checking overlaps, each thread reads the eight neighboring cells. Typical GPU kernels with this memory access pattern use shared memory as a managed cache to avoid multiple reads from the same cell. In our case, the data size per cell is large and would occupy a substantial fraction of the available shared memory, limiting parallelism. Instead, we read only the current cell into shared memory (line 8) and use hardware cached reads for the neighbor accesses (line 18).

A 128-byte wide cache line in K20 fits four full cells in a row. However, a row-major assignment of $[x, y, i]$ to a linear index is not ideal. Only one particle would be read at a time from alternating cells in a row, using 8 out of the 128 bytes in a delivered cache line. A carefully chosen mapping from $[x, y, i]$ indices to linear memory addresses leads to full utilization. Algorithm 3 implements that mapping. First, i is the slowest index so that memory instructions in loops over i in the kernel read contiguous data. The next fastest index is y which is handled in the traditional manner. The x index is the fastest, but it is rearranged so that all the odd, and similarly even, x values are contiguous in the linear space. We achieve further performance improvements by using texture reads (tex1Dfetch) in place of all global memory loads. On K20, the texture cache provides the most throughput for reused read-only data.

Thanks to these efforts, we achieve excellent utilization of the available memory bandwidth. Benchmarks with the NVIDIA visual profiler show a sustained bandwidth of $\sim$200 GB/s out of the texture cache. Achieving high occupancy and limiting divergence are just as important. For example, selecting the 16k/48k (L1/shared) mode increases occupancy and boosts performance by 50% compared to the 48k/16k mode. The loop on line 16 goes from 0 to $n_{\max}$ to boost performance by reducing divergent branches compared to looping over the number of particles currently in the cell. We set the values of the empty particle slots to $(-10, -10)$ so that they do not result in false overlaps. Early exit conditions upon finding the first overlap (not shown) cause additional divergent branches, but removing these checks reduces performance due to the increase in computations and memory accesses.

Kernel performance varies with block size [3]. Short benchmarks show that (32, 4) is the fastest, and we use it for all production runs. It outperforms the slowest by 55%, demonstrating the importance of performing this test.

Algorithm 4 implements the cell shift step on the GPU by equivalently translating the particles in the opposite direction. One thread per cell gathers all of the particles that belong in the new cell and builds the list in shared memory. It then writes out the cell to a separate memory area, disk_dbl and n_dbl, so that other running threads do not read updated

data. After the kernel completes, the double buffered data structures are swapped. Since the shift direction is one of either $+x$, $-x$, $+y$, or $-y$, only two old cells contribute particles to the new cell: the cell with the same index, and one neighbor. Lines 7–14 loop over the current cell, read in each particle, shift the cell, and if that particle is still within the cell boundaries it is added to the new list. Lines 18–28 perform the same operations on the neighbor cell, with the addition of a coordinate system transformation. These rely on the following logic: if the particle left its host cell, it must have entered this cell. In this manner, each particle is checked for inclusion by two threads and at two separate points in the code.

The floating point operations by both of these threads *must* be identical. Consider if line 20 were to perform the co-ordinate system translation $\vec{D} - \vec{f} \cdot d + \vec{s}$ and then check for particles that enter the current cell. Floating point round-off errors may result in both this check and the corresponding check on line 10 to fail, losing the particle. In a test simulation configured with 1024^2 particles, several hundred were lost after 10^6 sweeps. When implemented as shown in Algorithm 4, no particles are lost even after 10^9 sweeps.

3.3. Parameter tuning

The maximum move radius d and the number of trial moves performed in each cell update n_M are free parameters. At fixed n_M, we test $d \in [0.06, 0.08, \ldots, 0.20]$ and at fixed d, we test $n_M \in [1 \ldots 8]$. Each test measures the autocorrelation of the average orientational order parameter[45] over a long run of 10^9 sweeps for $N = 512^2$. We find that $d = 0.16$ and $n_M = 4$ minimize the autocorrelation time τ when measured in wall clock seconds.

4. Results

The hard disk system is a standard model system in statistical mechanics and the one which was originally used to pioneer the Monte Carlo computer simulation method [37]. Its phase behavior is completely determined by the equation of state, which is the relation between internal pressure P and packing fraction $\phi = \rho \pi (\sigma/2)^2$. Here, $\rho = N/V$ is density and V the volume of the simulation box. At packing fractions between $\phi = 0.7$ to 0.72 the system undergoes a first-order phase transition from the liquid phase to the hexatic phase, followed by a continuous transition to the solid phase [34]. In their paper, Bernard and Krauth showed with an event-chain simulation method [18] that only very large simulations ($N > 256^2$) have minimal finite size effects, and long equilibration times are necessary. This demonstrates the need for high performance MC code and explains why previous simulation studies of the hard disk system [46–51] have been unable to provide conclusive evidence for a first-order phase transition or the existence of an intermediate hexatic phase.

Phase transitions are extremely sensitive to the slightest programming error or inadvertent correlation of trial moves due to the appearance of (quasi-)long-range spatial correlations and long equilibration times. Structural fluctuations are more important in two dimensions than in three dimensions, and they are particularly large for the hard disk system. These properties, together with the availability of high-precision serial data to compare with, make hard disks a good system to test our algorithm.

In principal, computing the pressure for hard disks is simple: Estimate the radial distribution function $g(r)$ in the limit as r approaches the disk diameter σ from the right, and calculate [37]

$$P^* = \frac{P\sigma^2}{k_B T} = \sigma^2 \rho \left(1 + \frac{\pi}{2} \sigma^2 \rho \lim_{r \to \sigma+} g(r)\right), \tag{4}$$

where we introduce the dimensionless pressure P^*. In practice, obtaining an unbiased estimate requires special care. We estimate pressure using the following procedure. For each sampled configuration, a histogram of particle pair distances is computed over all N particles. $n[r_i]$ counts the number of particle pairs between r_i and $r_i + \delta r$. The pair distribution function $g(r)$ is evaluated by the equation

$$g(R_i) = \frac{n[r_i]}{N\rho \, \delta A} = \frac{n[r_i]}{N\rho \, 2\pi R_i \, \delta r} \tag{5}$$

at the sampling points $r = R_i$, where

$$R_i = \frac{2}{3} \frac{r_{i+1}^3 - r_i^3}{r_{i+1}^2 - r_i^2}. \tag{6}$$

The formula for R_i is derived assuming a linear dependence of $n(r)$ with r, which we observe to be valid close to σ. The minimum position R_0 is greater than σ, so direct evaluation of the limit is not possible. We fit $g(R_i)$ to a polynomial of degree d in the range $r \in (\sigma, \sigma + c]$ and extrapolate to $r = \sigma$. Extensive testing with a model distribution ($n(r) = 30e^{-30r}$) tunes the parameters to ensure that there is no systematic bias. We choose parameters in the middle of the flat region with systematic errors less than 10^{-5}. They are $\delta r = 10^{-4}\sigma$, $c = 0.02\sigma$, and $d = 5$.

Table 2 shows the average P^* data obtained by simulations with MPMC. Independent runs of $N = 256^2$, $N = 512^2$, and $N = 1024^2$ particles are performed at packing fractions between $\phi = 0.698$ and $\phi = 0.718$, which comprises the transformation from liquid to solid. Each run starts with a randomly generated configuration at low density and is quickly compressed to the target. The run then continues at constant density for 10^9 sweeps, which only takes 4 days for $N = 256^2$ (15 days for

Table 2

This table shows data for the equation of state $P^*(\phi)$ over the range where the liquid transforms into the solid. It includes runs by MPMC (this work) and by Bernard and Krauth (BK) with serial event chain simulation [34]. Error bars are shown at two standard errors of the mean, $\sigma = 2(\langle[P^* - \langle P^*\rangle]^2\rangle/N_{\text{samples}})^{1/2}$, where the number in parentheses is the error in the last digit shown. Our data is averaged over 8 independent runs of 10^9 sweeps (64 runs for $\phi = 0.718$, $N = 256^2$). See Ref. [34] for a description of the BK data averaging scheme. Differences in the pressures are shown with propagated error bars $(\sigma_{\text{MPMC}}^2 + \sigma_{\text{BK}}^2)^{1/2}$.

System size	Method	Dimensionless pressure P^* in the hard disk system at packing fraction					
		$\phi = 0.698$	$\phi = 0.702$	$\phi = 0.706$	$\phi = 0.710$	$\phi = 0.714$	$\phi = 0.718$
$N = 256^2$	MPMC	9.1709(1)	9.1920(2)	9.1854(1)	9.1792(1)	9.1758(1)	9.187(1)
	BK	9.1708(4)	9.1924(4)	9.1858(5)	9.1790(4)	9.1758(5)	9.186(1)
	Difference	0.0000(5)	0.0004(5)	0.0004(5)	0.0002(4)	0.0000(6)	0.001(1)
$N = 512^2$	MPMC	9.1699(5)	9.1900(2)	9.1861(1)	9.1828(1)	9.1800(1)	9.1930(4)
	BK	9.1700(2)	9.1899(6)	9.1856(6)	9.1821(5)	9.1803(4)	9.1937(2)
	Difference	0.0001(5)	0.0001(6)	0.0004(6)	0.0007(5)	0.0003(4)	0.0006(5)
$N = 1024^2$	MPMC	9.16934(4)	9.1882(3)	9.1859(3)	9.1842(3)	9.1819(4)	9.1951(4)
	BK	9.1693(1)	9.1880(2)	9.1855(2)	9.1843(2)	9.1822(2)	9.1949(3)
	Difference	0.0000(1)	0.0002(4)	0.0003(4)	0.0001(4)	0.0002(5)	0.0001(5)

Table 3

Benchmark hardware and software configurations.

	K20	E5540
Hardware	custom built	HP DL2x170h
Mainboard	ASUS Sabertooth	
Chipset	AMD 990FX	Intel 5520
CPU	Athlon II X4 630	2x Intel Xeon E5540
CPU clock	2.8 GHz	2.53 GHz
RAM	16 GB DDR3	24 GB DDR3
RAM clock	1333 MHz	1333 MHz
GPU	Tesla K20	
Core clock	706 MHz	
Memory clock	5200 MHz	
DRAM	5 GB GDDR5	
ECC	off	
OS	Gentoo	RHEL 6
Architecture	x86_64	x86_64
CPU compiler	GCC 4.5.3	GCC 4.7.0
flags	-O3 -funroll-loops	-O3 -funroll-loops
GPU compiler	CUDA 5.0	
flags	*default*	
GPU driver	310.19	

$N = 512^2$) to complete on a Tesla M2070 GPU. The pressure is averaged every 200 sweeps in each run after an equilibration period of $3 \cdot 10^8$ sweeps.

Our data confirms the equation of state reported in Ref. [34]. All values overlap within error bars. We do not analyze positional order or orientation order in the dense phase emerging from the phase transition. Such an analysis would be necessary to distinguish a hexatic phase from a solid phase and is left for a separate work [45]. We refer to Ref. [45] for an in-depth comparison of the phase diagram of hard disks obtained with various algorithms, including MPMC.

5. Performance

We test the performance of the MPMC hard disk code using CUDA on a single Tesla K20 GPU (Kepler). For comparison, we also implement the MPMC algorithm on the CPU and run it on our cluster nodes. The CPU implementation uses OpenMP to parallelize across all cores on a node. Each thread processes a single horizontal strip of the simulation domain and the innermost loop follows Algorithm 2. We make no attempt to parallelize across multiple GPUs or multiple CPU nodes using MPI. All tests are performed using single precision floating point format to store particle coordinates. Table 3 lists complete specifications of the hardware and software configurations of our test machines.

5.1. Scaling with number of particles

We perform benchmarks at a fixed packing fraction $\phi = 0.698$ and analyze the performance scaling with N. Fig. 4(a) plots the results and Table 4 lists selected numerical values.

Fig. 4. Benchmarks are performed with a varying number of disks, N, at a packing fraction $\phi = 0.698$. Each benchmark runs 100 sweeps to warm up and then measures the time it takes to run another 100 sweeps as well as the number of attempted trial moves to compute $\eta = N_{\text{moves}}/t$. 11 separate runs are performed and the median result is plotted here. In all cases, the error bars (one standard deviation) are smaller than the symbol size. Subfigure (a) plots the efficiency of the computation η vs. N on various hardware configurations. Subfigure (b) plots η normalized by the maximum obtained on each hardware configuration. Subfigure (c) plots the speedup obtained over a serial execution. The dashed line marks a speedup of 8.

Table 4
Numerical values for selected benchmarks from Fig. 4.

		253^2	760^2	1520^2	3040^2
K20	η	$6.23 \cdot 10^8$	$9.83 \cdot 10^8$	$1.09 \cdot 10^9$	$1.09 \cdot 10^9$
	$\eta/\max(\eta)$	0.567	0.895	1.00	1.00
	$\eta/\eta_{\text{serial}}$	83.6	134	148	148
E5540	η	$5.77 \cdot 10^7$	$5.78 \cdot 10^7$	$5.42 \cdot 10^7$	$5.26 \cdot 10^7$
(8 cores)	$\eta/\max(\eta)$	0.900	0.900	0.860	0.820
	$\eta/\eta_{\text{serial}}$	7.65	7.88	7.39	7.16
E5540	η	$7.54 \cdot 10^6$	$7.33 \cdot 10^6$	$7.33 \cdot 10^6$	$7.35 \cdot 10^6$
(1 core)	$\eta/\max(\eta)$	0.880	0.855	0.855	0.858

The algorithm has a running time $t \in O(N)$, so the number of trial moves per unit time, η, should be constant under ideal circumstances. In practice, there are overheads that cause deviation from constant efficiency. Fig. 4(b) collapses the individual η plots to a relative efficiency metric for each separate configuration. When running on a single CPU core, efficiency is flat with only a slight downward drift for large N. This is likely because the largest systems no longer fit in on-chip cache. The 8-core benchmarks show the same behavior at large N, although it starts a factor of 2 higher because the same data is now split over 2 chips' caches. The 8-core results also have a slight dip for $N < 10^5$ due to the overhead of managing worker threads. On the GPU, the kernel launch overhead is significant enough that for N less than 10^5, efficiency is less than 70%.

Despite this inefficiency, the GPU still outperforms the single CPU core runs by a factor of 83 for systems as small as $N = 253^2$. At peak efficiency, the GPU speedup over a single core is a factor of 148. We prefer thinking about speedups compared to a single core, but recognize that there are other ways of evaluating it. On a per socket basis, a single GPU is 37x faster than a quad core E5540. On a per-node basis, it is still an order of magnitude (20x) faster. In terms of aggregate performance per price, an 8-core node with 8 externally attached GPUs (the configuration we use) is 148x faster than just the host, but only costs 5.5x as much for a benefit of 27x more sweeps/unit time given a fixed budget. We do not have power monitoring equipment on our cluster, so we are unable to provide actual measurements of energy savings. Instead, we obtain an estimate using the manufacturer's TDP specifications, which report 80 W for the E5540 and 225 W for the K20. Based on these numbers and the per-socket speedup, a CPU simulation would use 13 times more energy than if it were run for the same number of sweeps on the GPU.

5.2. Limitations

The number of parallel threads depends on the cell size, interaction range, and particle shape. In this work, MPMC executes one parallel thread per active cell. In terms of the N particles, the number of active cells is $N/(\langle n \rangle 2^d)$ where $\langle n \rangle$ is the average cell population and d is the dimensionality of the system. It is approximately $N/8$ for hard disks ($d = 2$) at high density ($\langle n \rangle = 2$). Hard spheres ($d = 3$) decrease the number of active cells to $N/16$ ($\langle n \rangle = 2$). Expanding the interaction range to a truncated and shifted Lennard-Jones potential with $r_{\text{cut}} = 2.5\sigma$ in three dimensions drops the number to $N/144$ ($\langle n \rangle = 18$ determined by MD simulation). GPUs operate at peak efficiency only when running more than 10 thousand threads, establishing minimum practical system sizes of 80 thousand, 160 thousand, and 1.44 million particles for hard disks, hard spheres, and Lennard-Jones beads, respectively. Depending on the desired application, these sizes may be unnecessary or prohibitively large.

Alternate thread assignment schemes for MPMC can ameliorate this problem. One possibility is to execute 3^d threads per active cell, where each computes potential overlaps between the trial moves in the active cell with particles in one neighboring cell. This technique greatly increases the number of parallel threads, reducing the theoretical minimum practical system sizes to 9 thousand, 6 thousand, and 53 thousand for hard disks, hard spheres, and Lennard-Jones beads, respectively. These sizes are much more in line with typical simulation sizes in soft matter research.

Large density fluctuations present in Lennard-Jones systems should not pose a problem. Thread divergence is not increased because the loop over n_M trial moves in each active cell is identical across the entire system. The GPU's fine-grained scheduler replaces blocks as soon as they finish executing, so the load is automatically balanced despite blocks in low density regions completing sooner than those in high density areas. Memory is wasted in the sparse cell data structure, but modern GPUs have large amounts of memory (6 GB), enough to store tens of millions of particles even with wasted space. Alternate data structures could always be employed.

6. Conclusions

Efficient parallel algorithms are essential for the application of Monte Carlo particle simulations on current and future computer hardware. Building on prior works utilizing checkerboard domain decomposition, we detailed an algorithm for massively parallel MC and implemented it on the GPU and the CPU. The GPU speedup that we obtain is comparable to what has been achieved in MD simulations. Our findings demonstrate that GPUs are well-suited for running large-scale MC simulations. Medium scale simulations could also make good use of the GPU after some work to increase parallelism in the implementation.

Future work to parallelize on multiple GPUs using MPI will enable even larger scale simulations. However, a straightforward implementation of the MPMC algorithm onto multiple GPUs poses one major problem. Communication with neighboring MPI ranks is necessary after every sub-sweep update, so slow PCIe interconnect bandwidths would limit performance. A second level of parallelization that introduces inactive regions between the MPI ranks significantly decreases communication needs. Such a technique only requires communication at the end of a sweep when the grid shift is performed. The number of trial moves attempted per grid cell n_M could be further increased to lessen communication even more.

During the course of this work we learned that it is surprisingly difficult to implement MC in a parallel environment. Every slight violation of the balance conditions leads to incorrect sampling, and with parallel update moves it is not simple to determine which schemes do not obey balance. Sometimes the effect on our simulations was so small that it would be easy to miss in a typical complex practical application. It is important to test any new parallel algorithm thoroughly in a situation where reliable results are available from serial simulations. Hard disks are such a system. We computed to high precision the hard disk equation of state over the density range where the liquid transforms into the solid. Our results agree perfectly with serial event-chain Monte Carlo simulations and are compatible with the presence of a first-order liquid-hexatic phase transition [45] − a finding that is scientifically significant by itself.

Acknowledgements

References

[1] K. Asanovic, R. Bodik, B.C. Catanzaro, J.J. Gebis, P. Husbands, K. Keutzer, D.A. Patterson, W.L. Plishker, J. Shalf, S.W. Williams, K.A. Yelick, The landscape of parallel computing research: a view from Berkeley, Technical Report UCB/EECS-2006-183, EECS Department, University of California, Berkeley, 2006.

[2] J.E. Stone, D.J. Hardy, I.S. Ufimtsev, K. Schulten, GPU-accelerated molecular modeling coming of age, Journal of Molecular Graphics & Modelling 29 (2010) 116–125.

[3] J.A. Anderson, C.D. Lorenz, A. Travesset, General purpose molecular dynamics simulations fully implemented on graphics processing units, Journal of Computational Physics 227 (2008) 5342–5359.

[4] HOOMD-blue, http://codeblue.umich.edu/hoomd-blue, 2012.

[5] W.M. Brown, P. Wang, S.J. Plimpton, A.N. Tharrington, Implementing molecular dynamics on hybrid high performance computers − short range forces, Computer Physics Communications 182 (2011) 898–911.

[6] A.W. Götz, M.J. Williamson, D. Xu, D. Poole, S. Le Grand, R.C. Walker, Routine microsecond molecular dynamics simulations with AMBER on GPUs. 1. Generalized born, Journal of Chemical Theory and Computation 8 (2012) 1542–1555.

[7] S. Le Grand, A.W. Götz, R.C. Walker, SPFP: Speed without compromise − A mixed precision model for GPU accelerated molecular dynamics simulations, Computer Physics Communications 184 (2) (2013) 374–380, http://dx.doi.org/10.1016/j.cpc.2012.09.022.

[8] J.E. Stone, J.C. Phillips, P.L. Freddolino, D.J. Hardy, L.G. Trabuco, K. Schulten, Accelerating molecular modeling applications with graphics processors, Journal of Computational Chemistry 28 (2007) 2618–2640.

[9] P. Eastman, V.S. Pande, Efficient nonbonded interactions for molecular dynamics on a graphics processing unit, Journal of Computational Chemistry 31 (2010) 1268–1272.

[10] N. Ganesan, B.A. Bauer, T.R. Lucas, S. Patel, M. Taufer, Structural, dynamic, and electrostatic properties of fully hydrated DMPC bilayers from molecular dynamics simulations accelerated with graphical processing units (GPUs), Journal of Computational Chemistry 32 (2011) 2958–2973.

[11] P.H. Colberg, F. Höfling, Highly accelerated simulations of glassy dynamics using GPUs: Caveats on limited floating-point precision, Computer Physics Communications 182 (2011) 1120–1129.

[12] D.C. Rapaport, Enhanced molecular dynamics performance with a programmable graphics processor, Computer Physics Communications 182 (2011) 926–934.

[13] GROMACS, http://www.gromacs.org/, 2013.

[14] ACEMD, http://www.acellera.com/, 2013.

[15] R. Swendsen, J.-s. Wang, Nonuniversal critical dynamics in Monte Carlo simulations, Physical Review Letters 58 (1987) 86–88.

[16] J. Liu, E. Luijten, Rejection-free geometric cluster algorithm for complex fluids, Physical Review Letters 92 (2004) 1–4.

[17] S. Whitelam, P.L. Geissler, Avoiding unphysical kinetic traps in Monte Carlo simulations of strongly attractive particles, Journal of Chemical Physics 127 (2007) 154101.

[18] E. Bernard, W. Krauth, D. Wilson, Event-chain Monte Carlo algorithms for hard-sphere systems, Physical Review E 80 (2009) 5–9.

[19] G. Pawley, K. Bowler, R. Kenway, D. Wallace, Concurrency and parallelism in MC and MD simulations in physics, Computer Physics Communications 37 (1985) 251–260.

[20] R. Ren, G. Orkoulas, Acceleration of Markov chain Monte Carlo simulations through sequential updating, Journal of Chemical Physics 124 (2006) 64109.

[21] T. Preis, P. Virnau, W. Paul, J.J. Schneider, GPU accelerated Monte Carlo simulation of the 2D and 3D Ising model, Journal of Computational Physics 228 (2009) 4468–4477.

[22] T. Levy, G. Cohen, E. Rabani, Simulating lattice spin models on graphics processing units, Journal of Chemical Theory and Computation 6 (2010) 3293–3301.

[23] G.S. Heffelfinger, M.E. Lewitt, A comparison between two massively parallel algorithms for Monte Carlo computer simulation: An investigation in the grand canonical ensemble, Journal of Computational Chemistry 17 (1996) 250–265.

[24] A. Uhlherr, S.J. Leak, N.E. Adam, P.E. Nyberg, M. Doxastakis, V.G. Mavrantzas, D.N. Theodorou, Large scale atomistic polymer simulations using Monte Carlo methods for parallel vector processors, Computer Physics Communications 144 (2002) 1–22.

[25] R. Ren, G. Orkoulas, Parallel Markov chain Monte Carlo simulations, Journal of Chemical Physics 126 (2007) 211102.

[26] C.J. O'Keeffe, G. Orkoulas, Parallel canonical Monte Carlo simulations through sequential updating of particles, Journal of Chemical Physics 130 (2009) 134109.

[27] B. Sadigh, P. Erhart, A. Stukowski, A. Caro, E. Martinez, L. Zepeda-Ruiz, Scalable parallel Monte Carlo algorithm for atomistic simulations of precipitation in alloys, Physical Review B 85 (2012) 1–11.

[28] B.D. Lubachevsky, Efficient parallel simulations of asynchronous cellular arrays, Complex Systems 1 (1987) 1099–1123.

[29] G. Korniss, M. Novotny, P. Rikvold, Parallelization of a dynamic Monte Carlo algorithm: a partially rejection-free conservative approach, Journal of Computational Physics 153 (1999) 488–508.

[30] E. Martínez, J. Marian, M. Kalos, J. Perlado, Synchronous parallel kinetic Monte Carlo for continuum diffusion–reaction systems, Journal of Computational Physics 227 (2008) 3804–3823.

[31] G. Arampatzis, M.A. Katsoulakis, P. Plecháč, M. Taufer, L. Xu, Hierarchical fractional-step approximations and parallel kinetic Monte Carlo algorithms, Journal of Computational Physics 231 (2012) 7795–7814.

[32] K. Esselink, L. Loyens, B. Smit, Parallel Monte Carlo simulations, Physical Review E 51 (1995) 1560–1568.

[33] L. Loyens, B. Smit, K. Esselink, Parallel Gibbs-ensemble simulations, Molecular Physics 86 (1995) 171–183.

[34] E. Bernard, W. Krauth, Two-step melting in two dimensions: first-order liquid-hexatic transition, Physical Review Letters 107 (2011) 1–4.

[35] R. Durstenfeld, Algorithm 235: random permutation, Communications of the ACM 7 (1964) 420.

[36] W. Krauth, personal communications, 2012.

[37] N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, E. Teller, Equation of state calculations by fast computing machines, Journal of Chemical Physics 21 (1953) 1087.

[38] V.I. Manousiouthakis, M.W. Deem, Strict detailed balance is unnecessary in Monte Carlo simulation, Journal of Chemical Physics 110 (1999) 2753.

[39] NVIDIA, CUDA C programming guide, v4.2, 2012.

[40] D.B. Kirk, W.-m.W. Hwu, Programming Massively Parallel Processors: A Hands-on Approach, Morgan Kaufmann, 2010.

[41] R. Farber, CUDA Application Design and Development, Morgan Kaufmann, 2011.

[42] J. Sanders, CUDA by Example, Addison–Wesley Professional, 2010.

[43] S. Worley, Saru, personal communication, 2008.

[44] C.L. Phillips, J.A. Anderson, S.C. Glotzer, Pseudo-random number generation for Brownian dynamics and dissipative particle dynamics simulations on GPU devices, Journal of Computational Physics 230 (2011) 7191–7201.

[45] M. Engel, J.A. Anderson, S.C. Glotzer, M. Isobe, E.P. Bernard, W. Krauth, Hard-disk equation of state: First-order liquid-hexatic transition in two dimensions with three simulation methods, Physical Review E 87 (2013) 042134.

[46] B. Alder, T. Wainwright, Phase transition in elastic disks, Physical Review 127 (1962) 359–361.

[47] J. Lee, K. Strandburg, First-order melting transition of the hard-disk system, Physical Review B 46 (1992) 11190–11193.

[48] J. Zollweg, G. Chester, Melting in two dimensions, Physical Review B 46 (1992) 11186–11189.

[49] H. Weber, D. Marx, K. Binder, Melting transition in two dimensions: A finite-size scaling analysis of bond-orientational order in hard disks, Physical Review B 51 (1995) 14636–14651.

[50] A. Jaster, Computer simulations of the two-dimensional melting transition using hard disks, Physical Review E 59 (1999) 2594–2602.

[51] C. Mak, Large-scale simulations of the two-dimensional melting of hard disks, Physical Review E 73 (2006) 1–4.