

The development and expansion of HOOMD-blue through six years of GPU proliferation

Joshua A. Anderson^a, Sharon C. Glotzer^{a,b,*}

^aDepartment of Chemical Engineering, University of Michigan, Ann Arbor, Michigan, 48109, USA

^bDepartment of Material Science and Engineering, University of Michigan, Ann Arbor, Michigan, 48109, USA

Abstract

HOOMD-blue is the first general purpose MD code built from the ground up for GPU acceleration, and has been actively developed since March 2007. It supports a variety of force fields and integrators targeted at soft-matter simulations. As an open source project, numerous developers have contributed useful feature additions back to the main code. High performance is the first design priority in HOOMD-blue development. Over the years, we have rewritten core computation kernels several times to obtain the best possible performance on each GPU hardware generation. Ease of use is the second priority. Python job scripts allow users to easily configure and control their simulations. Good object-oriented software engineering practices ensure that HOOMD-blue's code is easy to extend, both for the users and developers. Extensive validation tests show that HOOMD-blue produces correct results.

Keywords: GPU, GPGPU, CUDA, Molecular Dynamics, MD, Kepler, Fermi, Tesla, GK110, GF100, G200, G80, NVIDIA

1. Overview

HOOMD-blue is a Molecular Dynamics (MD) simulation package written for graphics processing units (GPUs). It is a flexible code capable of executing many types of simulations with a variety of force fields. Every step of the computation is performed on the GPU and tuned for maximum performance. HOOMD-blue executes typical benchmarks on a single GPU at the same performance as 64–128 CPU cores on a cluster. Many researchers take advantage of this to replace their clusters with inexpensive workstations, while others are using clusters of GPUs to explore larger regions of phase space or compute better statistics than was previously possible.

HOOMD-blue is open-source software. Anyone may download the code [1] and use it under a permissive license. Recently, it has been used to study polymer self-assembly [2–10], polymer scaling laws [11–14], nanoparticle self-assembly [15–19], nanoparticle formation [20–22], patterns formed by ligands grafted to nanoparticles [23–25], complex fluids [26–32], DNA-directed self-assembly [33–36], dynamics of non-equilibrium systems [37–40], to determine thermodynamic properties of nanoparticles [41, 42], as a testbed for new temperature sampling methods [43], to study the properties of droplets on solid surfaces [44], and protein folding [45].

HOOMD-blue began development in March, 2007 shortly after the release of NVIDIA CUDA version 0.8 beta. After one summer's work, an initial working implementation was benchmarked and the algorithms published [46]. The initial code release, version 0.6.0, is little more than a benchmark code, with all routines in C++ and only basic support for the Lennard-Jones pair potential, harmonic bonds, and Nosé-Hoover thermostat. Since then, it has grown into a large codebase with

many features and the associated documentation. [Figure 1](#) shows how the number of lines of code and documentation have increased over time.

Today, HOOMD-blue version 0.11.3 is a capable general purpose MD tool geared toward soft-matter and coarse-grained simulations. It supports

- Lennard-Jones, Gaussian, CGCMM, Morse, Yukawa, EAM, and interpolated table pair forces
- Long range electrostatics via PPPM
- Harmonic, FENE, and interpolated bond forces
- Harmonic and CGCMM angle, dihedral, and improper forces
- NVE, Nosé-Hoover and Berendsen NVT, NPT, NPH, Langevin dynamics, DPD, and energy minimization for particles
- NVE, NVT, NPT, NPH, and energy minimization for rigid bodies
- 2D and 3D simulations
- Constraint that fixes a subset of particles onto the surface of a sphere
- Built-in initialization of random polymer systems
- IMD communication with VMD for visualization of simulations in real-time
- XML, DCD, MOL2, and PDB output file formats

A thriving community makes use of and develops for HOOMD-blue. From April 4 2011 through Jun 6 2013 a total

*Corresponding author

Email address: sglotzer@umich.edu (Sharon C. Glotzer)

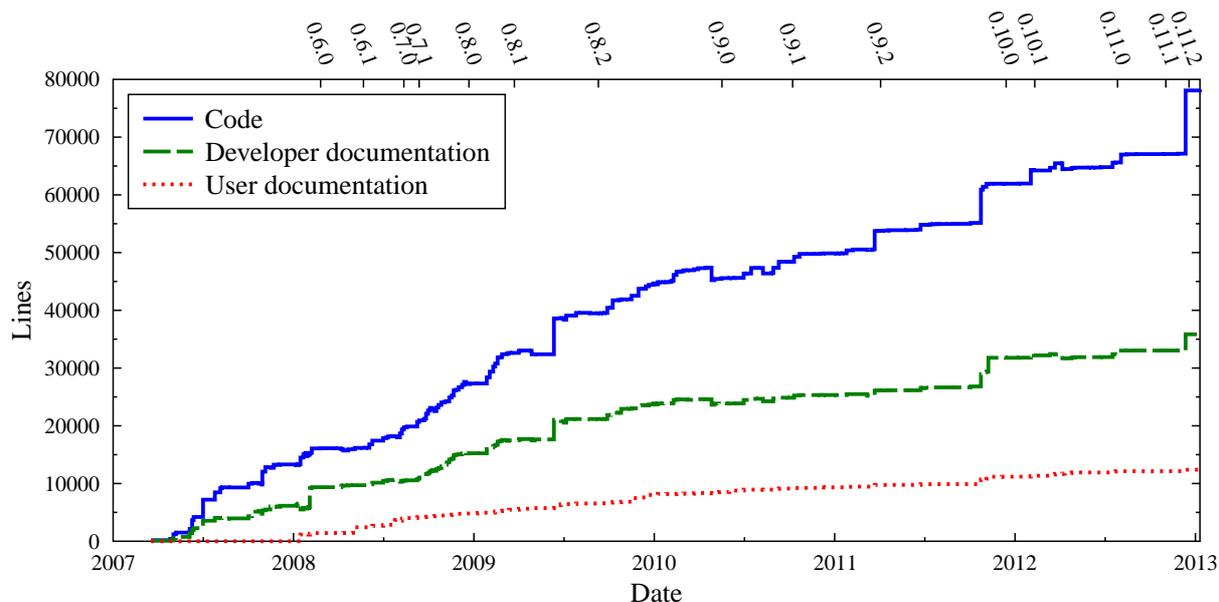


Figure 1: Lines of code in the master branch of HOOMD-blue as a function of date. *Code* counts the total lines of C++ and python code, not counting comments and blank lines, *Developer documentation* counts the lines of comments in the C++ code, and *User documentation* counts the number of comment lines processed into the user manual. Lines wrap at 120 characters and the spikes on the graph are the result of merging large feature branches or patches into the master branch.

of 2568 unique IPs have downloaded the software (see Figure 2 for a complete history); that is an average of 3 new downloads per day. Of those, 516 came back and downloaded updated versions. As of this writing, 45 peer-reviewed articles have been published that utilized HOOMD-blue to perform simulations and there are 198 members on the mailing list and 408 threads of conversation since its creation.

In the next release, version 1.0, we plan to add optimizations for the latest hardware generation, MPI scaling to multiple GPUs in a cluster, improved documentation, fast access to particle data for analysis tools, an improved binary dump file format, full double precision computations, triclinic unit cells, NPT integration that varies box angles, and interpolated table angle and dihedral potentials.

The predecessor of HOOMD-blue went by the name HOOMD and was developed at Iowa State University by Joshua A. Anderson and Alex Travesset. While at Iowa State, three major versions of the code were released: version 0.6.0 in February 2008, 0.7.0 in August 2008, and 0.8.0 in December 2010. Version 0.6.0 is the initial release of the benchmark code used in preparing ref. [46]. Version 0.7.0 is the first that implements Python job scripts. Version 0.8.0 is the first to include specific optimizations for G200 GPUs.

Development moved to the University of Michigan in August 2009 and the code changed its name to HOOMD-blue. At Michigan, Joshua A. Anderson continues as the lead developer for the project, now working with Sharon C. Glotzer and her research group. Six other groups have contributed to the project. Three major releases of HOOMD-blue have been made out of Michigan: Version 0.9.0 in May 2010, 0.10.0 in December 2011, and 0.11.0 in July 2012. Version 0.9.0 is the first to include optimizations for the GF100 GPUs. Version

0.10.0 adds the capability of simulating rigid bodies [47]. Version 0.11.0 adds many usability enhancements, NPH integration, and computation of the pressure tensor. Nine other minor releases along the way offer new features, bug fixes, and performance enhancements.

Over the course of the six years of HOOMD-blue development, NVIDIA has released four generations of hardware. Each introduces new features that are important to HOOMD-blue. G80 is the first processor capable of performing general purpose computations (GPGPU) with NVIDIA CUDA. G200 expands on G80 with more registers, atomic operations, and relaxed coalescing rules. GF100 is a radical redesign, the first that fully caches all memory loads in a L1/L2 cache hierarchy. GK110 expands on GF100 with increased parallelism, more registers, and a host of other capabilities.

Each new generation expands on the maximum throughput capabilities of its predecessors. G80 is capable of 350 GFLOPs, G200 boosts that to 622, GF100 pushes 1500, and GK110 attains an impressive 3950 GFLOPs [48]. Smaller steps have been made in between the performance doubling major generations, such as G92 and GF104. These chips are cut down versions of the flagship for the lower cost market segment.

HOOMD-blue is an easy to use (section 2), flexible, and fast simulation tool. This paper discusses all of the work that has been put into HOOMD-blue. Much of the work has implemented capabilities needed to enable science (section 3). Just as much of a focus is put on maintaining and increasing performance (section 4). When using a fast simulation engine, many researchers find that the bottleneck to project completion is no longer the computation itself, but rather the time it takes to configure the runs and analyze the output. As such, the usability of HOOMD-blue (section 5) also receives a great deal of attention.

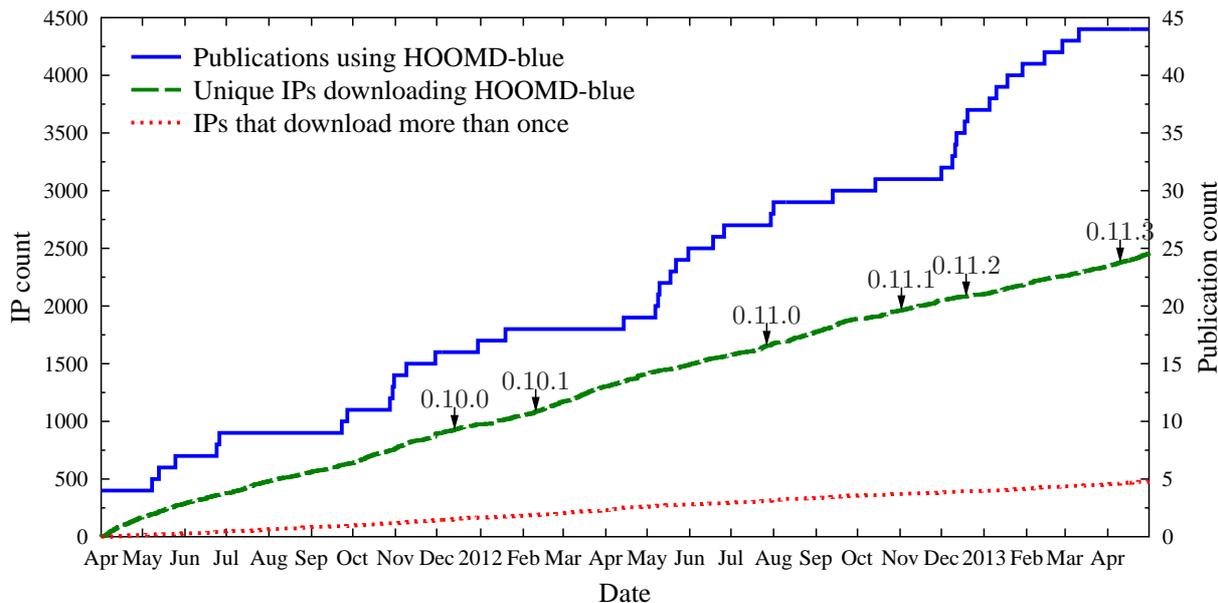


Figure 2: Count of publications using HOOMD-blue and downloads since April 2011. Downloads are tracked via web server logs and the counts include all downloads of HOOMD-blue source code releases, binary packages, and git repository accesses. Accesses by web crawlers and bots are excluded. The green dashed line shows the number of unique IPs that have downloaded HOOMD-blue. The red dotted line counts the number of IPs that have downloaded it two or more times. The solid blue line indicates the number of peer-reviewed publications that have used HOOMD-blue for science.

Over time, newer versions of software and hardware break compatibility. A certain level of maintenance (section 6) must be done to keep everything functioning correctly. Efforts are made so that HOOMD-blue’s core design is extendable (section 7) so that users may add their own functionality in plugins, and so that contributing features is easy. Last, but of no less importance than the first, extensive validation tests (section 8) are performed to ensure that every single line of code operates correctly.

2. HOOMD-blue features and usage

HOOMD-blue is a script driven simulation engine. Users write Python job scripts that configure which capabilities are enabled, set parameters, and control the progression of the run. Algorithm 1 is a simple job script that executes the benchmark used in section 4. Users can enable and configure all HOOMD features through job script commands.

Scripts are run from the command line, or via a job queue submission, through the HOOMD-blue executable `$ hoomd script.py`. When the job starts, it will scan the computer, select the fastest available GPU and execute the simulation. If no GPU is available, the job will run on the CPU. Command line options and/or script commands can be used to force a particular mode.

Job script commands are designed to be simple and concise so that users unfamiliar with Python object oriented code can learn to use HOOMD-blue quickly. Power users can take advantage of Python and write advanced job scripts that provide custom initialization, complex simulation protocols, or analyze data during the simulation.

Algorithm 1 Lennard-Jones liquid job script

```

from hoomd_script import *

# read initial configuration
init.read_xml(filename='input.xml')

# set force field
lj = pair.lj(r_cut=3.0)
lj.pair_coeff.set('A', 'A',
                  epsilon=1.0,
                  sigma=1.0)
nlist.set_params(check_period=5)

# set NVT integration
all = group.all()
integrate.mode_standard(dt=0.005)
integrate.nvt(group=all, T=1.2, tau=0.5)

run(500) # warm up
run(5000) # time this run

```

So, who should use HOOMD-blue? The first consideration is system size. Peak GPU performance is attained around 5000 to 10000 particles for short ranged forces. Below that, users should benchmark and see if CPU or GPU runs are more appropriate. For example, a hypothetical simulation of 10000 particles might have a speedup of 64 that decreases to 32 at 5000, and down again to 10 at 1000. That same simulation with only 100 particles might perform faster on the CPU than the GPU. On the other end of the spectrum, performance scaling is linear in N , so relative performance saturates from 10000 particles on up. However, GPU memory is limited. Typical simulations with ~ 100 neighbors per particle can only fit ~ 3000000 particles on a single 3 GB GPU. Users that want to run simulations with many millions of particles can use a different MD code that scales, or wait for HOOMD-blue version 1.0 which will include MPI domain decomposition [62].

The second consideration is simulation type. HOOMD-blue is a general purpose code, though it primarily targets soft-matter and coarse-grained simulations. Users that want to run all-atom models should use one of the many other MD codes specialized for that purpose. We will accept patches that add all-atom force fields to HOOMD-blue, but we currently have no plans to implement any ourselves.

If HOOMD-blue doesn't have the exact potential, file output type, or analysis needed, there is a plugin mechanism by which users can write their own modules and activate them in job scripts, see [section 7](#) for more information. In the interests of improving HOOMD-blue as a whole, we ask that those who write modules of general interest to the community please either make their plugins available open source or submit them to the developers as patches.

HOOMD-blue is not the only GPU MD code available. Readers may want to check out the following codes which may better fit their research. AMBER [49, 50], NAMD [51], FENZI [52, 53], and GROMACS [54] are specialty all-atom MD codes with GPU capabilities. OpenMM [55] is a GPU-accelerated library that can be called by user code to build MD simulations. LAMMPS [56] and HOOMD-blue share many similarities, as we modeled our design on the spirit of LAMMPS. LAMMPS is more established and has a much larger feature set that includes GPU acceleration. However, HOOMD-blue is significantly faster as our entire design is built around GPU performance from the start. HALMD [57] implements a variation on HOOMD-blue's algorithms with double-single arithmetic. The authors use the code for studying glasses, and the code supports a few simple LJ-type pair potentials, NVE and NVT ensembles.

3. Science enabling features

The very first release of HOOMD was little more than a benchmark code. It was fully functional, for simulations of bead-spring polymers, but lacked many capabilities that are expected of a production quality MD code. Namely, the capabilities present in version 0.6.0 include the core particle data structures, cell list, neighbor list, Lennard-Jones pair force, and harmonic bond force computations, a Velocity Verlet NVE, and

a Nosé-Hoover NVT integrator. It assumed all particles had a mass of 1 and computed only forces, not energies.

3.1. Building on a benchmark code

The first features added to HOOMD (while still developed at Iowa State) resolved these deficiencies and enabled different types of science to be performed with the code.

Energies do not need to be computed for correct dynamics, but production use often requires computing them for analysis [8, 15]. Similarly, NVE and NVT do not need the virial pressure, but is often useful in analysis. HOOMD version 0.7.0 computes these quantities for every particle interaction. They do not decrease performance because the force computations are bandwidth bound and there are "free" FLOPs available on the compute units. To go along with this change, a logging module tracks the values of thermodynamic quantities over time and saves them to a file.

The initial benchmark code used a fixed size cell and neighbor list data structure, and simulations would quit with an error if the size were exceeded. HOOMD version 0.7.0 adapts to any parameters the user chooses by detecting overflows and resizing the data structures as needed.

Many studies use Langevin or Brownian dynamics instead of Nosé-Hoover [4, 8, 17, 24, 26, 38, 39]. A new integrator class enables this as an option in HOOMD version 0.8.0. It is no more computationally demanding than Nosé-Hoover, but the needed parallel random number generation requires care to implement correctly and with high performance [58].

Researchers often use mean squared displacement to determine diffusion coefficients and identify crystallization or vitrification events [33–35]. Particle image flags must be tracked in order to compute the mean squared displacement properly. Every integrator in HOOMD version 0.8.0 is modified to update the image flags whenever a particle crosses the box boundary. It also includes a simple analysis module that computes the mean squared displacement and logs it to a file.

Simulations of polymer systems often use FENE bonds and set the pair cutoff radius for each type pair (instead of globally) [4, 11, 12, 14, 15, 22]. HOOMD version 0.8.0 enables these functionalities with an additional bond force class and modifications to the pair force code.

In many studies, not all particles are the same size [15, 41]. HOOMD version 0.8.1 allows a mass and diameter to be specified for each particle. It also improves energy conservation with shifting and smoothing options for every pair force.

3.2. Additional features

As a general purpose code, most features in HOOMD-blue are optional and only enabled when needed. We include in the main code capabilities that are useful for a range of simulations.

HOOMD-blue includes a Gaussian pair potential for general use [22], a shifted Lennard-Jones potential for use in systems with varying particle sizes [15, 41], and a table based pair potential for use when there is no functional form. Additional potentials are easily added to the code, and users can implement specialized potentials in plugins (see [section 7](#)).

Various types of constraints are also available. Simulations can optionally run in a 2D plane [39], or for a group of particles constrained to the surface of a sphere [24, 25]. Particles may also be constrained together into rigid bodies [8, 16, 26, 28, 33–35]. See our publication ref. [47] for implementation details.

Integration methods are applied to selected groups of particles, and any number may be active at a time. This enables mixed ensembles, fixed particles [23, 37], moving boundary layers [4, 39] and more.

FIRE [59] allows energy minimization runs to be performed with only minimal amount of code built on top of NVE integration, and it works for rigid bodies [47] and free particles.

Some simulations [7, 8], have a large number of exclusions, but the original method of processing them was limited to four due to performance concerns. We remove the limit with a generic solution that processes exclusions in batches, so it gracefully degrades in performance as more are added.

HOOMD-blue also performs Dissipative Particle Dynamics, commonly used for polymer simulations [13, 21, 22, 25], and sometimes used in other types of simulations [37]. See our publication ref. [58] for implementation details.

3.3. Contributions

HOOMD-blue is open source code and contains many contributions from outside of Michigan. The following people have submitted their modifications back to the developers with permission to release their work in HOOMD-blue. The Angle, Dihedral, and Improper and CGCMM coarse-grained potentials are contributed by Axel Kohlmeyer, David LeBard and Ben Levine from the ICMS group at Temple University [3]. EAM functionality is contributed by Igor Morozov, Andrey Kazenkov, Roman Bystryi, from the Joint Institute for High Temperatures of RAS [40]. PPPM electrostatics is contributed by Stephen Barr from Princeton University and Philipp Mertmann from Ruhr University Bochum [7]. The Yukawa and Morse potentials are contributed by Rastko Sknepnek from Northwestern, as is the original NPT integrator. The Berendsen thermostat is contributed by Brandon Denis Smith from the Non-equilibrium Gas and Plasma Dynamics Laboratory at the University of Michigan. Smith also contributed a patch that enables full double precision computations in HOOMD-blue GPU code. Jens Glaser from the Morse group at the University of Minnesota contributed NPH integration, modified the code to compute the full pressure tensor, enabled the dynamic insertion and removal of bonds, a new NPT integrator, and is currently working on MPI domain decomposition and triclinic simulation boxes. Jens Glaser and Pavani Medapuram from the Morse group contributed the external potential framework.

4. Performance improvements

Code does not exist in a vacuum. Performance of a code does not remain fixed at levels published years ago. Code exists in an ecosystem of hardware, compilers, and developers. All three contribute to performance improvements over time. To examine these factors, we run a single benchmark for all versions of

HOOMD version	G92 (9800 GTX)	G200 (S1070)	GF100 (GTX 580)	GK110 (GTX Titan)	NVCC version
0.7.0	167	212	461		2.2
0.7.1	167	212	465		2.2
0.8.0	190	243	399		2.2
0.8.1	196	278	417		2.2
0.8.2	198	278	420		2.2
0.9.0	205	295	544		3.0
0.9.1	206	290	725		3.0
0.9.2	195	302	787		3.2
0.10.0	194	301	757		4.0
0.10.1	194	300	757		4.0
0.11.3				1270	5.0

Table 1: Absolute performance of the LJ liquid benchmark (in time steps per second) as a function of HOOMD-blue version and on three major GPU generations. Bold numbers indicate the first version of HOOMD-blue to officially support each GPU architecture. Note that even the oldest version of HOOMD compiled with an old version of CUDA still runs on recent hardware, though upgrading to the latest version of HOOMD-blue offers significant performance improvements. GK110 results were only recorded on the latest version.

HOOMD-blue on three generations of GPUs: G92, G200, and GF100.

The benchmark is a $N = 64000$ particle Lennard-Jones fluid at packing fraction $\phi_P = 0.20$, with $r_{\text{cut}} = 3.0$ at a reduced temperature $k_b T = 1.2$ (Algorithm 1). This is the same benchmark used in ref. [46] to show a directly comparable improvement over time. Some may criticize the choice of $r_{\text{cut}} = 3.0$ and prefer the more popular value of 2.5. HOOMD-blue is certainly capable of such a simulation and will run faster as a result, but the trends discussed here do not change significantly. Readers wishing to reproduce these performance results for comparison are welcome to. Simply download the latest version of HOOMD-blue and run the accompanying benchmark script. By the time this article is published, the included performance numbers will already be out of date. We post additional benchmarks to our website [1], you may find a benchmark more relevant to your research there.

Table 1 lists the results of our benchmarking. First, look at how far we have come. Today, version 0.11.3 executes this benchmark at 1270 steps per second on a GK110 (GTX Titan). That is 6.35 times faster than the 200 steps per second obtained on the 8800 GTX (G80) GPU in ref. [46], or 7.6 time faster than the 9800 GTX (G92). A G80 GPU is no longer available for testing, so we make all further comparisons against G92, a slightly smaller version of the same generation.

4.1. Faster hardware

The biggest performance increases come from the GPU hardware itself. Each new generation is capable of processing more work in parallel and adds new capabilities that enable new and faster algorithms. Fixing the code version to 0.10.1 and comparing just GPUs in Table 1, G200 is 1.54 times faster than G92 and GF100 is 2.52 times faster than G200. Moore’s law is alive

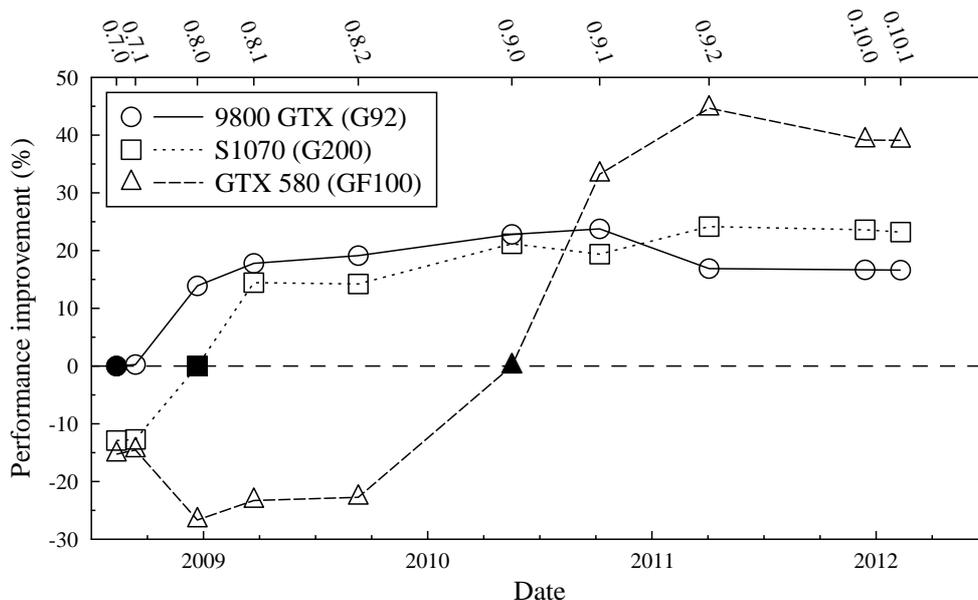


Figure 3: Relative performance of the LJ liquid benchmark as a function of date. Filled symbols indicate the first version of HOOMD-blue to officially support each GPU architecture, and are also the 1.0 relative baseline. Points after this release show improved relative performance due to code optimizations and compiler improvements for each architecture. Points before the baseline release show the relative performance degradation from running unoptimized and untuned code on new GPUs. The lines serve to guide the eye. [Table 1](#) lists the absolute values for the data points in this plot.

and well in GPUs and still equates to a doubling in performance each generation.

Some of these performance boosts come “for free”. For example, version 0.7.1 was released before G200 cards existed, yet it runs 27% faster on G200 than on G92. Similarly, Version 0.8.2 came out prior to GF100 hardware, and runs 51% faster on it when compared to G200.

[Figure 3](#) is an alternate view of the data in [Table 1](#). It chooses the first HOOMD-blue release that supports a given hardware generation as a baseline and plots relative performance improvements from that. The figure shows how performance on fixed hardware has varied over time due to other factors.

As discussed in ref. [46], tuning the kernel block size and r_{buff} values do not change the correctness of the simulation, but do have a large effect on the performance. In each case, the first release to specifically support a new hardware generation has these parameters tuned specifically for that hardware. For example, the 30% performance boost on GF100 from version 0.8.2 to 0.9.0 is due entirely to parameter tuning.

4.2. Better compilers

New compilers are also a source of performance improvement. There have been fifteen public releases of CUDA to date. Each new release supplies support for new hardware features and comes with improvements to the optimizer. These improvements are small but measurable, and add up over time. For example, recompiling and retuning version 0.10.1 with CUDA 4.1 nets a 3.1% rise in performance. Unfortunately, not every change boosts performance. The 4% performance loss from version 0.9.2 to 0.10.0 (on GF100) results from the CUDA version upgrade.

4.3. Improving code

Compiler optimizers can make only small tweaks to the order of machine instructions, register allocation, and so forth. They are far more limited than the programmer, who can reorganize data structures, make changes to thread work assignment, or implement whole new algorithms. Improvements to the code can have a profound effect on performance. Every new release of HOOMD-blue contains numerous changes that boost performance in small ways. Much of the slow upward track of G92 and G200 performance comes from code changes, totaling 23% up from the baseline.

The biggest of these changes is a complete rewrite of the neighbor list generation code. In the original version [46], one block is executed per cell and shared memory is used to stage memory reads from other cells. There is a big problem with this for a general purpose code. User chosen parameters can result in anywhere from two up to hundreds of particles in each cell. In the small cases, there are many threads left doing nothing. In the large cases, the cells do not even fit in shared memory. The solution is to break up the computation further to run one thread per particle and rely on the texture cache for the neighboring cell reads. This change alone provides the 14% performance increase from 0.7.1 to 0.8.0 on G92. Rapaport independently arrived at the same improved algorithm [60].

G200 is only a small architectural tweak from G92. Most of the changes tested on one also benefitted the other, as can be seen by the parallel tracks in [Figure 3](#). In contrast, the cache hierarchy on GF100 makes a big leap forward. A second complete rewrite of the cell and neighbor list code, with this cache in mind, grants the 33% performance increase from version 0.9.0 to 0.9.1.

4.3.1. Neighbor list optimizations

Algorithm 2 Generate cell list

Require: gdim is $\lceil \frac{N}{\text{bdim}} \rceil$
Require: Particles are in a box extending from \vec{b} to $\vec{b} + \vec{L}$
Require: $(s_x, s_y, s_z) = (m_x, m_y, m_z) / \vec{L}$
Require: celln is cleared to zero prior to the kernel launch

- 1: $i \leftarrow \text{bdim} \cdot \text{bid} + \text{tid}$
- 2: **if** $i \geq N$ **then**
- 3: **return**
- 4: **end if**
- 5: $\vec{R} \leftarrow \text{pos}[i]$
- 6: $\vec{r} \leftarrow \vec{r} + \vec{b}$
- 7: $c_x \leftarrow \lfloor r_x \cdot s_x \rfloor$
- 8: $c_y \leftarrow \lfloor r_y \cdot s_y \rfloor$
- 9: $c_z \leftarrow \lfloor r_z \cdot s_z \rfloor$
- 10: $n \leftarrow \text{atomicInc}(\text{celln}[c_x, c_y, c_z])$
- 11: $(\vec{R}, i) \Rightarrow \text{celldata}[n, c_x, c_y, c_z]$

Two aspects of GF100’s cache enable this large performance gain. First, atomic operations are performed in on-chip L2 cache and thus execute fast. This enables a simple cell list kernel, Algorithm 2. It runs one thread per particle (lines 1–4). The assigned particle is read from global memory (line 5), and its cell index computed (lines 6–9). An atomic increment adds 1 to the count of particles in that cell, celln . The atomic operation reserves a space in the cell list and ensures that all threads inserting into the same cell are resolved correctly, though the order of entries is undefined. Finally, both the position and index are written into the cell data structure celldata at the reserved location.

Algorithm 2 on GF100 is about 10 times faster than computing the cell list on the host. It provides further performance gains because it removes the need for host/device memory copies, and mitigates the performance lost to the serial bottleneck, as it accelerates one of the two steps that ref. [46] left on the CPU. HOOMD-blue can now perform hundreds of steps fully on the GPU with no host interaction. The last step left on the CPU is the SFCPACK sort [46]. It requires very little processing, but is becoming a serial bottleneck as GPUs get faster. We have not taken the time to implement it on the GPU yet, but Colberg has [57].

Algorithm 3 takes the output celldata and uses it to compute the neighbor list. It also executes only one thread per particle (lines 1–4) and starts by computing the particle’s cell index (lines 5–9). Then it loops over all 27 adjacent cells (line 11). Cell adjacency may be simple to compute on the fly, but doing so requires storing and looping over all three cell indices. The increased register pressure reduces performance significantly. Instead, we read precomputed adjacent cell indices on the GPU (line 12). Next, the kernel loops over all particles in the adjacent cell (line 14), reads in the position and index (line 15), and tests if it passes the cutoff (line 16–17). Neighboring particles are appended to the neighbor list data structure nlist (lines 18–19). Finally, the total number of neighbors in the list is written to nlistn (line 22). In this simplification of the actual

Algorithm 3 Generate neighbor list

Require: gdim is $\lceil \frac{N}{\text{bdim}} \rceil$
Require: Particles are in a box extending from \vec{b} to $\vec{b} + \vec{L}$
Require: $(s_x, s_y, s_z) = (m_x, m_y, m_z) / \vec{L}$

- 1: $i \leftarrow \text{bdim} \cdot \text{bid} + \text{tid}$
- 2: **if** $i \geq N$ **then**
- 3: **return**
- 4: **end if**
- 5: $\vec{R}_i \leftarrow \text{pos}[i]$
- 6: $\vec{R}_{\text{tmp}} \leftarrow \vec{R}_i + \vec{b}$
- 7: $c_x \leftarrow \lfloor R_{\text{tmp},x} \cdot s_x \rfloor$
- 8: $c_y \leftarrow \lfloor R_{\text{tmp},y} \cdot s_y \rfloor$
- 9: $c_z \leftarrow \lfloor R_{\text{tmp},z} \cdot s_z \rfloor$
- 10: $n_{\text{neigh}} \leftarrow 0$
- 11: **for** $a \in [0 \dots 27]$ **do**
- 12: $(d_x, d_y, d_z) \leftarrow \text{celladj}[a, c_x, c_y, c_z]$
- 13: $n_d \leftarrow \text{celln}[d_x, d_y, d_z]$
- 14: **for** $o \in [0 \dots n_d]$ **do**
- 15: $(\vec{R}_j, j) \leftarrow \text{celldata}[n, d_x, d_y, d_z]$
- 16: $\vec{r} \leftarrow \text{minimage}(\vec{R}_i - \vec{R}_j)$
- 17: **if** $|\vec{r}| < r_{\text{max}}$ **and** $i \neq j$ **then**
- 18: $j \Rightarrow \text{nlist}[i, n_{\text{neigh}}]$
- 19: $n_{\text{neigh}} \leftarrow n_{\text{neigh}} + 1$
- 20: **end if**
- 21: **end for**
- 22: $n_{\text{neigh}} \Rightarrow \text{nlistn}[i]$
- 23: **end for**

code, we do not show the overflow detection or other specific features such as exclusions.

Data structures and their layout in memory are critically important to the performance of the code. Cell adjacency is stored in a 4D matrix, $\text{celladj}[a, c_x, c_y, c_z]$, where a is the fastest index. The cell data itself is stored in the same layout. Due to the SFCPACK sort [46], threads within the same block are likely to be assigned to particles within the same cell, or nearby cells. So most of the threads in a block are likely accessing the same indices in celladj and celldata , utilizing the L1 cache effectively. Furthermore, threads are likely to find data already in cache as they loop over the fast index. Combined, these result in a 96% cache hit ratio. We performed benchmarks with different memory layouts and determined this is the fastest performing one.

The neighbor list nlist is stored as a 2D sparse matrix with the particle index as the fast index, the same as in ref. [46]. Threads in a warp will find neighbors at different points in the loop so their n_{neigh} counters will not be synchronized. The write on line 19 is thus a scattered write and a big bottleneck in Algorithm 3. Unfortunately, GF100 does not have enough shared memory to regularize the write.

5. Usability enhancements

We put as much effort into designing usable interfaces as we do to performance. When simulations run fast, the speed at

which the researcher can interpret the output and prepare new simulations becomes a limiting factor. Every new feature incorporated into HOOMD-blue is carefully planned and all potential use-cases are considered so that the interface exposed to users is as simple and elegant as possible.

5.1. Python job scripts

HOOMD-blue is a script driven simulation engine. Users can combine any number of its features together and control them with a single script file. Scripts can also read and write all particle, rigid body, and bond connectivity properties.

Job scripts are Python scripts. Python, a fully featured object oriented programming language, enables users to perform advanced tasks from within their job. For example, the initial configuration of the system can be set using Python code. Complex analysis can also be performed while the simulation is running.

In the design phase we first chose to have this flexibility, then we examined various options that would enable that. A custom script language was not considered because of the development effort needed to implement it. Python is our tool of choice because of its widespread adoption, beautiful syntax, and excellent integration with C++.

All of the high performance code in HOOMD-blue is written in C++ and CUDA, and all of the user interface elements in Python. This clear delineation proves very useful and enables a rapid development cycle. Core C++ data structures and compute kernels are written with only performance in mind and do not need a nice or stable interface. Python control code can be prototyped quickly and insulate users from any underlying changes. Over HOOMD-blue's history, the internal C++ interface has changed drastically numerous times, but users have been able to run existing job scripts relatively unmodified.

5.2. Documentation

Documentation is important both to users and developers. Users need to know what the software is capable of and how to instruct it to perform their simulations. Developers need to understand the overall design of the code, how data structures are laid out, and the details of every computation. Both types of documentation exist as comments in the code. Doxygen [61] generates release quality documents from the comments that are cross-referenced and easily navigable. We use it to create PDF and web page documentation, which are posted online [1].

Figure 1 shows how the amount of documentation has increased over time. In the current code base, there are approximately 80000 lines of code, 36000 lines of developer documentation and 12000 lines of user documentation.

HOOMD-blue's user documentation is extensive. It contains everything a user needs to know to get started, including a list of required hardware and software, installation instructions, and a tutorial introduction to job scripts. A large set of example scripts show how various features operate. The examples are cross-linked with the detailed descriptions for each script command. All commands are listed in an index with links to each command's documentation page where the commands syntax, parameters, and all options are described.

5.3. Operating system support

HOOMD-blue is supported on most flavors of Linux and Mac OS X. We provide binary packages that users can download and easily install on their machine. The documentation also includes detailed instructions on compiling HOOMD-blue so that users can build from source if they choose.

5.4. Integrate with external tools

HOOMD-blue integrates well with other tools. Trajectories are written in the DCD file format. DCD is a common format and many analysis and visualization codes, such as VMD, read it. Python job scripts can import any external tools written in Python and directly interface with them for simulation set up and analysis.

VMD is capable of connecting directly to a running simulation over a simple network protocol called Interactive Molecular Dynamics (IMD). HOOMD-blue provides an IMD module so that VMD users can make use of this functionality. VMD also has a file reader plugin, written by Axel Kohlmeyer, that can read and write HOOMD-blue's native XML file format.

6. Maintenance

Not all of the changes that go into HOOMD-blue enable science, enhance performance, or boost usability. Constant maintenance updates are also performed. Maintenance is needed to fix bugs and keep software running correctly as libraries and compilers evolve. We also continually review the code base and modify it to keep it clean and readable so that future developers can understand it. Changes are also needed to support new hardware.

7. Code extensions

HOOMD-blue's simulation engine is written with a modular and object oriented design. Common software engineering practices are employed to enable a flexible design with minimal dependencies between separate modules. This makes it easy for developers to change or add new features without requiring significant changes in the code.

Users can write plugins that implement specific functionality needed for their research. Plugins exist outside of HOOMD-blue and are imported into a job script using python's dynamic module loading capabilities. Users can easily update to a new version of HOOMD-blue without needing to maintain and update a patch.

The most common user task is adding a pair potential specific to their research. HOOMD-blue uses functors and templates to make this as easy as possible. All a user needs to do is write a small functor class that computes $V(r)$ and $-\frac{1}{r} \frac{dV}{dr}$, and the template pair potential class does the rest.

8. Validation tests

Extensive validation tests are performed on HOOMD-blue to ensure that the code is correct when written and remains correct as it is maintained. The importance of the latter is not overlooked as code changes in unrelated modules can sometimes come together to cause incorrect behavior. Validation tests must be repeatable and have a well-defined answer computed by some means external to the code.

We use two types of tests common in software engineering practice to validate HOOMD-blue.

8.1. Unit tests

A white box unit test is written for each module in HOOMD-blue. It is designed with knowledge of how the module is implemented and sufficient cases are tested so as to exercise all of the possible code paths that the module can take.

For example, a unit test for the Lennard-Jones pair force class creates a system with six particles in it. Some of these interact across periodic boundary conditions while others interact directly. It then sets the potential parameters ε , σ , and r_{cut} to different values and verifies that the correct force is computed each time. The correct force is determined in each case by hand with a calculator. As an added test, a large random configuration of particles is generated and the force computed by the CPU and GPU compared. If any of the values is outside of a tolerance from the correct value, the test fails.

Each and every class in HOOMD-blue is tested in this way. In total, there are more than 2000 independent tests in version 0.10.1's unit test suite. These cover more than 100 classes. All these tests run nightly on the latest development version of the code. They are run on five different build configurations each on Mac, various flavors of Linux, and all supported GPU generations. One batch of tests is even run through valgrind to detect any uninitialized memory reads or out of bounds accesses. Test results are uploaded to a website dashboard where they can be monitored. Any test failure is counted as a bug and we work to fix them as soon as they appear.

8.2. System integration tests

Unit tests show that each individual component works independently. To show that the entire system is valid when put together, we also perform black box system integration tests. Black box tests assume no knowledge about how the system works. Simulation parameters are entered into a job script and output files analyzed, the same way a user would run HOOMD-blue.

In this paper, we choose to validate HOOMD-blue using the simple Lennard-Jones liquid. We plan to perform more system integration testing using other features, such as polymer and rigid body systems. Future test results will be published on the HOOMD-blue website [1] and used to validate each new release version.

In a complicated system of interacting code modules, there is no way to prove that a given code is correct. Instead, we compare against a baseline. LAMMPS [62] is an actively developed general purpose MD code. We run the same simulations in HOOMD-blue and LAMMPS and compare the results.

We use the same methodology to also validate that HOOMD-blue running on the GPU produces the same results as on the CPU. Specifically, in versions 0.11.x and earlier, HOOMD-blue only computed values in single precision on the GPU. For some applications, double precision may be warranted citeColberg2011.

8.2.1. Test setup

Our test procedure is as follows. First, N particles are randomly placed in a cubic simulation box at number density ρ . The shifted Lennard-Jones pair potential

$$V(r) = V_{\text{ij}}(r) - V_{\text{ij}}(r_{\text{cut}})$$

is applied between every pair of particles, where

$$V_{\text{ij}}(r) = \begin{cases} 4\varepsilon \left[\left(\frac{\sigma}{r} \right)^{12} - \alpha \left(\frac{\sigma}{r} \right)^6 \right] & r < r_{\text{cut}} \\ 0 & r \geq r_{\text{cut}} \end{cases}$$

A run of 10^6 steps in the NVT ensemble at temperature T thermalizes the system. Then, any net momentum is removed from the system and an input file is written.

Next, four separate simulations are all started from the same input file. One uses the Nosé-Hoover NVT integrator at temperature T , one uses NPT at pressure P_{set} , one is performed in the NVE ensemble and one is performed with Langevin dynamics. Thermodynamic quantities are logged every 500 time steps and trajectory configurations saved every 100000. The parameters have values $N = 5000$, $\rho = 0.4/\sigma^3$, $r_{\text{cut}} = 2.5\sigma$, $k_{\text{B}}T/\varepsilon = 1$, $P_{\text{set}} = 0.06/(\sigma^3/\varepsilon)$ and a step size of $\delta t = 0.005\tau$. The time constants for NVT and NPT are all set to 1.0. They are run for an additional $20 \cdot 10^6$ steps.

Finally, a short NVE run of only 10τ is performed at various values of δt to show energy conservation. These runs log thermodynamic quantities at every step.

This process is repeated eight times each with different random number seeds. Statistics are computed over all eight runs. The mean values are reported with error bars of two standard errors of the mean (two standard deviations divided by $\sqrt{8}$).

In all runs, the basic thermodynamic quantities of kinetic energy K , total energy E , pressure P , and density ρ are computed instantaneously at the time of logging and written to a file. Additionally, the mean squared displacement [63]

$$\langle r(t)^2 \rangle = \frac{1}{N} \sum_{i=1}^N [\vec{r}_i(t) - \vec{r}_i(0)]^2$$

is computed and saved.

An approximate diffusion coefficient D is determined by the coefficient to the best line fit of $\langle r(t)^2 \rangle$ vs. t . Heat capacity is easily computed from fluctuations in the energy

$$\frac{C_V}{k_B} = \frac{\langle E^2 \rangle - \langle E \rangle^2}{(k_B T)^2}.$$

Lastly, the pair distribution function $g(r)$ is determined in the usual manner from particle positions.[63]

Short term energy fluctuations are characterized by the standard deviation over the mean energy in the short runs. Energy and momentum drifts in the long runs are determined from a best line fit

$$X(t) = X(0) + X_{\text{drift}} \cdot t,$$

where X is either the total system energy E or the average system momentum p .

8.2.2. Test results

The comparison of thermodynamic quantities are posted on the HOOMD-blue website[1] and we plan to update it with validation studies on each new release version of HOOMD-blue.

Table 2 shows the conservation tests determined during NVE simulations. Simulations with the shifted Lennard-Jones potential show identical energy fluctuation and drift values within error bars. Momentum conservation is the first quantity where values differ. Single precision runs drift at approximately $10^{-10}m\sigma/\tau^2$, and double precision runs drift 10 orders of magnitude slower. While the drift in single precision is detectable, it is still very small. It would take $2 \cdot 10^{10}$ time steps (231 days of run time at 1000 TPS) for it to reach a value of $0.01m\sigma/\tau$. The energy drift is worse, taking only $2 \cdot 10^7$ steps to accumulate to a value of 0.01ϵ .

However, the energy drift is not a result of the code, hardware, or the choice of precision, it is the same for all builds tested, including double precision. Rather, it is the choice of r_{cut} , δt and the shifted Lennard Jones potential. Cutting off and shifting the potential introduces a discontinuity in the derivatives of $V(r)$. This breaks the assumption made in the numerical methods underlying MD that $V(r)$ is continuous and infinitely differentiable everywhere. Choosing a larger r_{cut} or smaller δt decreases the drift by an order of magnitude or more. We select $r_{\text{cut}} = 2.5$ and $\delta t = 0.005$ because these are values commonly used in the literature for research problems. Energy conservation can also be improved significantly by smoothing $V(r)$ so that its first derivative goes to 0 at $r = r_{\text{cut}}$. Smoothing is often employed in studies of dense liquids. HOOMD-blue offers the XPLOR smoothing function (also used in NAMD) which users can choose to enable. XPLOR modifies the potential to

$$V(r) = V_{\text{LJ}}(r) \cdot S(r)$$

where

$$\begin{aligned} S(r) &= 1 & r < r_{\text{on}} \\ &= \frac{(r_{\text{cut}}^2 - r^2)^2 \cdot (r_{\text{cut}}^2 + 2r^2 - 3r_{\text{on}}^2)}{(r_{\text{cut}}^2 - r_{\text{on}}^2)^3} & r_{\text{on}} \leq r \leq r_{\text{cut}} \\ &= 0 & r > r_{\text{cut}} \end{aligned}$$

The bottom half of Table 2 shows the results of the conservation test with XPLOR smoothing enabled at $r_{\text{on}} = 3/2\sigma$. Energy fluctuations and momentum conservation are almost identical within error bars to runs with the shifted potential. The single precision builds have an energy drift one order of magnitude smaller than with the shifted potential. And the double

precision build gains another order of magnitude improvement down to $10^{-10}\epsilon/\tau$.

Colberg and Höfling perform a similar validation study on their GPU MD code, HALMD [57]. They report momentum and energy drift rates of approximately $10^{-10}m\sigma/\tau^2$ and $10^{-7}\epsilon/\tau$ for a Lennard-Jones system at $\rho = 0.75/\sigma^3$, $k_{\text{B}}T/\epsilon = 1.12$, $r_{\text{cut}} = 2.5\sigma$, $\delta t = 0.001\tau$, and using a different smoothing function. These match the values we obtain for the shifted potential system at a much higher δt , and our XPLOR smoothed runs exhibit two orders of magnitude smaller energy drifts.

Colberg and Höfling find $10^{-7}\epsilon/\tau$ too large for their detailed study of diffusion in glasses and implement a double-single precision mode. Double-single arithmetic uses two single precision floating point values to increase precision, though not to the same level as full double precision. They keep the force computation in single precision, but use double-single in the force accumulation and integration steps. Using it drops their momentum drift rate down two orders of magnitude to $10^{-12}m\sigma/\tau^2$, which is much higher than what we observe in full double precision builds. Using double-single also lowers their energy drift rate to undetectable levels smaller than $10^{-8}\epsilon/\tau$, however they only run a very short simulation of 10^5 steps from which it is hard to draw conclusions about slow drifts. We extend our CPU double precision XPLOR run out to $200 \cdot 10^6$ steps to more accurately evaluate the drift. Figure 4 plots one selected run for each drift, and Table 2 includes the updated values.

9. Conclusions

HOOMD-blue is the first general purpose MD code built from the ground up with GPU acceleration, and has been actively developed since March 2007. It is a powerful research tool that makes the computational power of a CPU cluster available on the desktop, and for a lower price tag. Many researchers are taking advantage of that to run more, bigger, or longer simulations than was previously feasible.

Throughout these last six years, the developers have added numerous features that enable research simulations while at the same time optimizing the code to run faster and take the most advantage of each new GPU hardware generation. The new features are important to those that need them, but degrading everyone's simulation performance for the benefit of a few is not acceptable. Each enhancement to HOOMD-blue is performed in such a way that performance only changes when optional features are enabled. This is evident in that the historical benchmarks show performance regressions in two cases only, one resulting from the CUDA compiler, and one from a conscious choice to no longer optimize for old G92 hardware.

Occasionally a complete rewrite of performance critical kernels is needed so that it maps well to the latest hardware, but this effort is well worth it, resulting in 20–50% performance improvements in typical benchmarks. New GPU hardware is released every year, and the HOOMD-blue developers are committed to reaching the highest possible simulation performance on each.

		$ \Delta E/E $	$E_{\text{drift}}/(N/\varepsilon/\tau)$	$p_{\text{drift}}/(m\sigma/\tau^2)$
shift	CPU-DP	$2.0 (1) \cdot 10^{-5}$	$1.2 (3) \cdot 10^{-7}$	$0.6 (4) \cdot 10^{-19}$
	CPU-SP	$2.0 (1) \cdot 10^{-5}$	$1.2 (3) \cdot 10^{-7}$	$0.8 (2) \cdot 10^{-10}$
	G200-SP	$2.03 (8) \cdot 10^{-5}$	$0.9 (2) \cdot 10^{-7}$	$1.0 (3) \cdot 10^{-10}$
	GF100-DP	$2.03 (8) \cdot 10^{-5}$	$1.3 (1) \cdot 10^{-7}$	$0.8 (2) \cdot 10^{-19}$
	GF100-SP	$2.0 (1) \cdot 10^{-5}$	$0.9 (3) \cdot 10^{-7}$	$1.0 (3) \cdot 10^{-10}$
xplor	CPU-DP	$1.90 (9) \cdot 10^{-5}$	$0.3 (1) \cdot 10^{-9}$	$1.6 (10) \cdot 10^{-20}$
	CPU-SP	$1.93 (7) \cdot 10^{-5}$	$1.1 (2) \cdot 10^{-8}$	$0.7 (1) \cdot 10^{-10}$
	G200-SP	$1.90 (3) \cdot 10^{-5}$	$3.8 (2) \cdot 10^{-8}$	$1.0 (3) \cdot 10^{-10}$
	GF100-DP	$1.94 (6) \cdot 10^{-5}$	$0.8 (3) \cdot 10^{-9}$	$1.1 (4) \cdot 10^{-19}$
	GF100-SP	$1.89 (5) \cdot 10^{-5}$	$2.0 (2) \cdot 10^{-8}$	$0.7 (1) \cdot 10^{-10}$

Table 2: Energy conservation, energy drift, and momentum drift validation test results for the LJ liquid in the NVE ensemble. The simulation parameters are $N = 5000$, $\rho = 0.4/\sigma^3$, $r_{\text{cut}} = 2.5\sigma$, $k_{\text{B}}T/\varepsilon = 1$, and $\delta t = 0.005\tau$. Several short runs of 10τ are made at various δt from an equilibrated configuration to evaluate the energy fluctuations. Drifts are measured as a best line fit over $10 \cdot 10^6$ steps. Results are averaged over 8 independent sets of runs with different initial conditions. The number in parentheses indicates 2 standard errors of the mean in the last displayed digit. SP indicates a single precision execution and DP indicates double precision.

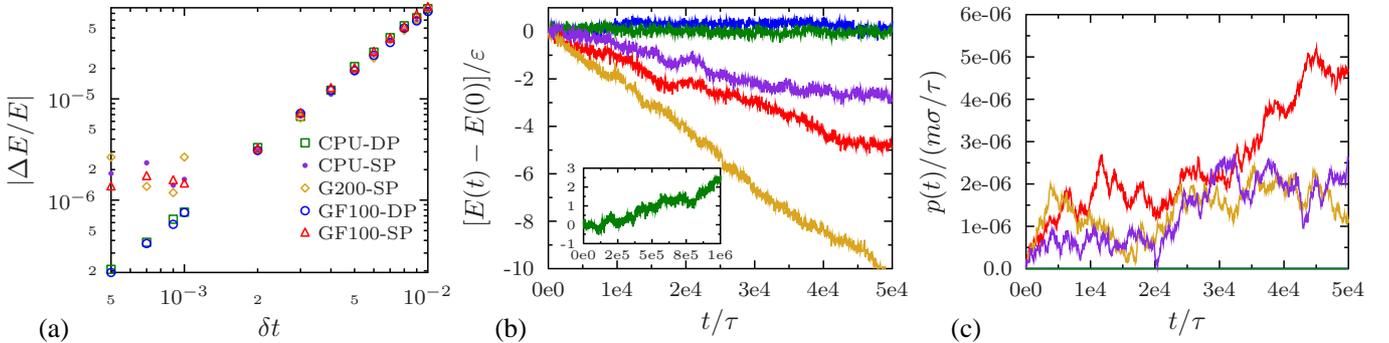


Figure 4: Plots for selected xplor conservation runs used in Table 2. (a) Energy fluctuations in short runs. (b) Energy drift in long runs (c) Momentum drift in long runs.

We value easy to use and extendable code design as much as we do performance. A significant amount of effort is put into the class structure and use cases of every piece of HOOMD-blue so that future development can easily modify and expand on it. Similarly, the command syntax is carefully planned so that users can set any option with a simple, succinct job script. A simple, elegant interface reduces the time it takes a user to configure their simulations. Complexity only leads to time consuming work and more opportunities for errors.

Simulation results are no good if they are not correct. HOOMD-blue developers perform extensive validation testing on every part of the code. Automated tests run each night to alert us when we make a change that creates a bug.

System integration tests demonstrate that HOOMD-blue in any build configuration, and on any hardware, produces the same results to a high degree of accuracy. No detectable difference in thermodynamic quantities is seen between any build configuration of HOOMD-blue, including the default mode of all single precision computations on the GPU. In simulations with typical parameters used in current research projects, the only noticeable change from single to double precision is that a barely noticeable momentum drift disappears. These tests provide a high degree of confidence that HOOMD-blue performs

correct results and is a production quality code.

Despite all the work that has been done, there is always room for improvement. As careful code is tested, there are still rare corner cases that trigger incorrect behavior. Continued code maintenance is needed to fix these bugs and to add support for new hardware, linux flavors, library and compiler versions. The latest GPU generation, Kepler GK110, is now available, and it is a moderate change from GF100. The current version of HOOMD-blue runs on it, of course, but performance tuning and code optimization needs to be done to obtain the fastest performance.

When one GPU is not fast enough, two or more are an option. So far, all HOOMD-blue releases run only on a single GPU. Jens Glaser from the University of Minnesota has submitted a patch that enables domain decomposition simulations in HOOMD-blue using MPI. We see strong scaling down to 20,000 particles per GPU.

We have usability improvements and other new features planned as well. The most extensive is a new binary file format that will enable efficiently saving trajectories containing any properties of particles and rigid bodies. Users will be able to use this for offline analysis of simulations where particle types change over time, or to easily compute order parameters of rigid

body orientations, for example.

Expect to see all of these changes, as well as many more to enable new types of simulations, in future versions of HOOMD-blue.

10. Acknowledgements

We thank all of the HOOMD-blue contributors for their help and support in making the code what it is today. We acknowledge discussions with Carolyn L. Phillips and Karen J. Coulter in outlining this work, and Ryan L. Marson for help editing. NVIDIA provided the GTX Titan used in some benchmarks reported in this work. We acknowledge support by the Assistant Secretary of Defense for Research and Engineering, U.S. Department of Defense [DOD/ASD(R&E)](N00244-09-1-0062). Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the DOD/ASD(R&E).

References

- [1] HOOMD-blue, <http://codeblue.umich.edu/hoomd-blue>, 2012.
- [2] J. A. Anderson, R. Sknepnek, A. Travesset, Design of polymer nanocomposites in solution by polymer functionalization, *Physical Review E* 82 (2010) 1–11.
- [3] B. G. Levine, D. N. LeBard, R. DeVane, W. Shinoda, A. Kohlmeyer, M. L. Klein, Micellization studied by GPU-accelerated coarse-grained molecular dynamics, *Journal of Chemical Theory and Computation* 7 (2011) 4135–4145.
- [4] A. Chremos, P. M. Chaikin, R. a. Register, A. Z. Panagiotopoulos, Shear-induced alignment of lamellae in thin films of diblock copolymers, *Soft Matter* 8 (2012) 3803.
- [5] A. Chremos, P. M. Chaikin, R. A. Register, A. Z. Panagiotopoulos, Sphere-to-Cylinder Transitions in Thin Films of Diblock Copolymers under Shear: The Role of Wetting Layers, *Macromolecules* 45 (2012) 4406–4415.
- [6] D. N. LeBard, B. G. Levine, R. DeVane, W. Shinoda, M. L. Klein, Premicelles and monomer exchange in aqueous surfactant solutions above and below the critical micelle concentration, *Chemical Physics Letters* 522 (2012) 38–42.
- [7] D. N. LeBard, B. G. Levine, P. Mertmann, S. A. Barr, A. Jusufi, S. Sanders, M. L. Klein, A. Z. Panagiotopoulos, Self-assembly of coarse-grained ionic surfactants accelerated by graphics processing units, *Soft Matter* 8 (2012) 2385–2397.
- [8] B. Li, Y.-L. Zhu, H. Liu, Z.-Y. Lu, Brownian dynamics simulation study on the self-assembly of incompatible star-like block copolymers in dilute solution., *Physical chemistry chemical physics : PCCP* 14 (2012) 4964–70.
- [9] J. Zhang, Z.-Y. Lu, Z.-Y. Sun, Self-assembly structures of amphiphilic multiblock copolymer in dilute solution, *Soft Matter* 9 (2013) 1947.
- [10] Y. Wang, B. Li, Y. Zhou, Z. Lu, D. Yan, Dissipative particle dynamics simulation study on the mechanisms of self-assembly of large multimolecular micelles from amphiphilic dendritic multiarm copolymers, *Soft Matter* 9 (2013) 3293.
- [11] D. Reith, A. Milchev, P. Virnau, K. Binder, Anomalous structure and scaling of ring polymer brushes, *EPL (Europhysics Letters)* 95 (2011) 28003.
- [12] D. Reith, L. Mirny, P. Virnau, GPU Based Molecular Dynamics Simulations of Polymer Rings in Concentrated Solution: Structure and Scaling, *Progress of Theoretical Physics Supplement* 191 (2011) 135–145.
- [13] J. Glaser, J. Qin, P. Medapuram, M. Mueller, D. Morse, Test of a scaling hypothesis for the structure factor of disordered diblock copolymer melts, *Soft Matter* 8 (2012) 11310–11317.
- [14] D. Reith, A. Milchev, P. Virnau, K. Binder, Computer Simulation Studies of Chain Dynamics in Polymer Brushes, *Macromolecules* 45 (2012) 4381–4393.
- [15] C. L. Phillips, C. R. Iacovella, S. C. Glotzer, Stability of the double gyroid phase to nanoparticle polydispersity in polymer-tethered nanosphere systems, *Soft Matter* 6 (2010) 1693.
- [16] T. D. Nguyen, E. Jankowski, S. C. Glotzer, Self-Assembly and Reconfigurability of Shape-Shifting Particles., *ACS nano* 5 (2011) 8892–8903.
- [17] J. Zhang, Z.-Y. Lu, Z.-Y. Sun, Self-assembly of amphiphilic patchy particles with different cross-linking densities, *Soft Matter* (2012) 3073–3080.
- [18] C. Knorowski, A. Travesset, Nanorods in functionalized block-copolymer gels: Flexible ladders and liquid crystalline order in curved geometries, *EPL (Europhysics Letters)* 100 (2012) 56004.
- [19] K. L. Kohlstedt, M. Olvera de la Cruz, G. C. Schatz, Controlling Orientational Order in 1-D Assemblies of Multivalent Triangular Prisms, *The Journal of Physical Chemistry Letters* 4 (2013) 203–208.
- [20] B. Buesser, A. J. Grohn, S. E. Pratsinis, Sintering Rate and Mechanism of TiO₂ Nanoparticles by Molecular Dynamics, *The Journal of Physical Chemistry C* 115 (2011) 11030–11035.
- [21] J. R. Spaeth, I. G. Kevrekidis, A. Z. Panagiotopoulos, Dissipative particle dynamics simulations of polymer-protected nanoparticle self-assembly., *The Journal of chemical physics* 135 (2011) 184903.
- [22] J. Zhang, Z.-Y. Lu, Z.-Y. Sun, A possible route to fabricate patchy nanoparticles via self-assembly of a multiblock copolymer chain in one step, *Soft Matter* 7 (2011) 9944–9950.
- [23] P. Guo, R. Sknepnek, M. Olvera de la Cruz, Electrostatic-Driven Ridge Formation on Nanoparticles Coated with Charged End-Group Ligands, *The Journal of Physical Chemistry C* 115 (2011) 6484–6490.
- [24] S. a. Barr, a. Z. Panagiotopoulos, Conformational transitions of weak polyacids grafted to nanoparticles., *The Journal of chemical physics* 137 (2012) 144704.
- [25] I. C. Pons-Siepermann, S. C. Glotzer, Design of patchy particles using quaternary self-assembled monolayers., *ACS nano* 6 (2012) 3919–3924.
- [26] Y. V. Kalyuzhnyi, C. R. Iacovella, H. Docherty, M. Holovko, P. T. Cummings, Network Forming Fluids: Yukawa Square-Well m-Point Model, *Journal of Statistical Physics* 145 (2011) 481–506.
- [27] A. Keys, L. Hedges, J. Garrahan, S. Glotzer, D. Chandler, Excitations Are Localized and Relaxation Is Hierarchical in Glass-Forming Liquids, *Physical Review X* 1 (2011) 021013.
- [28] D. Fragiadakis, C. Roland, Molecular dynamics simulation of the Johari-Goldstein relaxation in a molecular liquid, *Physical Review E* 86 (2012) 1–5.
- [29] L. Wan, C. Iacovella, T. Nguyen, H. Docherty, P. Cummings, Confined fluid and the fluid-solid transition: Evidence from absolute free energy calculations, *Physical Review B* 86 (2012) 214105.
- [30] N. Gnan, C. Maggi, G. Parisi, F. Sciortino, Generalized Fluctuation-Dissipation Relation and Effective Temperature Upon Heating a Deeply Supercooled Liquid, *Physical Review Letters* 110 (2013) 035701.
- [31] M. A. a. Barbosa, E. Salcedo, M. C. Barbosa, Multiple liquid-liquid critical points and density anomaly in core-softened potentials, *Physical Review E* 87 (2013) 032303.
- [32] N. a. Mahynski, B. Irick, A. Z. Panagiotopoulos, Structure of phase-separated athermal colloid-polymer systems in the protein limit, *Physical Review E* 87 (2013) 022309.
- [33] C. Knorowski, S. Burleigh, A. Travesset, Dynamics and Statics of DNA-Programmable Nanoparticle Self-Assembly and Crystallization, *Physical Review Letters* 106 (2011).
- [34] C. Knorowski, A. Travesset, Materials design by DNA programmed self-assembly, *Current Opinion in Solid State and Materials Science* 15 (2011) 262–270.
- [35] T. I. N. G. Li, R. Sknepnek, R. J. Macfarlane, C. a. Mirkin, M. Olvera de la Cruz, Modeling the crystallization of spherical nucleic Acid nanoparticle conjugates with molecular dynamics simulations., *Nano letters* 12 (2012) 2509–14.
- [36] T. I. N. G. Li, R. Sknepnek, M. Olvera de la Cruz, Thermally active hybridization drives the crystallization of DNA-functionalized nanoparticles., *Journal of the American Chemical Society* 135 (2013) 8535–41.
- [37] M. Bertrand, B. Joós, Extrusion of small vesicles through nanochannels: A model for experiments and molecular dynamics simulations, *Physical Review E* 85 (2012) 1–8.
- [38] N. Nguyen, E. Jankowski, S. Glotzer, Thermal and athermal three-dimensional swarms of self-propelled particles, *Physical Review E* 86 (2012) 1–9.
- [39] D. Wilms, P. Virnau, S. Sengupta, K. Binder, Langevin dynamics simula-

- tions of a two-dimensional colloidal crystal under confinement and shear, *Physical Review E* 85 (2012) 1–10.
- [40] I. Morozov, A. Kazennov, R. Bystryi, G. Norman, V. Pisarev, V. Stegailov, Molecular dynamics simulations of the relaxation processes in the condensed matter on GPUs, *Computer Physics Communications* 182 (2011) 1974–1978.
- [41] G. I. Guerrero-García, P. González-Mozuelos, M. Olvera de la Cruz, Potential of mean force between identical charged nanoparticles immersed in a size-asymmetric monovalent electrolyte., *The Journal of chemical physics* 135 (2011) 164705.
- [42] C. R. Iacovella, G. Varga, J. Sallai, S. Mukherjee, A. Ledeczi, P. T. Cummings, A model-integrated computing approach to nanomaterials simulation, *Theoretical Chemistry Accounts* 132 (2012) 1315.
- [43] P. Zhao, L. J. Yang, Y. Q. Gao, Z.-Y. Lu, Facile implementation of integrated tempering sampling method to enhance the sampling over a broad range of temperatures, *Chemical Physics* 415 (2013) 98–105.
- [44] N. Tretyakov, M. Müller, D. Todorova, U. Thiele, Parameter passing between molecular dynamics and continuum models for droplets on solid substrates: the static case., *The Journal of chemical physics* 138 (2013) 064905.
- [45] A. Kapoor, A. Travasset, Folding and stability of helical bundle proteins from coarse-grained models., *Proteins* 81 (2013) 1200–11.
- [46] J. A. Anderson, C. D. Lorenz, A. Travasset, General purpose molecular dynamics simulations fully implemented on graphics processing units, *Journal of Computational Physics* 227 (2008) 5342–5359.
- [47] T. D. Nguyen, C. L. Phillips, J. A. Anderson, S. C. Glotzer, Rigid body constraints realized in massively-parallel molecular dynamics on graphics processing units, *Computer Physics Communications* 182 (2011) 2313–2307.
- [48] NVIDIA web page, <http://www.nvidia.com>, 2012.
- [49] 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.
- [50] 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* (2012).
- [51] 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.
- [52] 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.
- [53] M. Taufer, N. Ganesan, S. Patel, GPU enabled Macromolecular Simulation: Challenges and Opportunities, *IEEE Computing in Science and Engineering (CiSE)* 15 (2013) 56–64.
- [54] GROMACS, <http://www.gromacs.org/>, 2013.
- [55] P. Eastman, V. S. Pande, Efficient nonbonded interactions for molecular dynamics on a graphics processing unit, *Journal of Computational Chemistry* 31 (2010) 1268–72.
- [56] 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.
- [57] 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.
- [58] 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.
- [59] E. Bitzek, P. Koskinen, F. Gähler, M. Moseler, P. Gumbsch, Structural Relaxation Made Simple, *Physical Review Letters* 97 (2006).
- [60] D. C. Rapaport, Enhanced molecular dynamics performance with a programmable graphics processor, *Computer Physics Communications* 182 (2011) 926–934.
- [61] Doxygen, <http://www.doxygen.org>, 2012.
- [62] S. Plimpton, Fast Parallel Algorithms for Short-Range Molecular Dynamics, *J. Comp. Phys.* 117 (1995) 1–19.
- [63] D. Frenkel, B. Smit, *Understanding Molecular Simulations*, Academic

Press, 2 edition, 2001.