

An implementation of the smooth Particle-Mesh Ewald method on GPU hardware

M. J. Harvey^{1,*} and G. De Fabritiis^{2,†}

¹*High Performance Computing Service, Information and Communications Technologies,
Imperial College London, South Kensington, London, SW7 2AZ, UK*

²*Computational Biochemistry and Biophysics Lab (GRIB-IMIM),
Universitat Pompeu Fabra, Barcelona Biomedical Research Park (PRBB),
C/ Doctor Aiguader 88, 08003 Barcelona, Spain*

The smooth particle-mesh Ewald summation method is widely used to efficiently compute long-range electrostatic force terms in molecular dynamics simulations and there has been considerable work in developing optimised implementations for a variety of parallel computer architectures. We describe an implementation for Nvidia graphical processing units (GPUs) which are general purpose computing devices with a high degree of intrinsic parallelism and arithmetic performance. We find that, for typical biomolecular simulations (eg DHFR, 26k atoms), a single GPU-equipped workstation is able to provide sufficient performance to permit simulation rates of ≈ 50 ns/day when used in conjunction with the ACEMD molecular dynamics package¹ and exhibits an accuracy comparable to that of a reference double-precision CPU implementation.

Keywords: Smooth particle-mesh Ewald, GPU, molecular dynamics

I. INTRODUCTION

Atomistic molecular mechanics simulations are widely used in the study of a range of biomolecular and inorganic molecular systems. However, the $O(N^2)$ scaling of the electrostatic interactions, combined with the explicit timestep scheme, make it challenging to access the microsecond timescale. To reduce the computational cost, the contribution of longer-range interactions is typically approximated by using a scheme with more favourable scaling properties, such a derivative of the Ewald summation method.² The smoothed particle-mesh Ewald summation (SPME),³ which scales with $O(N \log N)$ is perhaps the most widely used of these methods in popular MD packages.

Production molecular dynamics (MD) simulations are typically performed on highly parallel multi-processor or cluster supercomputers. Parallelisation of the directly calculated non-bonded electrostatic interactions may be efficiently achieved by using spatial decomposition methods, such as the neutral territory schemes of Bowers *et al*⁴ as only local neighbour communication is required. Parallelisation of the PME method is more challenging because of the long-range communications pattern required by the 3D Fourier transform step and, consequently, the development of efficient implementations remains an active topic of research.^{5,6}

Due to the computing resources required to perform MD simulations, there has been considerable research into accelerating the computation by using, for example dedicated, specialised hardware⁷ or commodity, high-performance accelerator processors.⁸ In recent years, commodity graphics processing units (GPUs) have acquired non-graphical, general-purpose programmability and have undergone a doubling of computational power every 12 months. Of the devices currently available on the market, those produced by Nvidia offer the most mature programming environment, the so-called compute

unified device architecture (CUDA),⁹ and have been the focus of the majority of investigation in the computational science field.

Several groups have lately shown results for MD codes which are accelerated by use of CUDA-capable GPUs, for example^{10–14} and it has been recognised that high-performance, low-cost GPU-accelerated molecular dynamics simulations could have a significant practical impact in the field of *in silico* drug discovery.¹⁵ In recent work we presented ACEMD, an MD package designed to run on high-performance Nvidia GPUs.^{1,16} ACEMD exhibits performance comparable to clusters of $\approx 100 - 200$ CPUs and enables microsecond-scale simulations on a single GPU-equipped workstation, representing a significant improvement in the accessibility and cost-efficiency of MD simulation. A principal feature of ACEMD is its SPME treatment of long-range electrostatic terms as this is generally considered a requirement for production-quality biomolecular simulations.

In this paper we focus on the GPU implementation of the SPME scheme used in ACEMD and discuss the resulting performance.¹ The SPME calculation is performed on a single GPU but utilises the device's massive intrinsic parallelism, necessitating significant modification of the implementation in comparison to a serial or parallel block-decomposed version for conventional CPUs.

II. ALGORITHM AND IMPLEMENTATION

A. The Smooth Particle-Mesh Ewald Method

The smooth particle-mesh Ewald (SPME) method of Essman and coworkers¹⁷ is an efficient scheme for computing long-range electrostatic forces which exhibits improved $O(N \log N)$ scaling in comparison to the $O(N^2)$ scaling of standard Ewald summation, differing from the

latter in the use of interpolation and the 3D FFT in computing the reciprocal space summation.

The electrostatic interaction energy of an N particle system is given by the sum of the pairwise Coulombic interactions:

$$E = \sum_{i=1}^{N-1} \sum_{j>i}^N \frac{q_i q_j}{r_{ij}}, \quad (1)$$

where q_i is the charge of the i th particle and r_{ij} the distance between the i th and j th particles. The traditional Ewald summation method re-expresses E as the sum of three potentials $E = E_{dir} + E_{rec} + E_{self}$, where E_{dir} is summed over all pairs in real space, E_{rec} is summed in reciprocal space and E_{self} is a corrective term representing the self-energy of the system. These three potentials are defined as:

$$E_{dir} = \sum_{i=1}^{N-1} \sum_{j>i}^N \frac{q_i q_j \operatorname{erfc}(\beta r_{ij})}{r_{ij}}, \quad (2)$$

$$E_{rec} = \frac{1}{2\pi V} \sum_{\vec{m} \neq 0} \frac{\exp(-\pi^2 \vec{m}^2 / \beta^2)}{\vec{m}^2} S(\vec{m}) S(-\vec{m}), \quad (3)$$

$$E_{self} = -\frac{\beta}{\sqrt{\pi}} \sum_i^N q_i^2, \quad (4)$$

where $\operatorname{erfc}(x)$ is the complementary error function $\operatorname{erfc}(x) = 1 - \operatorname{erf}(x)$, β is the Ewald parameter, $\vec{m} = (m_1, m_2, m_3)$ are the reciprocal-space lattice vectors and V is volume of the unit cell in reciprocal space. $S(\vec{m})$ is the lattice structure factor, given by

$$S(\vec{m}) = \sum_j q_j \exp(\vec{m} \cdot \vec{r}_j). \quad (5)$$

The SPME method approximates this with:

$$S(\vec{m}) \approx \tilde{S}(\vec{m}) = b_1(m_1) b_2(m_2) b_3(m_3) F(Q)(m_1, m_2, m_3), \quad (6)$$

where $b_i(m_i)$ are Euler exponential splines given by

$$b_\alpha(m_\alpha) = \exp(2\pi i(n-1)m_\alpha/K_\alpha) \times \left[\sum_{k=0}^{n-2} M_n(m+1) \exp(2\pi i m_\alpha k / K_\alpha) \right]^{-1}, \quad (7)$$

for $\alpha = 1, 2, 3$. $F(Q)$ is the discrete Fourier transform of the 3D matrix Q of dimension $K_1 \times K_2 \times K_3$ which is given by

$$Q(k_1, k_2, k_3) = \sum_{i=0}^N \sum_{p_\alpha=0}^{K_\alpha-1} q_i M_n(u_i^1 - k_1 - p_1 K_1) \times M_n(u_i^2 - k_2 - p_2 K_2) \cdot M_n(u_i^3 - k_3 - p_3 K_3), \quad (8)$$

where \vec{u}_i is the scaled fractional coordinate of particle i within the bounds $0 \leq u_i^\alpha \leq K_\alpha$. M_n is the cardinal B-spline of order n . The summations are over all integers p_α in the range $0 \leq p_\alpha < K_\alpha$. Cardinal B-splines have the following recursive form:

$$M_n(u) = \frac{u}{n-1} M_{n-1}(u) + \frac{n-u}{n-1} M_{n-1}(u-1) \quad (9)$$

for $n \geq 2$. For $n = 2$, $M_2(u) = 1 - |u - 1|$ for $0 \leq u \leq 2$ else $M_2(u) = 0$. B-splines have compact support and the effect of applying Equation 8 is to ‘spread’ each charge q_i out over a cubic volume of n^3 elements of Q .

The approximate reciprocal space energy sum \tilde{E}_{rec} can be re-expressed as the convolution:

$$E_{rec} \approx \tilde{E}_{rec} = \frac{1}{2} \sum_{m_\alpha=0}^{K_\alpha-1} Q(m_1, m_2, m_3) \cdot (\theta_{rec} * Q)(m_1, m_2, m_3), \quad (10)$$

where

$$\theta_{rec} = F(B \cdot C), \quad (11)$$

and B and C are the arrays given by

$$B(m_1, m_2, m_3) = \prod_{i=1}^3 |b_i(m_i)|^2, \quad (12)$$

$$C(m_1, m_2, m_3) = \frac{1}{\pi V} \frac{\exp(-\pi^2 \vec{m}^2 / \beta^2)}{\vec{m}^2}. \quad (13)$$

From which the per-particle forces may be derived by differentiation with respect to \vec{r}_i . As Q is continuously differentiable $n - 2$ times with respect to the particle positions and θ_{rec} is independent of the particle positions, the atomic force terms are given by

$$\frac{\delta \tilde{E}_{rec}}{\delta r_i^\alpha} = \sum_{m_\alpha=0}^{K_\alpha-1} \frac{\delta Q}{\delta r_i^\alpha}(m_1, m_2, m_3) \cdot (\theta_{rec} * Q)(m_1, m_2, m_3) \quad (14)$$

B. Nvidia GPU architecture

Contemporary Graphics Processing Units (GPU) produced by Nvidia Corp. are general purpose, single-program multiple data (SPMD) processing units with a high degree of intrinsic parallelism and arithmetic performance: for example, over 30k execution threads may run concurrently on the Tesla C1060 device, with a peak arithmetic rate of ≤ 933 single precision GFLOPS. These devices are composed of a set of independent *multiprocessors*, each of which contains 8 scalar cores. A scalar program fragment, known as a *kernel* may be run concurrently in a *block* of *threads* in parallel *warps* of 32

threads. Each block is restricted to executing on a single multiprocessor. Multiple independent blocks may be executed concurrently across the device in a *grid*.

Threads within a common block may intercommunicate via a small 16KB region of in-core shared memory. Each multiprocessor has a large register file, of which each thread receives a private, static allocation. The total number of threads which may execute on a multiprocessor – and thus the degree of parallelism – is dependent on the register resources required by each individual thread and also the shared memory required by the block. Current devices support up to 1024 threads and up to 8 blocks per multiprocessor.

Multiprocessors are grouped together on a single package and share a high-bandwidth link to external DRAM, known as *global memory*. Access to this memory is uncached, and so is particularly costly, although the interleaved execution of warps acts to mitigate the cost of this latency.

The current programming model for these GPUs is CUDA, a C-like language with type qualifier extensions for indicating data locality and a special function-call syntax for specifying the parallelism of a kernel invocation. An associated API provides functions for managing GPU-host memory allocation and transfer. CUDA code is compiled to an intermediate byte-code which is interpreted at runtime, providing forward compatibility with future device architectures.

For further details on device capabilities and programming model, the reader is referred to the CUDA SDK documentation.⁹

C. SPME Implementation

The implementation of SPME has two distinct components: the real-space evaluation of the pairwise E_{dir} term and the reciprocal space evaluation of \tilde{E}_{rec} . The evaluation of the former is trivially implemented into an existing non-bonded electrostatic force computation program with a change of potential function¹ and we do not discuss this further. \tilde{E}_{rec} , however, involves several different computational steps and its calculation frequently becomes the limiting step in parallel molecular dynamics simulations.^{6,18}

In this section we describe in detail the implementation of the computation of \tilde{E}_{rec} using CUDA. No part of the computation is performed by the host CPU, nor is it necessary to transfer any state between the GPU and host memory. In this way the performance of the code is entirely dependent upon the specification of the GPU.

The calculation of \tilde{E}_{rec} may be divided into 5 steps, which we summarise below before describing their implementation:

1. *Charge spreading*: Spreading of charges on to the array Q (eq 8). The scaled fractional coordinates of each particle are calculated and the Q array popu-

lated with n^3 terms computed from the B-spline coefficients $M_n(u_i^\alpha - j)$, where $i = 1, \dots, N$, $\alpha = 1, 2, 3$ and $j = 0, \dots, n$.

2. *3DFFT*: real-to-complex 3DFFT in-place transform of Q into reciprocal space.
3. *Energy computation*: application of eq 10 in reciprocal space, to compute terms for \tilde{E}_{rec} . Q is replaced by the product of itself with array B and C , as defined in equations 12 and 13.
4. *3DFFT*: computation of the convolution θ_{rec} (eq 11) through a complex-to-real 3DFFT of the array resulting from step 3.
5. *Force computation*: real-space computation of per-atom force terms $\delta\tilde{E}_{rec}/\delta r_i^\alpha$ by multiplication by $\delta Q/\delta r_i^\alpha$ (eq 14).

Charge spreading. In this step, each charge q_i is mapped to a site on the real-space PME grid at a location determined by particles scaled fractional coordinates \tilde{r}_i . The charge is then distributed over points in a surrounding volume according to eq 8. The volume of the spreading region is dependent on the order n of the cardinal B-spline. For $n = 4$, typically used in production simulations, each charge is spread over $n^3 = 64$ grid points.

This spreading is straightforward to do in a serial implementation but poses a significant challenge when performed with fine-grained parallelism. A naïve implementation which used one thread to map the spread of each individual charge onto the grid would encounter synchronisation problems when different threads attempt to accumulate charge to the same grid location, thus necessitating the use of thread-safe atomic memory operations²³. However, as only integer atomic operations are supported on current hardware, either floating-point atomic operations must be emulated using an atomic compare-and-set loop construct, or the charge spreading altered to use fixed-precision arithmetic. Furthermore, in order to achieve acceptable performance from the uncached GPU memory subsystem, it is necessary for the threads within a block to perform memory accesses to contiguous address ranges. The naïve approach has an essentially unordered memory access pattern, leading to poor performance.

Instead, rather than performing a charge spreading, a per grid-point gather is used. This is conducted in three steps which we term *placement*, *accumulation* and *overflow*. First, each particle is mapped to a grid location and its charge and position are recorded in a three dimensional array. At the PME grid sizes and system densities typical of biomolecular simulation, this array is 90% sparse. Nevertheless, because the placement is performed in parallel, one particle per thread, the setting of elements within this array must still be performed atomically, using the global memory atomic operation primitives, in order to prevent access conflicts. Each grid site

is permitted to hold a single charge; additional charges are placed in a simple list referred to as the overflow list.

Second, a separate kernel is used to sum the charge at each grid point by finding the contributions from all charges within the surrounding n^3 points (Figure 1). This kernel is executed in a block of K_1 threads which operates on a single full-grid-width stencil along the m_x vector. Each thread i accumulates the charge for the i th grid point in the row. Each thread loads in turn the n^2 surrounding points in the $y-z$ plane and, if containing a particle, the associated $3n$ spline terms $M_n(u_i^\alpha - j)$ (for $\alpha = 1, 2, 3, j = 0, \dots, n$) are computed and placed into shared memory from where they may be accessed by other threads.

This use of shared memory reduces the number of global memory accesses by a factor of n and has a memory access pattern favourable for the hardware. Although this method requires each set of spline terms to be computed n^2 , the arithmetic cost of the calculation is sufficiently low to make their re-computation preferable to loading them from a precomputed table in global memory.

Lastly, the charges in the overflow list are spread onto the grid using the naïve method described initially and using floating-point atomic memory operations synthesised with an atomic integer compare-and-set loop construct. Although this processing is slow, it is much faster than a second iteration of the gathering kernel.

One consequence of the use of the per-grid-point gather is that the scaling characteristic of the charge-gridding phase changes from $O(N)$ to $O(K^3)$.

3DFFT. 3D FFTs are performed using the CUDA cuFFT library.¹⁹ Although this library provides 3D FFT routines, the current version (version 2.2) shows poor performance. Consequently, we implement the multidimensional FFT explicitly using three sets of orthogonal 1D transforms. The transforms are interleaved with in-place transpositions of the charge array Q which are necessary to satisfy the contiguous memory layout of the input data required by the FFT library.

Energy computation. $b(m_i)$ and exponential terms for the Euler exponential splines (eq 7) are precomputed once and stored in the GPU's *constant* memory. This is a small (64KB), read-only region of global memory that has cached read-access for reduced access latency. Calculation of the product of the transformed Q with $B \cdot C$ (eq 11) is then performed in-place using a kernel that computes one column of values along m_z per thread. When required for diagnostic output, energy terms can be accumulated in a separate buffer and the sum transferred back to the host.

Force computation. The computation of $\delta \tilde{E}_{rec} / \delta r_i^\alpha$ is performed using a kernel which operates on a single particle per thread. The $3n$ spline terms M_n are again computed, along with the first derivatives M'_n , once per particle and n^3 force terms are computed as per eq 14. Because of the particle distribution is essentially unordered memory accesses are uncoalesced, but the kernel

is sufficiently simple to permit a large number of threads per block, effectively hiding the memory access latency.

III. DISCUSSION

A. Performance

The performance was also measured on production simulations of models of dihydrofolate reductase (DHFR) ($62 \times 62 \times 62 \text{ \AA}^3$, 23,558 atoms) and apoA1 ($108 \times 108 \times 78 \text{ \AA}^3$, 92,224 atoms), two protein systems commonly used for benchmarking exercises.¹ The DHFR and apoA1 models are typical of the scale of system routinely simulated using all-atom biomolecular mechanics, with 1 million atom simulations of complete virions (eg the satellite tobacco mosaic virus, STMV)²⁰ at the far extreme, representing a range of simulation cell sizes from $\approx 60^3 - 220^3 \text{ \AA}^3$. We determine the practical performance in timesteps/day achievable with production simulation parameters and ACEMD, as shown in Table I.

ACEMD is a parallel code able to distribute all other aspects of an MD simulation (bonded, nonbonded force terms, etc) over multiple GPUs and, in practice, the PME computation represents the performance-limiting critical path in simulations over this range of sizes making the optimisation of the PME implementation essential for best performance.

To more closely assess the performance characteristics of the code, it was tested over a range of synthetic cubic input systems constructed with linear dimension in the range $16 \leq l_{sim} \leq 256 \text{ \AA}$ and with a random uniform distribution of particles with number density ρ_N of $0.1 \text{ atoms \AA}^{-3}$, representative of the sizes and density of biomolecular systems of interest.

Execution times, $T_{kernels}$, for each individual kernel were recorded using the standard CUDA runtime profiling tool. The total execution wall-time, T_{PME} , was measured from start to finish and so includes all additional host(CPU)-side overhead.

K_α were initially chosen to be equal to $\text{int}(l_{sim})$, as is commonly used in production MD simulations. For each synthetic system size, 100 iterations were performed, each with a different random particle distribution in order to characterise the sensitivity of the implementation to fluctuations in the distribution (Figure 2a)²⁴.

For large K , T_{PME} exhibits cubic scaling with respect to the PME grid dimension (dashed lines in Figure 3), corresponding to a broadly linear scaling with respect to particle number. However, for small grid sizes and particle counts, $K < 50$, there is insufficient available parallelism to fully occupy the GPU resources. This can be seen more clearly in the component kernel timings in Figures 2a and 2b, most prominently in the co-ordinate scaling, and energy calculation and overflow list kernels. In the case of the first two, the kernels are computationally simple and a high thread occupancy may be achieved: for example, the co-ordinate scaling ker-

nel may run 15360 threads (512 threads/multiprocessor) simultaneously, corresponding to a test system size of $K = 53$. For smaller systems, there is insufficient parallelism to fully occupy the GPU and the execution time of the kernel is weakly dependent on system size: all blocks may be spread out across multiprocessors and executed completely in parallel. At $K = 53$, the GPU is fully subscribed and a step in the execution time can be seen, indicating that some multiprocessors have now to process a second block. After this point, the execution time scales approximately linearly with particle count, as subsequent step changes are much less significant owing to divergence between the work executing on the separate multiprocessors.

In the case of the overflow and charge placement kernels, a significant variation in runtime is observable — this arises from collisions between atomic operations from different threads. The synthesised floating-point atomic operations used in the overflow kernel are much more sensitive to collisions as they use a looping construct that may make several access to global memory. Because a warp of threads is executed synchronously, it requires only one thread to suffer an unfavourable memory access for the execution time of the entire block to increase.

For the charge accumulation and energy calculation kernels, the number of threads per block is directly related to K_1 . Initially, several blocks may run concurrently per multiprocessor but as K_1 increases the occupancy falls.

We use the Nvidia-supplied CUDA FFT library. As is characteristic of the FFT algorithm, the computational cost is highly sensitive to the prime factorisation of the input dimension (Figure 2a). For this reason, some PME implementations, such as that of DESMOND²¹ are designed to use only specific grid sizes (2^n in that case). However, for the relatively small systems ($O(100k)$ atoms) that ACEMD is intended to simulate, this would be quite restrictive. Therefore, for each l_{sim} , we select a modified K in the range $l_{sim} \leq K \leq 1.5l_{sim}$ such that T_{PME} is minimised. By selecting for the minimum T_{PME} , rather than just minimum FFT cost, we compensate for the increased runtime of components of the implementation which are sensitive to the ρ_N and K of the input. $T_{kernels}$ for these optimised dimensions are shown in Figure 2b and T_{PME} in Figure 3. The saw-tooth pattern of the charge placement and overflow kernels are indicative of atomic operation collisions, the probability of them occurring dropping as K becomes larger than $int(l_{sim})$.

Because the accumulation and energy kernels operate preferentially along one dimension of the PME grid, we test the sensitivity of the implementation to the orientation of the input for cuboid systems. Figure 4 shows the execution time for synthetic input systems of constant ρ_N , varying each dimension l_α (for $\alpha = 1, 2, 3$) in turn in the range $0.5l_{sim} \leq l_\alpha \leq 2l_{sim}$ for $l_{sim} = 64, 128$ and using the predetermined ‘optimised’ grid dimensions.

Whilst there are differences between the three variations, reflecting differing grid sizes and multiprocessor occupancies, we conclude that these are insufficiently large to require care over the orientation of the input system.

Finally, we consider the effect of the spline interpolation order on performance. When optimising the SPME parameters for production MD simulations, it is not unusual to reduce grid size and increase the spline interpolation order to optimise performance or control accuracy. We tested the performance impact of these changes using two commonly used biomolecular benchmarks - DHFR (26k atoms), and apoA1 (92k atoms). In each case, the PME grid size was varied between $0.5l_{sim} < K_\alpha < 2l_{sim}$ and timings taken (Figure 3, green(squares) and blue(diamonds) lines). It can be seen that T_{PME} is minimal at approximately $K_\alpha = 0.6l_{sim}$, increasing for smaller grids as the overflow list processing becomes significant. However, increasing the interpolation order to 6 significantly increases the execution time of the charge accumulation (Figure 3, magenta line) and no performance increase is yielded.

B. Accuracy

The GPU code exclusively uses single-precision floating-point arithmetic and thus operates at a reduced precision in comparison to a typical double-precision CPU code. To quantify any reduction in accuracy we compared the force terms produced by the GPU code (using ACEMD)¹ with a reference double precision CPU code (NAMD)²² when performing the joint Amber-Charmm DHFR benchmark. The relative error of the reciprocal force terms, $(|F_{namd}| - |F_{acemd}|)/|F_{namd}|$, was found to be 10^{-5} , significantly below the 10^{-3} generally considered the acceptable maximum for relative error in the force terms for biomolecular simulation.²¹ As such, we consider single precision arithmetic to be adequate for production simulation.

For long simulations in the NVE ensemble in which energy conservation is important, the code could be converted to use the double precision arithmetic capabilities of the latest GPUs. As this would incur a significant performance penalty (up to $\approx 8\times$ slower) and is not presently necessary for our work this remains an area for future investigation. Furthermore, a double precision version of the cuFFT library is not yet available (as of CUDA 2.2).

Further data on the accuracy of benchmark MD simulations performed with PME in conjunction with ACEMD are given in Section IV of.¹

IV. CONCLUSIONS

We have implemented the smooth particle-mesh Ewald method on Nvidia GPU hardware using the CUDA programming language. The implementation has been in-

tegrated into ACEMD,¹ permitting all-atom biomolecular simulations with long-range electrostatics to be fully-accelerated on GPU devices with an accuracy comparable to a reference double-precision CPU implementation. On contemporary hardware, a high level of performance is observed, commensurate with molecular dynamics simulation rates between 100 ns/day for systems of a few thousand atoms and 1ns/day for million atom systems ,

when used within ACEMD .

Acknowledgements This work was partially funded by the HPC-EUROPA project (R113-CT-2003-506079). GDF acknowledges support from the Ramon y Cajal scheme and from the EU Virtual Physiological Human Network of Excellence. We gratefully acknowledge Nvidia Corporation (<http://www.nvidia.com>) for their hardware donations.

-
- * Electronic address: m.j.harvey@imperial.ac.uk
 † Electronic address: gianni.defabritiis@upf.edu
- ¹ Harvey, M. J.; De Fabritiis, G.; Giupponi, G. *J. Chem. Theor. Comp.* **2009**, in press.
 - ² Ewald, P. *Ann. Phys.* **1921**, *369*, 253–287.
 - ³ Essman, U.; Perera, L.; Berkowitz, M. L.; Darden, T.; Lee, H.; Pedersen, L. G. *J. Chem. Phys.* **1995**, *19*, 8577–8593.
 - ⁴ Bowers, K. J.; Dror, R. O.; Shaw, D. E. *J. Phys.: Conf. Series* **2005**, *16*, 300–304.
 - ⁵ Oh, K. J.; Deng, Y. *Comp. Phys. Commun.* **2007**, *177*, 426–431.
 - ⁶ Fang, B.; Martyna, G.; Deng, Y. *Comp. Phys. Commun.* **2007**, *177*, 362–377.
 - ⁷ Shaw, D. E. Anton, a Special-Purpose Machine for Molecular Dynamics Simulation. *Proceedings of the 34th annual international symposium on Computer architecture*, 2007.
 - ⁸ De Fabritiis, G. *Comp. Phys. Commun.* **2007**, *176*, 600.
 - ⁹ Nvidia Corp, *NVIDIA CUDA Compute Unified Device Architecture Programming Guide 2.0*; Technical Report, 2008.
 - ¹⁰ Anderson, J. A.; Lorenz, C. D.; Travesset, A. *J. Comp. Phys.* **2008**, *22*, 5342–5359.
 - ¹¹ Meel, J. A. V.; Arnold, A.; Frenkel, D.; Portegies, S. F.; Belleman, R. G. *Mol. Sim.* **2008**, *34*, 259,266.
 - ¹² Phillips, J. C.; Stone, J. E.; Schulten, K. Adapting a message-driven parallel application to GPU-accelerated clusters. *Proceedings of the 2008 ACM/IEEE Conference on Supercomputing*, 2008.
 - ¹³ Friedrichs, M. S.; Eastman, P.; Vaidyanathan, V.; Houston, M.; LeGrand, S.; Beberg, A. L.; Ensign, D. L.; Bruns, C. M.; Pande, V. S. *J. Comp. Chem.* **2009**, *30*, 864–872.
 - ¹⁴ Liu, W.; Schmidt, B.; Voss, G.; Müller-Wittig, W. *Comp. Phys. Commun.* **2008**, *179*, 634–641.
 - ¹⁵ Giupponi, G.; Harvey, M. J.; de Fabritiis, G. *Drug Discovery Today* **2008**, *13*, 1052.
 - ¹⁶ ACEMD home page, <http://multiscalelab.org/acemd>.
 - ¹⁷ Essman, U.; Perera, L.; Berkowitz, M. L.; Darden, T.; Lee, H.; Pedersen, L. G. *J. Chem. Phys.* **1995**, *103*, 8577.
 - ¹⁸ Toukmaji, A.; Paul, D.; Board Jr, J. *Distributed Particle-Mesh Ewald: A Parallel Ewald Summation Method*; Technical Report, 1996.
 - ¹⁹ NVidia Corp, *CUDA CUFFT Library, Document PG-00000-003 V2.0*; Technical Report, 2008.
 - ²⁰ Freddolino, P. L.; Arkhipov, A. S.; McPherson, S. B.; Schulten, K. *Structure* **2006**, *14*, 437–449.
 - ²¹ Bowers, K. J.; Chow, E.; Xu, H.; Dror, R. O.; Eastwood, M. P.; Gregersen, B. A.; Klepeis, J. L.; Kolosvary, I.; Moraes, M. A.; Sacerdoti, F. D.; Salmon, J. K.; Shan, Y.; Shaw, D. E. Scalable Algorithms for Molecular Dynamics Simulations on Commodity Clusters. *Proceedings of SuperComputing 2006, Tampa, US, 11-17th November*, 2006.
 - ²² Phillips, J. C.; Braun, R.; Wang, W.; Gumbart, J.; Tajkhorshid, E.; Villa, E.; Chipot, C.; Skeel, R. D.; Kale, L.; Schulten, K. *J. Comp. Chem.* **2005**, *26*, 1781–1802.
 - ²³ Atomic functions are memory operations that are designed to ensure that concurrent accesses to the same memory location by competing threads always leave the contents of that location in a consistent state by preventing race conditions. These functions require dedicated hardware support and have limited set of operations, as detailed in Appendix C of⁹
 - ²⁴ The test hardware was a single Nvidia Tesla C1060 on HP xw660 workstation (dual Xeon 5430, 4GB RAM, Red Hat Enterprise Linux 5.3, CUDA 2.1, Nvidia driver 180.22).

Figures




FIG. 1: Charge-spreading phase in the real-space charge grid. For efficient implementation, the charges are gathered to each grid point (shown as cells for clarity): for a spline order of n , each grid point receives charge terms from n^3 neighbouring points (red dashed region indicates the cubic charge import region for the yellow(light grey) point $Q(n, y, z)$). The charge gathering computational kernel operates in (K_y, K_z) blocks of K_x threads, each block calculating charge terms for a whole x row of the grid. Each thread loops over the n^2 neighbour cells in the $y-z$ plane (blue/dark grey region), calculating in turn the $3n$ B-spline terms M_n corresponding to any charge located in the cell and accumulating the charge contribution for its cell. The spline terms are held in shared memory, allowing each thread to reuse the values computed by its neighbours for computing charge contributions from the adjacent $(n-1)$ $y-z$ planes.

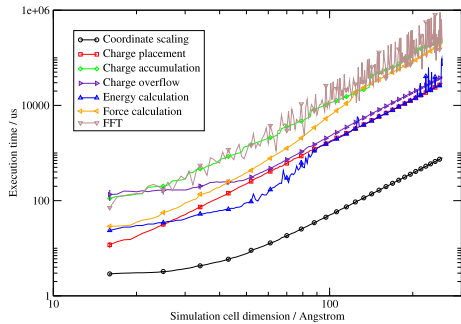
(a) $T_{kernels}$

FIG. 2: Execution times for the individual GPU kernels given as a function of the linear dimension of the simulation cell, for a constant system number density of 0.1\AA^{-3} . FFT timing (red dashed line) shows timings for the case of the PME grid having linear dimension $K = \text{int}(l_{sim})$, as determined by the method given in the main text. Diamonds indicate the

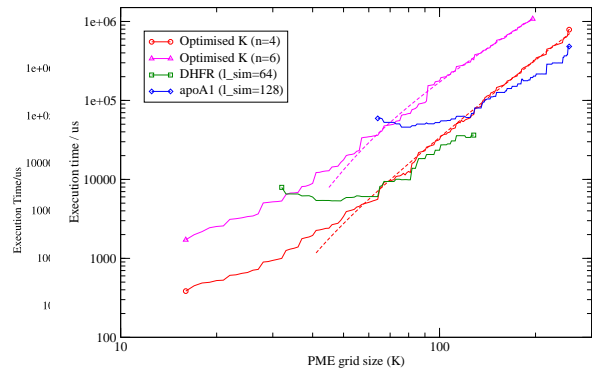


FIG. 3: T_{PME} timing showing the effect of varying the PME grid size for fixed system size K and of varying the B-spline order n . Magenta (circles) and red (triangles) lines show timings for the case of the PME grid having linear dimension $K = \text{int}(l_{sim})$, as determined by the method given in the main text. Diamonds indicate the simulated systems are of constant size (26k atoms DHFR and 92k apoA1 benchmarks respectively) and the ordinal indicates the PME grid dimension K .

Tables




FIG. 4: T_{PME} timing showing the minor sensitivity of the implementation to the orientation of the input for cuboid systems. The length of each simulation box dimension l_α is varied in turn within the range $l_{sim}/2 \leq l_\alpha \leq 2l_{sim}$ with the remaining sides held at a length of l_{sim} , where $l_{sim} = 64$ (solid lines) and 128 (dashed lines). A constant system density of $\rho_N = 0.1 \text{ \AA}^3$ is maintained.

Model	Atoms	PME grid size	10^6 timesteps/day	ns/day
DHFR	26k	65	12	49
ApoA1	92k	108	2.7	11
STMV	1M	220	0.17	0.69

TABLE I: Performance of the PME code in the parallel MD program ACEMD (running on 3 GTX280 GPUs) for a range of model sizes using production parameters of B-spline order $n = 4$ and timestep of 4fs. STMV performance estimated from timings given in Figure 3.