



Published in final edited form as:

J Chem Theory Comput. 2012 November 13; 8(11): 4624–4636. doi:10.1021/ct300324k.

Revisiting Molecular Dynamics on a CPU/GPU system: Water Kernel and SHAKE Parallelization

A. Peter Ruymgaart and Ron Elber*

Department of Chemistry and Biochemistry, Institute for Computational Engineering and Sciences, University of Texas at Austin, Austin, TX 78712

Abstract

We report Graphics Processing Unit (GPU) and Open-MP parallel implementations of water-specific force calculations and of bond constraints for use in Molecular Dynamics simulations. We focus on a typical laboratory computing-environment in which a CPU with a few cores is attached to a GPU. We discuss in detail the design of the code and we illustrate performance comparable to highly optimized codes such as GROMACS. Beside speed our code shows excellent energy conservation. Utilization of water-specific lists allows the efficient calculations of non-bonded interactions that include water molecules and results in a speed-up factor of more than 40 on the GPU compared to code optimized on a single CPU core for systems larger than 20,000 atoms. This is up four-fold from a factor of 10 reported in our initial GPU implementation that did not include a water-specific code. Another optimization is the implementation of constrained dynamics entirely on the GPU. The routine, which enforces constraints of all bonds, runs in parallel on multiple Open-MP cores or entirely on the GPU. It is based on Conjugate Gradient solution of the Lagrange multipliers (CG SHAKE). The GPU implementation is partially in double precision and requires no communication with the CPU during the execution of the SHAKE algorithm. The (parallel) implementation of SHAKE allows an increase of the time step to 2.0fs while maintaining excellent energy conservation. Interestingly, CG SHAKE is faster than the usual bond relaxation algorithm even on a single core if high accuracy is expected. The significant speedup of the optimized components transfers the computational bottleneck of the MD calculation to the reciprocal part of Particle Mesh Ewald (PME).

Keywords

computer simulations; constraints; DHFR; Optimization

Introduction

Extending time scales of molecular dynamics (MD) simulations makes it possible to address more complex molecular problems than was possible before. One approach to this problem utilizes advances in hardware to obtain faster propagation of the trajectory in time. Increases in computational speed of molecular dynamics algorithms have been obtained by utilizing multiple CPU's and/or specialized hardware such as graphics-processing units (GPU)^{1–5}. Parallelization is the dominant way nowadays to speed up MD using modern computer hardware. Speed-up is achieved while retaining the overall accuracy of the calculations. As simulations become more complex, and are more expensive and lengthy to conduct it becomes

Contact: Ron Elber, ron@ices.utexas.edu, phone: 512-232-5415, Fax: 512-471-8694.

Supporting Information: A code of the pre-conditioner conjugate gradient algorithm on the GPU (Appendix A), code for “in kernel parallel reduction” (Appendix B), and code for GPU water-water interactions (Appendix C) are provided. This information is available free of charge via the Internet at <http://pubs.acs.org>.

important to be able to assess their quality on the fly and avoid costly mistakes. One of the important measures of correctness is the conservation of the total energy (or other conserved entities). Moreover, if the focus of the calculations is on explicit dynamics and time scales then the NVE ensemble (N number of particles, V volume, and E energy) is the most natural choice. To sample from this ensemble correctly and to accurately compute the dynamics the energy must be conserved. This requirement is significant and not trivial to satisfy even for relatively short trajectories of only a few nanoseconds. If energy conservation is not adequately addressed the phenomenon of energy drift (energy change linear in time) is observed. Under these conditions, even if some physical observables are reasonably well behaved, or if the energy is adjusted empirically, it is difficult to accept the results without significant reservations. One of the focuses of our code MOIL^{5,6} is on the calculations of kinetics with the method of Milestoning⁷⁻¹¹. In Milestoning we launch a large number of short trajectories and use their statistics to estimate rates at long times. To ensure the correctness of our short trajectories we insist on relative energy conservation ($|\Delta E/E|$) which is at most 10^{-2} per microsecond for the standard benchmark of DHFR.⁵

At present our code is aimed at individual laboratories and typical laboratory hardware. Since our applications focuses on launching a very large number of short trajectories it makes sense for us to optimize the code of one trajectory for one node. We optimized our code for a CPU with 4 cores and a GPU. With this system size we can easily and efficiently study systems of about 100,000 atoms. Simulations of systems larger than this size will require further adjustments of the code to overcome current memory limitations.

We chose to split the computations between the GPU and the CPU since we wanted a stable programming model (to avoid re-writing significant portion of the code when new GPU hardware is introduced), and since we seek a programming platform that supports reasonably rapid code development. The GPU provides massively parallel computations for simple data and instructions and can be faster for specific tasks by a factor of several tens compared to a single core. However, it is not always the best choice and consideration of memory and accuracy are important. The GPU is most efficient if single precision is used. However, single precision accuracy is not always sufficient. While recent GPU versions support double precision computations there is a cost using this option (e.g. the variables may not fit in a cache or available registers). The evolving support level of GPU makes it difficult to retain a stable programming model, and extensive code changes may be required for new hardware variants, something that we wish to avoid to the extent possible. Having an option of computing in double precision when deemed necessary on the CPU is therefore an advantage. Moreover, GPU is significantly more difficult to code and for a program that we routinely add and change, programming on the GPU is costly in human and time resources. Complex code is better developed on the CPU. Finally, it is likely that exploitation of heterogeneous computing is possible for the CPU/GPU pair in which some of the calculations are done on the CPU and some on the GPU to obtain greater efficiency. Ignoring this possibility (by insisting on doing everything on the GPU) is limiting. We demonstrate the possibility of asynchronous computing in the paper. We show significant asynchronous speed up for SHAKE by SHAKing the water molecules on the CPU and the macromolecule on the GPU.

We believe the future of single node computing in general is heterogeneous. The AMD Fusion and Intel Sandybridge are already commonplace with AMD offering training for heterogeneous programming on its Fusion chips. Now also NVIDIA joins the heterogeneous CPU-GPU market with "Project Denver". We also think memory transfer between CPU/GPU will increase in performance in future systems. Calculations that are complex (require many registers), are few in number and involve numerous atoms per calculation are best suited for the CPU. Just like operations with many branch possibilities. PME, torsion and improper torsion calculations fall in this category and can be performed on CPU cores at the same time as the Non-bonded

calculations on the GPU. We believe in using them both, and using them both at the same time. Let the total computational cost be T and in our programming model a fraction $1-\alpha$ is computed on the GPU and a fraction α on the CPU. The compute time on the GPU is $T(1-\alpha)/G$ and on the CPU $T\alpha/C$ where the computational speeds of the CPU and the GPU are denoted C and G respectively. We also consider an overhead o of the communication between the CPU and the GPU. If we use just the GPU the calculation time is T/G . If we use both, the new time is $t=\max[T(1-\alpha)/G, T\alpha/C]+o$. The parameter α is a subject of optimization. Can we indeed argue that the optimal α is always 1? For reasonably small o and diverse computing tasks this is not the case.

The only argument supporting GPU-only computing is that the data transfer o is expensive. Indeed in the case of a simple Lennard-Jones fluid GPU-only simulation is sound. In this case memory transfer between host and device (the only pitfall of heterogeneous computation) is costly if done at each step. However, in the far more complex simulations of explicitly solvated biopolymers, there are components that are not suited for massive parallelization with single precision. In particular we illustrate in our first paper that particle meshed Ewald (PME) calculated in single precision increases the energy drift by a factor of about 10. Therefore, single precision PME should not be used in quantitative simulations. There are also terms such as torsions and improper torsions that run with similar efficiency on both the CPU and the GPU. The latter two require significant local resources (registers) and at the same time are relatively small in number reducing benefits from parallelism. Hence, these arguments and experimentations led to our choice of using heterogeneous computing in a laboratory setting to speed up molecular simulations.

An increase in the size of the integration time-step is another way to enhance the efficiency of the calculation. Enlarging the time step (while retaining numerical accuracy) can be achieved by removing fast degrees of freedom from the system. A widely used example is the SHAKE algorithm¹² which constrains bond lengths to their ideal values. Another approach is based on filtering high frequency modes¹³. Removal of bond vibrations, which are fast degrees of freedom, allows doubling the step of integration while retaining numerical precision. However, a combination of parallel architecture and the popular bond-relaxation SHAKE was proven hard to achieve. This is since bond-relaxation is inherently serial and requires many iterations to adjust coordinate positions. An alternative SHAKE formulation is based on a direct and iterative solution to a set of approximate linear equations for the Lagrange's multipliers. This alternative is easier to parallelize, and requires significantly fewer iterations of coordinate adjustments¹⁴⁻¹⁶. In the present manuscript we consider a novel implementation of the matrix formulation of SHAKE in which the entire algorithm is carried out on the GPU with some components in double precision. The Lagrange multipliers are found using conjugate gradient algorithm¹⁴ in single precision while the coordinate adjustments are conducted in double precision. This implementation brings more work to the GPU freeing the CPU to asynchronously conduct other computations. Our implementation on the GPU of parallel matrix SHAKE is more efficient than the use of CPU only in cases where the number of CPU cores is low compared to the system size. We use a typical laboratory tool of a CPU with four cores and a single GTX480. On a node with 12 cores and the same GPU, the constraints are better computed on the CPU. We therefore retain the option to run the same algorithm on multiple shared memory CPU's as defined by user input illustrating the flexibility of heterogeneous computing.

In addition to the improvement in SHAKE parallelization we also consider a special GPU routine (kernel) to handle non-bonded interactions between water molecules. We exploit the fact that most biological macromolecules simulated with an explicit solvent contain a large number of water molecules. Since all water - water interactions share Lennard Jones and electrostatic parameters we can significantly reduce memory access by keeping these

parameters in fast memory. Moreover, since we use neighbor lists, making a special case for water, one stores a water index rather than three atomic indices thereby reducing list size by up to three-fold. We recently reported a GPU implementation of the non-bonded force calculation using atomic neighbor lists⁵. Here we build on this code and present a specialized GPU implementation of water non-bonded force calculations.

These two additions led to speedup compared to our previous code (on the standard test case of DHFR) by about three-fold overall and four-fold in the GPU non-bonded force calculation. For each GPU implementation, an equivalent optimized function was implemented for execution on a multi-core shared memory CPU (using the parallel library Open Multi Processing (OpenMP) in case a GPU is not available).

Matrix SHAKE

In order to increase the time step while maintaining energy conservation we remove bond displacements that are fast degrees of freedom from the equations of motion. It is easy to parallelize the constraints of only the fastest bonds that include hydrogen atoms since the number of coupled bonds is always small and decoupled. Here we focus on the more general problem of parallelization of a network of bonds that is not factored to independent bond clusters. Examples are long peptide chains. Matrix SHAKE with conjugate gradient optimization to determine the Lagrange multipliers is an appropriate methodology to handle the above problem and was described in¹⁴. We briefly review it and discuss specific adjustments appropriate for the heterogeneous computing system of the present manuscript. Consider the algorithm of velocity Verlet¹⁷ with constraints (bold font denotes a vector):

$$\mathbf{x}(t+\tau)=\mathbf{x}_{i+1}=\mathbf{x}_i(t)+\mathbf{v}_i(t)\tau-1/2M^{-1}[\nabla U(\mathbf{x}_i)+\sum_a\eta_a\nabla\sigma_a(\mathbf{x}_i)]\tau^2 \quad (1.a)$$

$$\sigma_a=(\mathbf{x}_{a1}-\mathbf{x}_{a2})^2-r_{eq,a}^2 \quad (1.b)$$

Where τ is the time step, $\mathbf{x}, \mathbf{v} \in \mathbf{R}^{3N}$ are coordinate and velocity vectors of the whole system. \mathbf{x}_{a1} and \mathbf{x}_{a2} are the three dimensional coordinates of the two bonded atoms, and σ_a is the constraint on bond a , i.e. the distance square between the two points \mathbf{x}_{a1} and \mathbf{x}_{a2} is constrained to be equal to the square of the equilibrium distance $r_{eq,a}^2$. The Lagrange multipliers are η_a . M is a $3N \times 3N$ diagonal matrix where N is the number of atoms. The mass of a particle repeats three times in the matrix for the x, y , and z directions. For a corresponding equation for the velocities see¹⁴. We split (1.a) into a step without the application of the constraints; a Verlet step, which we call $\mathbf{x}^{(0)}$ and a correction when the constraints are taken into account

$$\mathbf{x}^{(0)}=\mathbf{x}_i(t)+\mathbf{v}(t)\tau-1/2M^{-1}\nabla U(\mathbf{x}_i(t))\tau^2 \quad (2.a)$$

$$\Delta\mathbf{x}_{constr}=1/2M^{-1}\sum_b\eta_b\nabla\sigma_b(\mathbf{x}_i)\tau^2 \quad (2.b)$$

$\Delta\mathbf{x}_{constr}$ is the adjustment to the coordinate enforced by the constraints such that $\mathbf{x}_{i+1}=\mathbf{x}^{(0)}-\Delta\mathbf{x}_{constr}$. The constraints at step $i+1$ are approximated by a first order Taylor expansion near the coordinates $\mathbf{x}^{(0)}$ which are now adjusted by $\Delta\mathbf{x}_{constr}$ to satisfy the constraint condition $\sigma(\mathbf{x}_{i+1})=0$:

$$\mathbf{0}=\sigma_a(\mathbf{x}_{i+1})\approx\sigma_a(\mathbf{x}^{(0)})-\nabla\sigma_a(\mathbf{x}^{(0)})\Delta\mathbf{x}_{constr} \quad \forall a \quad (3.a)$$

$$\sigma_a(\mathbf{x}^{(0)}) \approx \nabla \sigma_a(\mathbf{x}^{(0)}) \Delta \mathbf{x}_{\text{constr}} \quad (3.b)$$

We can write the right hand side of (3.b) $\nabla \sigma(\mathbf{x}^{(0)}) \Delta \mathbf{x}_{\text{constr}}$ in matrix form using (2.b) while factoring out the Lagrange multipliers and the time step:

$$\mathbf{A}_{ab} = [\nabla \sigma_a(\mathbf{x}^{(0)})]^t 1/2M^{-1} \nabla \sigma_b(\mathbf{x}_i) \quad (4)$$

This matrix is asymmetric because $\sigma_a(\mathbf{x}^{(0)})$ is not equal to $\sigma_a(\mathbf{x}_i)$. We can use the matrix \mathbf{A} to estimate the Lagrange multipliers $\lambda_b = \eta_b \tau^2$ of Eq. (2.b) using Eq. (3) and (4):

$$\mathbf{A}^{-1} \boldsymbol{\sigma} = \boldsymbol{\lambda} \quad (5)$$

The solution of this system is approximate because $\boldsymbol{\sigma}$ is not linear in $\boldsymbol{\lambda}$. The linear approximation made at (3.a) requires a Newton-like iteration until convergence. It converges provided that the initial point is sufficiently close to the right answer. Convergence is assumed when

$$|\sigma_a(\mathbf{x}_i^{(n)})| < \varepsilon \quad \forall a \quad (6)$$

where ε is a small number. The convergence is required for all constraints. For every iteration step we have

$$\mathbf{x}_{i+1}^{(n+1)} = \mathbf{x}_{i+1}^{(n)} - 1/2M^{-1} \sum_b \lambda_b^{(n)} \nabla \sigma_b(\mathbf{x}_i) \quad (7)$$

We define a corresponding matrix

$$\mathbf{A}_{ab}^{(n+1)} = \nabla \sigma_a(\mathbf{x}^{(n+1)}) 1/2M^{-1} \nabla \sigma_b(\mathbf{x}_i) \quad (8)$$

And we also have

$$\boldsymbol{\sigma}^{(n+1)} = \mathbf{A}^{(n+1)} \boldsymbol{\lambda}^{(n+1)} \quad (9)$$

It is easier to understand the computational aspects of the linear equation (Eq. (9)) for the constraints when the explicit expression for the bond constraints is used. Consider Eq. (1.b), taking a spatial derivative we have $\nabla \sigma_a = 2(\mathbf{x}_{a1} - \mathbf{x}_{a2})$ and a matrix element is therefore

$$\mathbf{A}_{ab}^{(n+1)} = 4(\mathbf{x}_{a1}^{(n+1)} - \mathbf{x}_{a2}^{(n+1)})^t 1/2M^{-1} (\mathbf{x}_{b1} - \mathbf{x}_{b2}) \quad (10)$$

Discussions on the properties of this matrix were given in the original paper of SHAKE^{12 and 14,18}. We consider three cases: (i) The constraints “a” and “b” are identical. In this case we have $\mathbf{A}_{aa}^{(n+1)} = 2(1/m_{a1} + 1/m_{a2})(\mathbf{x}_{a1}^{(n+1)} - \mathbf{x}_{a2}^{(n+1)})^t (\mathbf{x}_{a1} - \mathbf{x}_{a2})$. (ii) The constraints “a” and “b” are not the same but they share one atom (e.g. $a_2 = b_1$), we have $\mathbf{A}_{ab}^{(n+1)} = 2(\mathbf{x}_{a1}^{(n+1)} - \mathbf{x}_{a2}^{(n+1)})^t (1/m_{a2})(\mathbf{x}_{a2} - \mathbf{x}_{b2})$. (iii) No atoms are shared, $\mathbf{A}_{ab} = 0$.

From the explicit expression (Eq. (10)) and the cases discussed above it is clear that the matrix is sparse (bonds are coupled only if at least one atom is shared) and is recomputed every λ iteration. One alternative to the calculation of the matrix in the linear equation for the Lagrange multipliers is to approximate $A_{ab}^{(n+1)}$ by its diagonal, which is the bond relaxation algorithm⁶. The disadvantage of the bond-relaxation algorithm is that it requires many iterations; iterations that must be executed serially. Bond-relaxation is therefore not effective for parallelization and for GPU implementation. From now onwards we consider only options that employ non-diagonal matrices. Besides using the exact matrix we consider two approximate matrices (a) a symmetric matrix and (b) a constant matrix (Figure 1).

Consider first the symmetric matrix option, we write:

$$A_{ab}^s = \nabla \sigma_a(\mathbf{x}_i) \frac{1}{2M} \nabla \sigma_b(\mathbf{x}_i) = 4(\mathbf{x}_{a1} - \mathbf{x}_{a2})^t \frac{1}{2M} (\mathbf{x}_{b1} - \mathbf{x}_{b2}) \quad (11)$$

Since \mathbf{x}_{a1} is different from $\mathbf{x}_{a1}^{(n+1)}$ only by terms of order of τ^2 (the small time step), the difference between the exact and the symmetric matrix is small. The iterations of Eq. (9) converge with both matrices from Eq. (10) or Eq. (11) for sufficiently small τ . The number of iterations is however different. In practice a smaller number of iterations is required to converge the values of the Lagrange multipliers using the asymmetric matrix.

Another option is to use a constant matrix that does not depend on the coordinate at all. Hence a single matrix can be used throughout the simulations. The constant matrix was proposed in¹⁴ and an implementation was discussed in¹⁹. The basic idea is that the deviation of the constrained bond from ideality is expected to be small, so when constructing the symmetric matrix we have for non-zero elements

$$A_{aa}^s = 2(1/m_{a1} + 1/m_{a2}) r_{eq,a}^2 \quad A_{ab}^s = 2(1/m_{a2}) r_{eq,a} r_{eq,b} \cos(\theta_{ab}) \quad (12)$$

Where a and b bonds are sharing an atom ($a2=b1$), and θ is the angle between the two bonds. For the protocol to be exact the angle must be constrained as well. We can enforce it by adding an additional bond constraint between a1 and b2 atoms of equation (11)¹⁴. Alternatively we can pick the equilibrium, constant angle ($\theta_{ab,eq}$) and plug it into equation (12) to compute an approximate but time independent constraint matrix¹⁹. Since bond angles are usually restrained with a stiff harmonic potential their fluctuations are small suggesting that the approximation of a constant matrix is sound (even though it is no longer possible to argue that the solution is accurate with a particular order of τ). A significant advantage is that the matrix is fixed and requires computation only at the beginning of the calculations. The disadvantage (similar to the discussion about the symmetric matrix) is that the number of iterations to convergence that must be conducted serially may be significantly larger and more costly than the number of iterations for matrices that are updated during the calculation. Below we examine the options of symmetric and constant matrices.

Another complication that we need to consider is that we need to find the inverse of \mathbf{A} to determine the Lagrange multipliers (Eq. (9)). Calculation of an inverse is usually significantly more expensive than the construction of the matrix itself. How to estimate the Lagrange multipliers with explicit calculation of the inverse therefore attracted some attention. In LINCS²¹ an approximate inverse is computed. Here we follow the algorithm outlined in¹⁴ in which a conjugate gradient algorithm (CG) estimates the Lagrange multipliers. For a symmetric and non-negative matrix \mathbf{A} we seek the minimum of a target function T as a function of the Lagrange multiplier

$$T = \sigma^t \lambda + (1/2) \lambda^t \mathbf{A} \lambda \quad (13)$$

Note that since we are restricted to symmetric and non-negative matrices we cannot use the exact matrix Eq. (13) “as is” and we therefore use the symmetric versions. The symmetric matrix is indeed non-negative as was shown in ¹⁴. Differentiating the target function T with respect to λ recovers Eq. (9). So solving (13) is equivalent to solving (9). Formally, CG is guaranteed to find the global minimum in L steps where L is the number of constraints. In practice the number of iterations to determine the Lagrange multipliers is below 10. As is illustrated in the Results section in our most efficient implementation we fixed the number of iterations to seven. It provides good accuracy and at the same time avoids a convergence check and an “if” statement. Since the final decision about convergence is made by coordinate test, high accuracy of the Lagrange’s multipliers (which are calculated here in single precision) is less important.

We add one more twist to the straightforward application of CG using a pre-conditioner. The CG is solved trivially if \mathbf{A} is diagonal or close to identity. A good pre-conditioner will bring the matrix closer to the identity operator. A common practical choice for a pre-conditioner is the following. Consider the matrix \mathbf{D} that includes only the diagonal elements of \mathbf{A} ($\mathbf{D} = \text{diagonal}(\mathbf{A})$). We use \mathbf{D} to bring the matrix \mathbf{A} closer to identity and to reduce the number of iterations required.

Finally it is important to emphasize that we have two loops of iterations. We have iterations to solve the Lagrange multipliers and iterations to adjust the coordinate vectors. The latter must be done in double precision or bond energy will not converge to less than 1×10^{-7} or so. We offer the option to run the entire algorithm on the GPU (determine the Lagrange multipliers and adjust the coordinates in double precision) without memory transfer before convergence.

Finally we comment on the implementation on the GPU. The linear solver must be programmed from scratch manually, which is an additional challenge for non-symmetric matrices. As we argue below, linear solvers for the GPU, based on function libraries, are not effective.

The preconditioned conjugate gradient SHAKE algorithm

The resulting SHAKE Algorithm for constrained Molecular Dynamics ¹⁴ on the CPU or the GPU is below. Note the meaning of the three indices: time step i , SHAKE step n and conjugate gradient step k . Step 4 is the CG algorithm where \mathbf{r} (residual), \mathbf{b} (direction of search step) and \mathbf{x} are utility vectors. The scalars α (step magnitude) and γ are coefficients required in the parallel version of the conjugate gradient algorithm (PCG). In addition, `max_shk` limits the maximum number of SHAKE iterations, `max_CG` limits the maximum number of CG steps, ϵ is the maximum allowed constraint error (maximum of the error found in all bonds) and `CG_tolerance` is the conjugate gradient convergence tolerance for the determination of the Lagrange multipliers. The dashed lines indicate barriers. A barrier is a point in the code that enforces any threads that reached that point to wait. This is until all (other) threads have completed their tasks and reach the same point. Barriers are special breaks during parallel execution in which all the threads are forced to synchronize and are typically used at a time in which data sharing between the threads is required.

Algorithm: At time step i :

- 1) Compute the matrix non-zero elements A_{ab}
- 2) do while ($|\sigma_a(\mathbf{x}_i^{(n)})| > \epsilon$) and ($n < \text{max_shk}$)
3. Evaluate constraints. If $|\sigma_a(\mathbf{x}_i^{(n)})| \leq \epsilon \square_a$, stop.

```

4. Solve the linear system  $\Sigma_b A_{ab} \lambda_b = \sigma_a(\mathbf{x}_1^{(n)})$  to find  $\lambda^{(n)}$  by PCG:
 $\mathbf{r}_1 = \sigma - \mathbf{A}\lambda$ 
 $\mathbf{b}_1 = \mathbf{D}^{-1}\mathbf{r}_1$ 
 $\mathbf{x}_1 = \mathbf{b}_1$ 
- - - - (end of PCG segment 1)
do while (  $|\mathbf{r}_k| > \text{PCG\_tolerance}$  ) and (  $k < \text{max\_CG}$  )
 $\alpha = \mathbf{r}_k^T \mathbf{x}_k / \mathbf{b}^T \mathbf{A} \mathbf{b}$ 
- - - - (end of PCG segment 2) delete;  $\mathbf{r}_{k+1} = \mathbf{r}_k - \alpha \mathbf{A} \mathbf{b}_k$ 
 $\lambda_{k+1} = \lambda_k + \alpha \mathbf{b}_k$ 
 $\mathbf{x}_{k+1} = \mathbf{D}^{-1} \mathbf{r}_{k+1}$ 
 $\gamma = \mathbf{r}^T$ 
 $\mathbf{x}_{k+1}$ 
 $\mathbf{x}_{k+1} / \mathbf{r}$ 
 $\mathbf{x}_k$ 
 $\mathbf{x}_k$ 
- - - - (end of PCG segment 3)
 $\mathbf{b}_{k+1}$ 
 $= \mathbf{x}_{k+1} + \gamma \mathbf{b}_k$ 
 $\mathbf{b}_k$ 
- - - -
end do
5. Compute an adjusted coordinate set:
 $x_{(n+1)}^i = x_{(n)}^i - (1/2)M^{-1} \Sigma \lambda_{(n)} \beta$ 
 $x_{(n)}^i = \sigma_{(n)} \beta^i (\mathbf{x}_{(n)}^i)$ 
6.  $n = n + 1$ 
end do

```

Implementation of parallel conjugate gradient SHAKE

We consider a symmetric or a constant \mathbf{A} matrix. The PCG algorithm consists of matrix vector multiplications, vector dot products and the product of a pre-conditioner matrix with a vector. In our case, the pre-conditioner matrix is diagonal and is stored as a linear array. For efficient execution we exploit the sparsity of the matrix and minimize memory access. We store the sparse matrix in compressed row storage format^{22 14}. This is the same format used by the cusparse NVIDIA GPU linear algebra library [NVIDIA CUDA user manual]. This storage is well suited for a parallel matrix vector multiplication algorithm²³. The nonzero columns of each row are stored adjacent (row major order) in memory in the array **elem**. The length of **elem** is the number of non-zero matrix elements. The column number of each entry in **elem** is stored in the integer **col** array of the same length as **elem**. Similarly, the row number of each entry is stored in the integer **row** array of the same length as **elem**. The **cmprow** array is of length of the number of rows plus one. The elements of **cmprow** (except the very last) are the first non-zero column of \mathbf{A} for each row (as found in **elem**). The very last element in the array stores the would-be-first column of a non-existing additional row. This allows the creation of a loop such as shown in the matrix multiplication algorithm below to iterate over all rows without exceptions for the last row. Note that the indices follow C convention in which the

first element is **elem**(0). An example for the compressed representation of a symmetric matrix follows below:

Compressed row storage example:

$$\begin{array}{r|cccc|}
 & 3 & 2 & 0 & 0 & \\
 A = & 2 & 1 & 5 & 0 & \\
 & 0 & 5 & 6 & 0 & \\
 & 0 & 0 & 0 & 4 &
 \end{array}
 \quad
 \begin{array}{l}
 \mathbf{elem}=\{3, 2, 2, 1, 5, 5, 6, 4\} \\
 \mathbf{col}=\{0, 1, 0, 1, 2, 1, 2, 3\} \\
 \mathbf{row}=\{0, 0, 1, 1, 1, 2, 2, 3\} \\
 \mathbf{cmprow}=\{0, 2, 5, 7, 8\}
 \end{array}$$

The sparse matrix multiplication algorithm $\mathbf{b} = \mathbf{A}\mathbf{x}$ is using the same arrays as described above. \mathbf{b} is the result vector. All arrays are global variables. The integers `row_start` and `row_end` are the start and end row numbers assigned to the thread. In the GPU implementation each thread loops over one row and the outer loop disappears. In this case `row j` equals the thread number.

```

for (int j=row_start; j<row_end; j++) {
double elm = 0.0;
for (int i= cmprow [j]; i< cmprow [j+1]; i++) {
int c = col [i];
double Aij = elem [i];
elm += Aij * x [c];
//---- if use symmetry, we must add the following here in the inner loop:
if (c != row) b [c] += Aij * x [j]; //-- scattered write
}
b [j] = elm;
}

```

We do not take advantage of the symmetry of the \mathbf{A} matrix because it requires scattered memory access and, in parallel, necessitates the use of a reduction. Threads need to write to elements in the result vector other than the rows they are assigned to in the load balance phase. The latter also requires local (thread private) copies of the result vector.

Parallelization—Steps 1 and 5 of the algorithm are easily parallelized. In step 1, when re-computing the matrix we take advantage of the fact that the identity of the non-zero elements is fixed and that it was already stored in efficient sparse storage format. We consider only nonzero matrix elements that are divided precisely between open-MP threads on CPU and get the bond indices from the non-compressed column and row arrays. This can be done in parallel by dividing the number of matrix elements over the number of threads. Step 3 is divided between open-MP threads on the CPU or one constraint per thread on the GPU. On the CPU, each thread stores its own maximum σ error determined in a global array of length of the number of threads. Subsequently, on the CPU, all threads loop over this array and determine max error individually. On the GPU, a shared memory parallel reduction is done inside the function without global memory access. The GPU reduction is explained in more detail below. So on the CPU, the final reduction is not done in parallel as it is on the GPU. On the CPU, each thread performs its own final sum. In step 5 of the SHAKE algorithm (and equation 7) we sum up the adjustments of the coordinates. It is not possible to adjust the coordinates by looping independently over the constraints since it is possible that multiple constraints share an atom. Therefore we loop instead over all atoms involved with one or more constraints and divide these between the threads. On the GPU this is one atom per thread. To work simultaneously in the constraints and the atom representation we keep 3 additional arrays, one array storing each atom number involved in any constraint and one array with pointers into a third array with

the constraint numbers for that atom. The coordinate adjustments are in a nested loop, the outer loop over the atoms and inner loop (the sum of equation 7) over the constraints. This prevents the write conflicts between threads.

Parallel Preconditioned Conjugate Gradient—Dot products, while simple to calculate in series, are more challenging to do in parallel since communication is required to gather the results from all threads. To begin with each thread calculates its share of elements of the dot product and accumulates the partial sum. To perform the global sum all threads must complete their task and a barrier for synchronization is required. In case of a dot product on several CPU threads, each thread must complete its own partial sum and write the result into a global array of length of the number of threads before a global sum can be attempted. Then the final sum can be made. This final sum can also be done in parallel as we do on the GPU (see parallel reduction below) or as we do on the CPU; each thread can perform this final sum on its own whenever it is ready to use the data. The last option is not efficient on thousands of threads but is faster with a limited number of threads (say up to 128) because it does not require a second barrier. Therefore we implement the parallel reduction only in the GPU code (see parallel reduction below).

The main obstacles to efficient parallelization are the thread synchronizations required to calculate the coefficients α and γ and update of the vector \mathbf{b} . We therefore combined segments of the PCG algorithm (see algorithm above, dashed lines) in such a way that only three barriers remain and at the same time memory (array) access is minimized. In general, the PCG code can be combined into 3 main segments (Supporting Information Appendix A) that can be calculated without intermediate barriers in massive parallel calculation at a level of one thread per row of the matrix \mathbf{A} . To reduce (array) memory access, for example in PCG segment 1, we declare a local utility variable q :

```
double q = Sigma[n] - Ax[n] (see appendix A);
```

This variable q is then used 4 times in the product with D^{-1} , the assignment of vectors and the accumulating dot product. So this prevented us from accessing arrays **Sigma** and **Ax** again. Similarly, memory access is minimized in the other PCG segments (Supporting Information Appendix A). This process works the same way on the GPU.

GPU Implementation—The entire algorithm is ported to the GPU in the GPU implementation. The outer Newton-like iterations need to be evaluated in double precision. In particular, the error vector determination, and coordinate adjustment need to be done in double precision or the error will not converge to less than $\sim 10^{-10}$ as measured by relative error in bond lengths. The conjugate gradient itself is divided into the same segments as discussed above and all steps are computed entirely on the GPU. There is no requirement for memory transfer between host and device during the determination of the Lagrange multipliers calculation nor during coordinate corrections resulting in significant speedup.

GPU Parallel Conjugate Gradient (PCG)

The main difference between the CPU and GPU code is the granularity of parallelization. The CPU code is still serial in nature. On the GPU, each row of the vector \mathbf{x} (and the matrix \mathbf{A}) in $\mathbf{Ax} = \mathbf{b}$ is represented by its own single thread. This completely eliminates all outer loops. In the matrix multiplication code segment shown above, the loop over row j is eliminated and row j is the thread number. We notice right away now there are no partial sums at all. Each thread only has the single value calculated in that thread. We take as an example $\mathbf{x}^T \mathbf{x}$. We could take the same approach as in the serial implementation. We place a barrier that enforces a wait

until all threads have completed their single product $x_j * x_j$. And then we could have each thread sum all rows. This would give the correct result but with thousands of rows, very inefficiently. We therefore implement a shared memory parallel reduction algorithm at the end of each PCG segment, inside the kernel that is logarithmic in the number of array elements (Supporting Information Appendix B).

A quick review of CUDA instructions necessary for the text below follows. Threads are divided into blocks. Within a block threads can communicate through shared memory. Threads in a different block cannot access shared memory in another block. Individual threads also have local registers in which they can read/writer scalars. Only the owner thread can access local registers. Global memory can be accessed by all threads so it can also be used for communication between threads but is slow. Continuing with our example, we realize that each thread in a particular block has calculated its $x_j * x_j$ product; it can store this in fast and **shared** memory shared that can be seen by other threads in the same block. The shared memory array is the length of the fixed block-size which is a factor of 32 (we use 128). The parallel shared memory reduction is conducted as follows. Each thread stores its value in shared:

```
shared[j] = x_j * x_j;
```

Reduction is first conducted within the thread block (see Figure 1). We start with 128 values in separate threads.

```
if (j < blocksize/2) shared[j] += shared[j + blocksize/2];
```

With this single statement, we have reduced our block data by half. The reduction, which is done in parallel, is exponentially fast:

```
for (int n=6; n<=0; n--)
  if (j < 2n) shared[j] += shared[j + 2n];
```

For a block size of 128, 7 statements are needed to compute the sum of all elements of the block in the first position of the array **shared**. At this point we need to communicate between blocks. We have up to 128 blocks that each have the sum of the dot product in the first thread of that block. We perform the last sum using the slow global memory. We define a global array **global** of length blocksize.

```
If (thread == 0) global[block ID] = shared[0].
```

At this point we have a global array with up to 128 values and one more reduction of one final block is required. Therefore at this point we need a global thread barrier, which is done as follows (see also Appendix B in Supporting Information). A counter is kept for each block during the global writes. When this counter reaches the number of total blocks, each has written its value and the last block of threads to write is the one that will perform the final reduction. This last block will now load the global data back into shared:

```
shared[j] = global[j];
```

At this point the reduction is just like what is shown above, except that only one block is used. Repeat for $n=6$ to 0

```
if (j < blocksize/2n) shared[j] += shared[j + 2n];
```

When done:

```
global[0] = shared[0].
```

The final single scalar result is now in **global[0]**. All other threads were waiting for this result; the waiting is enforced by the global barrier. Thread zero of this last block will release this global thread barrier.

A reduction, (summation of array elements that are spread over a large number of threads) can be costly if communication is not handled efficiently. In the context of GPU programming, such reduction can be done in parallel if we stay within the boundary of a block – a subgroup of threads that share memory. For a block the reduction requires only a logarithmic number of serial steps as a function of the length of the array. This is illustrated in Fig. 1 for a block of threads on the GPU that share the same memory. A global barrier (for all blocks) is not provided currently in CUDA (only block-wise barriers are provided). The global barrier implemented between the block reductions and the final block reduction allowed calculation of an entire PCG step in one GPU kernel. The alternative is a separate kernel for each PCG segment. Calling separate kernels from the CPU introduces a global barrier anyway; however the last barrier is a bit slower. Local storage of Conjugate Gradient intermediate results minimizes array memory access altogether just as in the CPU code (see example, $\text{double } q = \text{Sigma}[n] - \text{Ax}[n]$ above and details in Supporting Information, Appendix B). We use a fixed block size of 128 allowing a maximum of 16384 constraints. The fixed block size also allows complete unrolling of the reduction loop. The block-size can be doubled to 256 (increasing the maximum to 65536 constraints) or increased further as needed.

Water-specific force kernels on the GPU

In molecular simulations with explicit solvent the number of water interactions is significant and a focus on optimizing this part, even with a specific code, seems like a reasonable approach. Indeed, the program MOIL⁶ has used water-specific code for a long time. Here we discuss the development of a new water code for the GPU that led to a dramatic speed up and was ported back to the CPU with a similar effect.

We start with the generation of atomic neighbor lists presented recently in our initial GPU implementation. Atomic neighbor lists are widely implemented in CPU based code. However, GPU memory limitations motivated a number of studies to seek different solutions^{24–26}. In our previous paper we show that atomic lists can be used accurately and efficiently on the GPU. The main simulation box is divided into rectangular cells or grid boxes, equal in size. Each atom is placed in a grid box according to its atomic coordinates. Since the neighboring boxes are predetermined, atomic neighbors are considered according to their own and neighboring boxes and are finally added to a list according to a cutoff distance criterion.

Force calculation can be done directly at the level of the grid boxes avoiding the use of lists. Why use atomic neighbor lists? The advantages of the atomic list approach are as follows. 1) Exclusion of non-bonded-interactions (of bonded atoms) is made at the list generation step, so we avoid exclusion tests during the force calculation. 2) Some parameter calculation can be done and stored during the list calculations (e.g. charge-charge multiplication) in a coalesced

way. This does not only save a set of floating point calculations in the force kernel but also prevents the need to use local registers of the threads to store parameters for the atoms. It also prevents the need for shared memory storage of its neighbor atom parameters. 3) Since the cutoff distance is spherical fewer interactions need to be calculated compared to direct box calculations.

A disadvantage of the neighbor list is that neighbor coordinates cannot be stored in shared memory (each atom has a unique list and the lists are too large) but this is compensated for by the significantly reduced number of interactions and frees the otherwise required shared memory for other uses. In particular, on Fermi based GPUs, a portion of this freed shared memory can be used as L1 cache. This cache memory significantly increases coordinate access speed. The remaining portion of the freed shared memory stores a lookup table for a quadratic interpolation of the electrostatic force. We also keep the option of storing larger tables in texture memory but we find the shared memory option significantly faster (almost 2 fold) even if the texture memory table access is completely cached. A third option is a medium-sized (up to 2000 knots or so) lookup table in constant memory; this option is still slower than the shared memory approach.

To conclude, the use of a neighbor list makes the resultant force kernel simpler. With the list it has more resources available for other types of calculations. The contribution of the present manuscript is the division of input particles into non-water atoms and water molecules resulting in multiple and specific force kernels, which are significantly faster. Our most complex force kernel (water-water interactions) allows the calculation of all 9 terms of a water-water interaction while remaining around the maximum limit of 32 registers. Satisfying the limit of 32 registers allows for up to 66% occupancy on Fermi GPU's. For a water-water interaction, each thread represents one water molecule. It requires 18 local registers to store 9 coordinates + 9 force accumulation registers. After separating the water molecules from the rest of the atoms, each grid box has two lists: a list of non-water atoms and a list of water particles. The lists are generated in the same way as described previously⁵. We do not take advantage of the symmetry in the neighbor list in the GPU implementation (i.e. we compute both interactions, particle i and j and particles j and i). Memory access considerations motivate the doubling of computations (also see discussion by Stone et al.²⁵). This means that we now have four respective kernels for neighbor list generation and four kernels for non-bonded force calculation. Each force kernel accesses its own neighbor list. The interaction types are: (1) a non-water atom interacting with another non-water atom. In this case each interacting pair can comprise 2 different atoms and the geometric averages of the pair Lennard Jones and product of charge parameters are stored in a single float4 of the neighbor list as previously described.¹⁶ (2) A water molecule interacting with a non-water atom. In this case we again still use a float4 neighbor list matrix but store the non-water atom parameters rather than the geometric pair averages stored in the previous case. We do not store products such as done in case (1) because additional memory would be needed. This is since there are two interaction types (non-water with water oxygen and water hydrogen). We almost always choose to do an extra calculation over global memory access since the performance is typically bound by memory access. This is true for both CPU and GPU. (3) Non-water atom with water particle. In this case, only the integer number of the neighbor water molecule is stored in the respective neighbor matrix in order to obtain its coordinates in the force kernel. Note that this function is not needed in the CPU version of the code since on the CPU we do take advantage of the symmetry and calculations of (3) are conducted in case (2). (4) Water interacting with another water molecule: We store the single integer of the number of neighbor water in the neighbor matrix.

All of the above mentioned functions are also implemented with Open Multi-Processing (OMP) for parallel execution on a shared memory system with multiple CPU threads in absence of a

GPU. The simplicity and performance of the GPU functions prompted us to port this code back to the CPU for multi-core execution in the absence of a GPU. On the CPU, we do take advantage of symmetry in the neighbor matrix so no pair interaction is calculated twice. The CPU has a limited number of threads and it makes sense to have a copy of the force vectors for each thread. This allows the scattered writes to neighbor atom force that we avoid on the GPU. Therefore on the CPU a reduction over all thread results is required after the completion of the non-bonded force calculation on all threads. This reduction is parallelized in a trivial way: each thread adds n vectors of force of its subset of atoms where n is the number of threads. On the GPU one thread represents either one atom in cases (1) and (3) above or one water molecule in cases (2) and (4) above while on the CPU one thread has an outer loop iterating over multiple non-water atoms or water molecules. Neighbor matrices are of the same types, and they can be generated on the CPU or on the GPU. Importantly however, for coalesced access, the CPU matrix versions are the transpose of the GPU version. In massive parallel computing, the 1st row of the neighbor matrix is the first neighbor of all atoms or water molecules. Switching the index generates the other type: for CPU serial execution, the first row contains all the neighbor atoms of the 1st atom or water. More discussion on this topic can be found in ^{4,7}.

Since all water molecules have the same composition and atomic parameters it makes sense to store these parameters in constant or local memory rather than include them as interacting pair information in the neighbor lists. We pre-calculate the 5 water hydrogen-hydrogen, oxygen-oxygen and hydrogen-oxygen pair interaction for charge ($q_i q_j$) and Lennard Jones ($A_i A_j$ and $B_i B_j$) products (see #define, appendix C). This significantly reduces memory access and storage requirements of the neighbor matrix in cases (3) and (4) above since a neighboring water is represented simply by one integer. In addition, since one GPU thread now represents a water molecule in two of the interaction types discussed above (cases (2) & (4)) one set of neighbor coordinates read can be used for the calculation of 3 interactions: one with each of the 3 atoms of the water belonging to the thread. This means 3 fold less coordinate access in these kernels compared to the non-water atom approach used exclusively before.

The NB force kernel in more detail (see also Supporting Information)—The non-bonded force kernel for non-water interactions is essentially identical to the one previously reported ⁵. The new water-water kernel is shown in its entirety in supporting information. It is an interesting case to illustrate the tradeoff of the computing versus memory access. The efficiency is dominated by memory accesses and less by computing, especially for the GPU. In the procedure described below we therefore emphasize reduced memory access at the expense of more calculation. This leads to a significantly more efficient code.

In the current implementation of the water-water kernel one thread represents one water molecule. The coordinates and forces of one water molecule are loaded into 18 local registers. Quadratic force lookup coefficients are read from texture memory and stored in shared memory ⁵. The neighboring water molecules are read from the global memory matrix in a coalesced way in which the next thread accesses the exact adjacent memory location of the neighbor list resulting in substantial performance gain. For a more extensive discussion on coalesced memory access with atomic neighbor matrix molecular dynamics see ^{5,27}. The neighboring water is identified only by its index since all the energy parameters are stored in registers. Information that we miss is the grid-box position of the neighboring water. This information is useful since it identifies neighbors that are related by translational symmetry. Nevertheless to avoid memory accesses we found it beneficial to re-examine the Cartesian displacements between the water oxygens instead. The translation only adds a logical test, which is not very expensive, leading to an overall computational gain. Once the single integer number denoting a water neighbor is obtained from the water-water neighbor matrix the coordinates of the neighboring oxygen are fetched in a single float4 read, from global (L1 cache) or optionally from texture memory as discussed in ⁵. The difference vector with respect

to the oxygen of the prime water molecule is computed. The elements of the difference vector are (rx, ry, rz) . If any of these elements (say ri) exceeds half box length the neighbor must result from a translational symmetry operation. The displacement is shifted to make the length of all the vector elements smaller than half the box size. The extra cost of the shift was found to be less significant than memory access (if we attempt to store and transmit the grid-box identity and/or the difference vector).

After accessing the neighbor oxygen coordinate the O-O interaction force can be calculated, as well as the H1-O and H2-O interaction forces. This is before further neighbor coordinates are read from global memory. Each set of coordinates (one of the three neighbor water atoms) allows for three pair calculations with each of the atoms of the water belonging to the thread (see Supporting Information, Appendix C, underlined comment for neighbor coordinate read). During these three calculations, only shared memory is accessed in the force lookup. The nine calculations have the effect of partial unrolling of the neighbor loop. In loop unrolling the statements inside a loop are written explicitly and thereby reduce the loop iterations. Loop iterations carry some overhead.

Results

Systems

The systems investigated and benchmarked in this study varied in their size (from 2690 atoms to 99,905) with consistent improvements in performance as outlined in this section. Important simulation parameters are given in Table 1 below:

Performance

The new matrix SHAKE implementation for satisfying the macro-molecule bond constraints is faster than the serial 'bond relaxation' approach even on a single core. Timing results for SHAKE on a single core are listed in Table 2 where the last row is the bond relaxation method. Scaling seems to improve with system size.

We comment that system sizes considered are small to overcome the parallelization overhead. A factor that leads to less than optimal scaling is the relatively low load per thread. Even in our largest linked system with 7691 non-separable bonds we only have 641 bonds per core on 12 cores. This is a relatively small amount of work to justify significant communication. On the GPU, at one constraint or atom per thread, this translates to low occupancy (only a portion of the card is being used). Thus, while we see significant speed up in the SHAKE calculations while adding more cores the speed-up is slower than linear. Hence the parallelization of a globally linked system is significantly more complex than parallelization of decoupled bonded molecules (e.g. water molecules). Decoupled systems can be divided trivially between cores, scale perfectly, and are not included in the simulation below.

An interesting feature of the combination of the heterogeneous environment of multiple CPU cores and a GPU is the possibility to conduct asynchronous calculation. For example the GPU may work on conducting the SHAKE algorithm for the coupled system of bonds of the macromolecule(s) while at the same time the CPU's may work on adjusting water coordinates. A possibility of partitioning the workload is given in Table 3. Using the GPU and exploiting asynchronous calculations as mentioned earlier are providing additional speed advantage. On the larger systems, the GPU constrained calculations are done at about the same speed as 5 CPU cores but at the same time free CPU cores for potentially picking up other tasks such as in this case water SHAKE.

More detailed benchmarks of SHAKE options and the use a fixed matrix, versus a matrix that is updated every time step is provided in Table 2. In all cases reported below convergence was

assumed when the relative errors of bond lengths were smaller than 10^{-10} and the corresponding errors for velocities did not exceed 10^{-8} . Significantly less accurate convergence criteria did not conserve energy to the threshold we desired (less than one percent change in energy for a microsecond simulation of DHFR in the NVE ensemble). Significant experimentation with the SHAKE parameters led to the following conclusions:

1. The use of a constraint matrix A which is updated every time step led to execution about ~30% faster than the use of a constant matrix which is pre-computed in the beginning of the calculations. The constant matrix is of course cheaper to compute but it requires a larger number of SHAKE iterations.
2. Conjugated gradient (CG) iterations on the GPU, to determine the Lagrange multipliers, are conducted in single precision. In principle CG iterations can continue until high precision Lagrange multipliers are obtained. This however is not necessary and less efficient. The final judge of convergence is the accuracy of the bond lengths (in double precision) so the algorithm is still successful if the (approximate) Lagrange multipliers are not the most accurate but are successful in generating accurate bonds.
3. The CG calculations are conducted with a fixed number of iterations (and not necessarily until exact convergence of the Lagrange multipliers is achieved). This choice was made following efficiency considerations. A fixed number of iterations converge faster to the correct bond length than CG calculations that accurately determine the Lagrange multipliers. We note that testing for convergence is inefficient, even on the CPU. Instead we conduct a fixed number of CG steps (typically seven, unrolled) that provide a sound next-approximation for the coordinate adjustments.
4. In principle velocity SHAKE requires no iterations and should be satisfied with a single round of Lagrange multipliers. However, our calculations are using approximate, single precision Lagrange multipliers and therefore sometimes iterations are required (until the above convergence criterion is satisfied).

As we discussed in our first paper on GPU/CPU implementation of MOIL⁵ our code conserves well the energy of the system and therefore provides adequate sampling of configurations and trajectories from the microcanonical ensemble. We illustrate the energy conservation in table 4 and figure 4. In principle the SHAKE algorithm allows the use of a larger time step than before. In Table 4 we report energy conservation for the full SHAKE implementation as well as for the option of SHAKE-L, which means that only bonds that include light atoms (like hydrogen atoms) are included in the set of constrained bonds. SHAKING all bonds in the DHFR system is showing an improved energy conservation compared to SHAKE-L. The example below is for DHFR

We consider next the enhancement in performance that we obtained in the calculation of real space non-bonded interactions due to the introduction of the water kernel. In Figure 5 we compare the time of calculation using the new system in milliseconds on a system that contains a GPU card (GTX480), parallel calculations on all four cores of the shared-memory system, and calculations on a single core. We made a choice of computational set-ups that can be purchased cheaply and can be found in a usual laboratory setting. The speed up of the GPU system compared to a typical four core machine is roughly a factor of ten for DHFR and DOPC and about a factor of twelve for 1IHF and myosin.

It is also useful to examine directly the benefit of the water kernel on the calculation of the non-bonded force on the GPU with a different measure. We consider the time (in nanoseconds) that is required per atom and the comparison is provided in Fig. 6 for the overall real-space non-bonded interaction using (black square) or not using (gray diamond) special water kernels.

The speed up is consistently more than a factor of three and up to just over a factor of four. This is with exception of the two smallest systems for which overhead associated with the multiple kernels plays a role.

Finally we provide an overall view of our code performance. We consider the same seven benchmark systems and compare a 4 core computer to a 4 core computer + GTX480. We measure the performance in nanoseconds per day emphasizing the optimization of different components in Figure 7. The best possible speedups are obtained for relatively small systems and the speedup is reduced for larger systems. The reason is the lack of successful optimization of the reciprocal summation of the non-bonded interactions (PME). In our earlier paper ⁵ we noted that to conserve energy the PME calculation is better conducted in double precision and we therefore keep it on the CPU. While we parallelized the PME code on 4 cores the performance is still lagging behind other speedup of the code. This optimization remains a topic for future work.

Figure 8 graphically illustrates the increased weight of the PME code and the necessity for future optimization to address this component or find an alternative to the calculation of long-range electrostatics. The figure is arranged such that the smallest system is at the center and the largest at the upper layer of the circle. It is evident that the yellow component of the calculation dominates for the largest system and is significantly less important for the smallest system. Future direction may consider parallelization on more cores, or doing parts of the PME calculations in single precision (like we are doing for the SHAKE algorithm, the Lagrange multipliers are computed in single precision and the coordinates in double).

Conclusions

In the present manuscript we revisited a Molecular Dynamics code (MOIL) that was ported to the heterogeneous environment of the CPU and GPU and further enhance the efficiency and versatility of the program. We have shown that a special kernel for water molecules speeds up the GPU calculations of real-space non-bonded forces by a factor of four for systems larger than 20,000 atoms. We have also illustrated that the usually difficult-to-parallelize SHAKE algorithm can run efficiently when the matrix form is used both on multiple CPU's or entirely on the GPU. The speedup is evident even on a single core compared to the bond relaxation approach. The GPU SHAKE implementation performance is similar to a calculation conducted on 5 cores. In our first CPU/GPU implementation of MD reported in ⁵ we emphasized the use of double precision coordinates and velocities to ensure energy conservation and correct dynamics in the microcanonical ensemble. Therefore we split the SHAKE algorithm such that the coordinate adjustments are done in double precision on both the GPU or CPU. We were further able to take advantage of asynchronous calculations in which the SHAKE of water (that is done in separate routines in matrix form) is conducted simultaneously with the GPU SHAKE of the macromolecular coupled bond system. We anticipate that the gain from the new version of SHAKE will be more significant as the number of coupled bonds increases. Future directions may include the study of more complex network of bonds (e.g. the inclusion of bond angle as Urey Bradely terms) and the expansion of special kernel to other molecular fragments.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This research was supported by NIH grant GM05976, NSF grant CCF-0833162, and by Welch grant F-1783. Helpful discussions and the advice of Professor Martin Burtcher are gratefully acknowledged.

References

1. Hess B, Kutzner C, van der Spoel D, Lindahl E. *J Chem Theory Comput.* 2008; 4:435–447.
2. Phillips JC, Braun R, Wang W, Gumbart J, Tajkhorshid E, Villa E, Chipot C, Skeel RD, Kale L, Schulten K. *J Comput Chem.* 2005; 26:1781–1802. [PubMed: 16222654]
3. Harvey MJ, Giupponi G, De Fabritiis G. *J Chem Theory Comput.* 2009; 5:1632–1639.
4. Shaw DE, Deneroff MM, Dror RO, Kuskin JS, Larson RH, Salmon JK, Young C, Batson B, Bowers KJ, Chao JC, Eastwood MP, Gagliardo J, Grossman JP, Ho CR, Ierardi DJ, Kolossvary I, Klepeis JL, Layman T, McLeavey C, Moraes MA, Mueller R, Priest EC, Shan YB, Spengler J, Theobald M, Towles B, Wang SC. *Commun ACM.* 2008; 51:91–97.
5. Ruymgaart AP, Cardenas AE, Elber R. *J Chem Theory Comput.* 2011; 7:3072–3082. [PubMed: 22328867]
6. Elber R, Roitberg A, Simmerling C, Goldstein R, Li HY, Verkhivker G, Keasar C, Zhang J, Ulitsky A. *Comput Phys Commun.* 1995; 91:159–189.
7. Kirmizialtin S, Nguyen V, Johnson KA, Elber R. *Structure.* 2012; 20:618–627. [PubMed: 22483109]
8. Kirmizialtin S, Elber R. *J Phys Chem A.* 2011; 115:6137–6148. [PubMed: 21500798]
9. Majek P, Elber R. *J Chem Theory Comput.* 2010; 6:1805–1817. [PubMed: 20596240]
10. West AMA, Elber R, Shalloway D. *Journal of Chemical Physics.* 2007; 126:145104. [PubMed: 17444753]
11. Faradjian AK, Elber R. *J Chem Phys.* 2004; 120:10880–10889. [PubMed: 15268118]
12. Ryckaert JP, Ciccotti G, Berendsen HJC. *Journal of Computational Physics.* 1977; 23:327–341.
13. Olender R, Elber R. *J Chem Phys.* 1996; 105:9299–9315.
14. Weinbach Y, Elber R. *J Comput Phys.* 2005; 209:193–206.
15. Hess B. *J Chem Theory Comput.* 2008; 4:116–122.
16. Elber R, Ruymgaart AP, Hess B. *Euro Phys J - Special Topics.* 2011; 200:211–223.
17. Verlet L. *Phys Rev.* 1967; 159:98–103.
18. Barth E, Kuczera K, Leimkuhler B, Skeel RD. *J Comput Chem.* 1995; 16:1192–1209.
19. Eastman P, Pande VS. *J Chem Theory Comput.* 2010; 6:434–437. [PubMed: 20563234]
20. Miyamoto S, Kollman PA. *J Comput Chem.* 1992; 13:952–962.
21. Hess B, Bekker H, Berendsen HJC, Fraaije J. *J Comput Chem.* 1997; 18:1463–1472.
22. Barrett, R.; Berry, M.; Chan, TF.; Demmel, J.; Donato, J.; Dongarra, J.; Eijkhout, V.; Pozo, R.; Romine, C.; der Vorst, HV. *Templates for the solution of Linear Systems: Building Blocks for Iterative Methods.* SIAM; Philadelphia PA: 1994.
23. Buatois, L.; Caumon, G.; Levy, B. *High Performance Computing and Communications, Proceedings.* Perrott, R.; Chapman, BM.; Subhlok, J.; DeMello, RF.; Yang, LT., editors. Vol. 4782. 2007. p. 358
24. van Meel JA, Arnold A, Frenkel D, Zwart SFP, Belleman RG. *Mol Simul.* 2008; 34:259–266.
25. Stone JE, Phillips JC, Freddolino PL, Hardy DJ, Trabuco LG, Schulten K. *J Comput Chem.* 2007; 28:2618–2640. [PubMed: 17894371]
26. Eastman P, Pande VS. *J Comput Chem.* 2010; 31:1268–1272. [PubMed: 19847780]
27. Anderson JA, Lorenz CD, Travesset A. *J Comput Phys.* 2008; 227:5342–5359.
28. Tuckerman M, Berne BJ, Martyna GJ. *J Chem Phys.* 1992; 97:1990–2001.

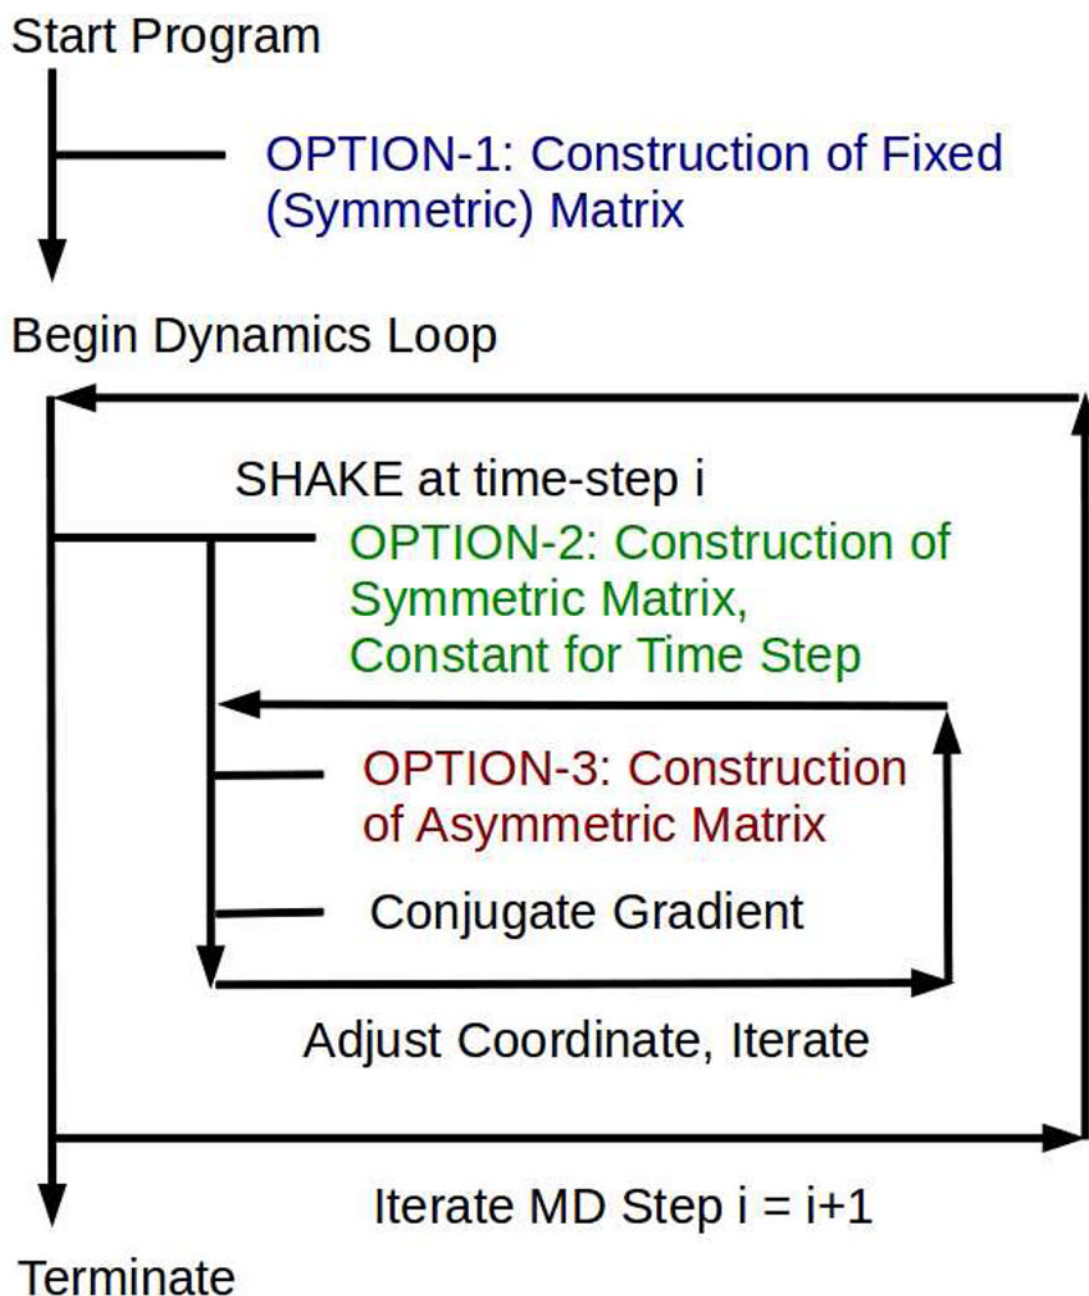


Figure 1.

An illustration of the different options in matrix SHAKE. Option-1 computes a fixed matrix that does not change throughout the calculations. It implies that the bond angles are exactly fixed¹⁴ or approximately so.¹⁹ The advantage is that the fixed matrix is computed only once throughout the simulation. The disadvantage is that the fixed matrix is approximate and a relatively large number of iterations are required to determine the Lagrange multipliers accurately. We determine the Lagrange multipliers using conjugate gradient (CG) algorithm. Building on the fixed matrix of reference¹⁴ an implementation that does not constrain the angles was proposed¹⁹. This variation is likely to require even larger number of iterations to converge to an accurate solution. In general it is not possible to retain high accuracy and to avoid iterative solutions of the Lagrange multipliers (and coordinates) with the exception of

very small systems for which analytical solutions are available (e.g. SETTLE²⁰). In the second option a symmetric matrix is constructed every time step. Since the matrix is sparse the cost of generating it each time step is small. The reduction in the number of CG iterations to achieve high accuracy is significant. This results in a more efficient algorithm. The third option of generating asymmetric matrix every CG step is not considered here since CG cannot use an asymmetric matrix without costly symmeterization, for example by using A^tA instead of A .

\$watermark-text

\$watermark-text

\$watermark-text




Figure 2. Parallel reduction (summation of array elements) in one block of 8 GPU threads. In this example, the values in the top row of rectangles are examples of vector entries we wish to add. Three operations (in parallel) are required to accomplish this reduction of 8 numbers. See also appendix B. Doing the block-wise shared memory parallel reduction in the CG kernel avoids the intermediate storage of partial results in global memory. Only the final block reduction requires global memory access.

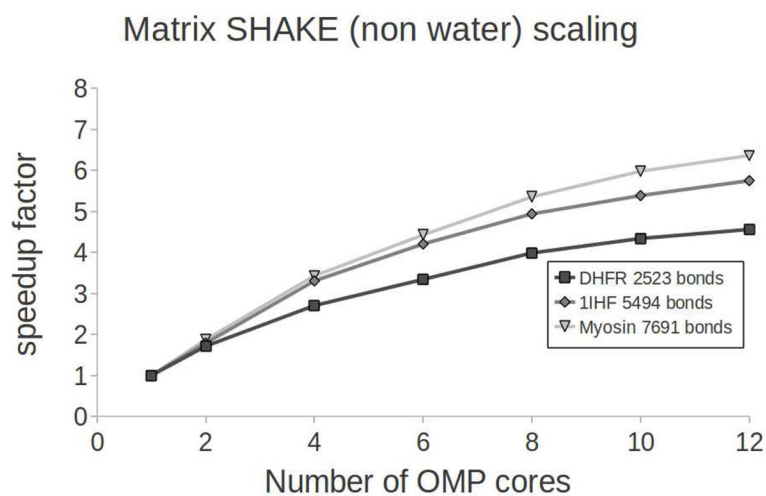


Figure 3. Speedup of the conjugate gradient SHAKE algorithm (coordinate + velocity) with the number of Open-MP cores used. Results are shown for DHFR, 1IHF, (protein + DNA) and myosin: the constraints for water are done in a separate calculation. These data were obtained on a 12-core node. The load of SHAKE bonds per thread in case of DHFR is only ~200 on the twelve cores. Scaling improves with larger systems but even our largest test system, myosin, balances to only 641 bonds per core making further improvement with system size likely. Fastest GPU results (one GTX 480) are five-fold faster than 1 CPU core so translate to at least 5 CPU cores for the largest systems.

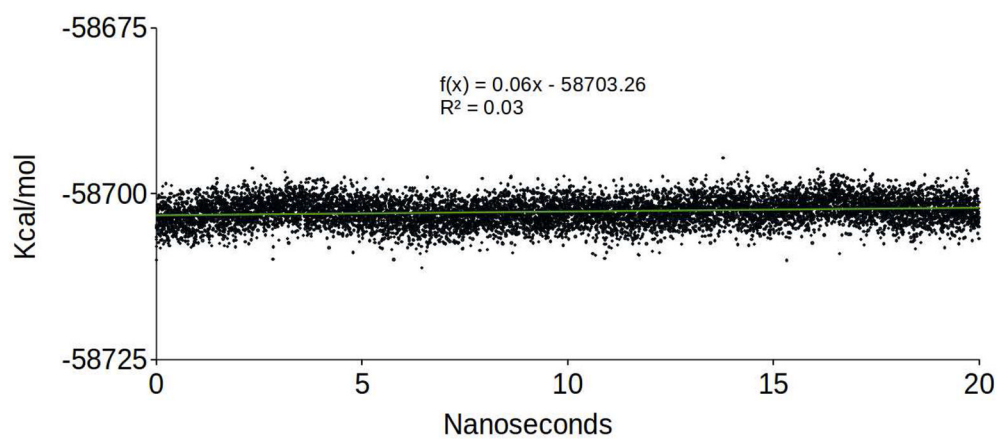


Figure 4. Energy conservation in the DHFR system. See legends of Table 1 and Table 4 for the simulation parameters. The total energy (in kcal/mole) is plotted as a function of time (in nanoseconds). A linear fit of the energy as a function of time suggests negligible energy drift, which we estimated as 0.1 percent relative error in a period of one microsecond.

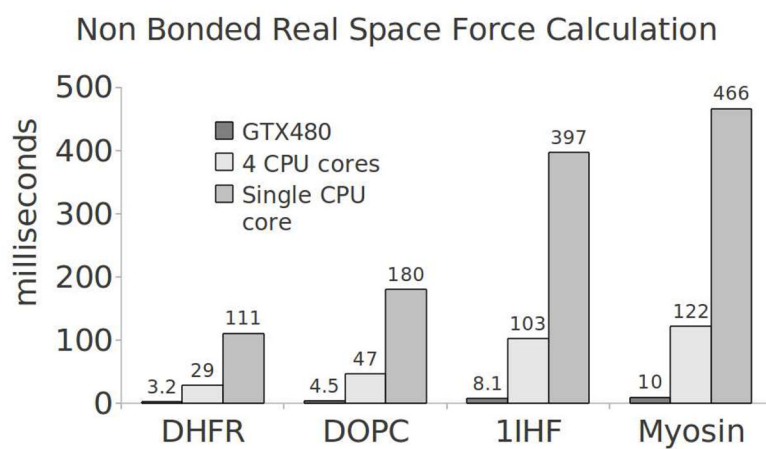


Figure 5. Calculation time of the non-bonded real space forces. Non-bonded real space force calculation on a single CPU core, 4 CPU cores and a GTX480. GPU Speedup factors versus a single CPU core are 37, 40, 49 and 47. For more details on the systems see Table 1.

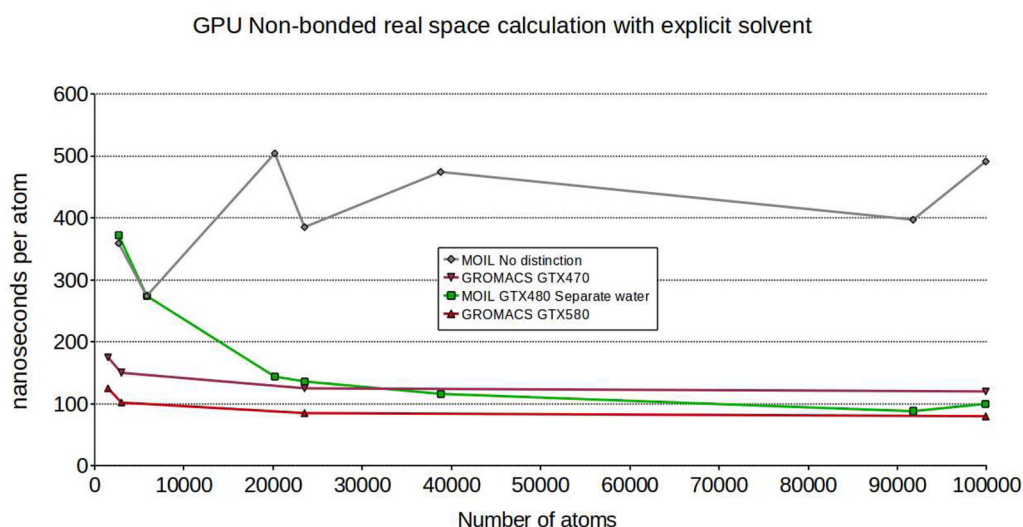


Figure 6.

Time of non-bonded real space force calculation as a function of the number of atoms (time per atom shown) on the GTX480 GPU. The same seven test systems (see Table 1) are used. The GPU non-bonded real-space calculation is done with and without utilization of separate water force kernels. The GPU performs better with an increasing number of atoms in the system. The system at ~39K atom is a DOPC membrane with relatively less water molecules compared to other systems. Other fluctuations in separate water trends can be due to ionic concentration and ion distribution (more non-water water interactions). The additional overhead of separate water kernels causes slower results in small systems. Comparing our results to those of GROMACS (cores flyer: www.binarybio.com) we note that they use one faster and a second slower GPUs (GTX470 and GTX580) in their benchmarks. We used an intermediate GPU (GTX480). For systems larger than 20,000 atoms the results are comparable.

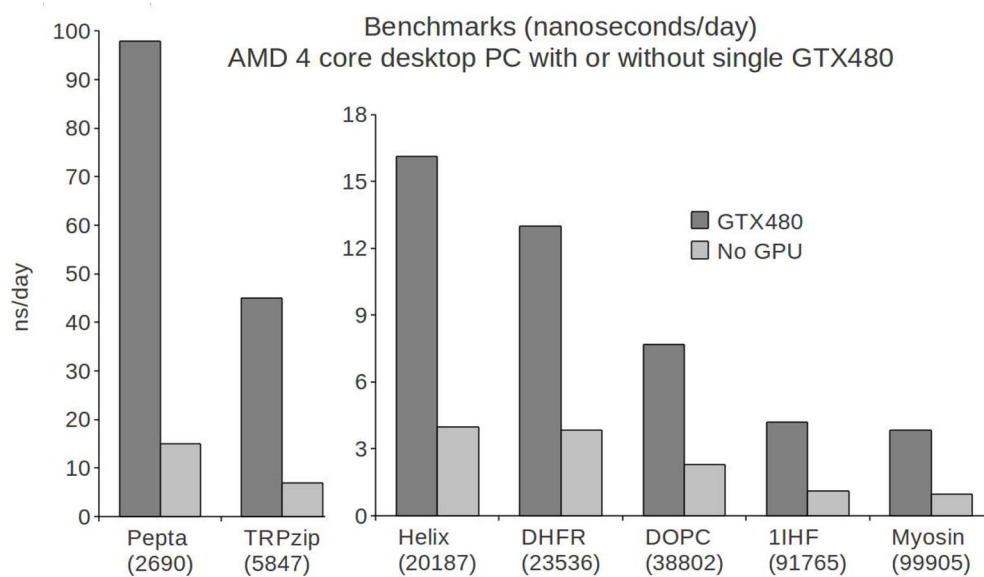


Figure 7. MOIL-opt with GPU overall performance. Although the speedup of real-space non-bonded force calculation does not degrade on the GPU with number of atoms (see Fig. 4), the overall speedup decreases with number of atoms due to newly exposed bottlenecks such as the PME (see Fig. 8). The overall performance on DHFR (13 ns/day) is similar to results reported for GROMACS 4.5 website from the PME results).

Fraction of total calculation time

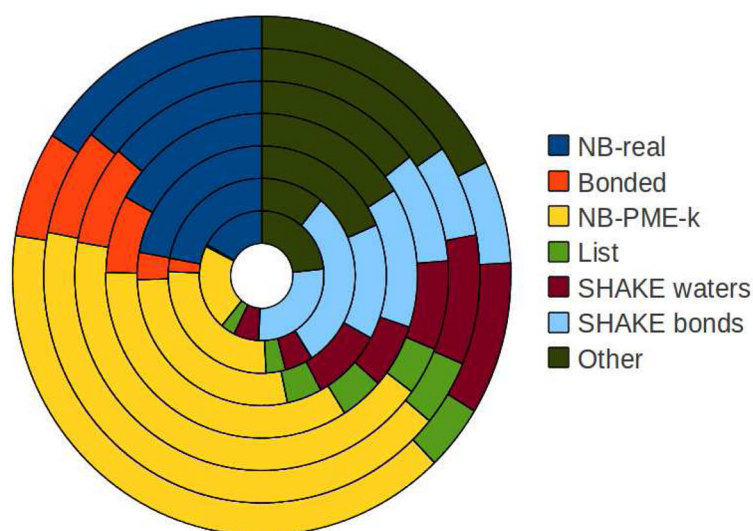


Figure 8. Fraction of computation cost of different parts of the molecular dynamics algorithm on a 4 core desktop PC equipped with a single GTX480 GPU. All 4 CPU cores and GPU are used here. The inner time of RESPA is 2fs and the outer time 4fs. In accord with our previous publication the PME reciprocal calculation is not repeated every step⁵. The seven test systems are shown as different layers in the pie-chart. The outermost layer is the largest system (myosin) and innermost pepta (see Table 1 for more details). Note that positive times per step are shown while actual time per step includes a “negative” pie slice contribution because some time is gained from asynchronous execution on the GPU. This gain is not shown here. It is coming from SHAKE (see Table 3). Another gain is the simultaneous calculation of bonded interactions and a part of the reciprocal sum calculation on the CPU while the GPU computes the real non-bonded interactions.

Table 1

Test systems. All seven systems listed below use explicit solvation. Unless otherwise noted, all results are obtained with tabulated conditions on a 4-core AMD Phenom 965 desktop PC equipped with a GTX 480 GPU. The solvation box sizes in angstroms are $30.1 \times 30.1 \times 30.1$ (Pepta), $40.0 \times 40.0 \times 40.0$ (TRPzipper), $59.0 \times 59.0 \times 59.0$ (Helix), $62.2 \times 62.2 \times 62.2$ (DHFR), $65.0 \times 65.0 \times 120.0$ (DOPC-NaCl), $99.8 \times 99.8 \times 99.8$ (IIHF) and $90.0 \times 112.5 \times 105.0$ (Myosin). Cutoff list (update frequency) calculation: “d-l” is the cutoff distance used in the atomic neighbor list generation, update “frequency” is the number of steps between list updates and “d-f” is the cutoff distance used in the force calculation. All the water molecules are constrained with a specialized version of Matrix SHAKE for individual water molecules (the inverse of the symmetric and constant **A** for a water molecule is pre-computed at the beginning of the calculations and used ever after).

	Nr. of atoms	Number of bond constraints	PME-DP grid	Cutoff d-l(freq)/d-f	Nr. of water molecules
Pepta	2690	43	$16 \times 16 \times 16$	10.0(8) 9.0	882
TRPzipper	5847	227	$32 \times 32 \times 32$	9.99(8) 9.0	1875
Helix	20187	388	$48 \times 48 \times 48$	9.8(8) 9.0	6588
DHFR	23536	2523	$48 \times 48 \times 48$	10.0(8) 9.0	7012
DOPC-NaCl	38802	6784	$48 \times 48 \times 96$	10.0(8) 9.0	10502
IIHF	91765	5494	$64 \times 64 \times 64$	9.97(8) 9.0	28763
Myosin	99905	7691	$64 \times 64 \times 64$	10.0(8) 9.0	30776

Table 2

Use of fixed constraint matrix **A** versus a matrix which is updated every time step (see also figure 1). Explanation of symbols: M-up is an **A** matrix which is updated every time step. M-fix is a fixed matrix which is generated once at the beginning of the calculations and used ever after. PCG-fix-it=7 is the use of a parallel conjugate gradient algorithm with a fixed number of iterations (in this case seven). Similarly PCG-var-it=7 means that the number of iterations was set to a maximum of seven, but the calculation was terminated if convergence was detected at a smaller number of iterations. This may be expected to save time, however the check (for convergence) is less efficient on the GPU and on the CPU. SHAKE-B is the standard use of SHAKE by the bond relaxation method¹². The numbers reported below each example name are reported in the format x.xx(yy.y)z.zz and means the following: x.xx is the average number of overall SHAKE iterations that are required for convergence. Note the significantly larger number of iterations required for coordinate adjustments when the matrix is fixed. yy.y is the total number of conjugate gradient iterations required per time step. If “*” is added to the left of the number then the number CG of iterations was fixed, otherwise the number of iterations was determined according to convergence of the conjugate gradient algorithm (“var” option). It is not surprising that the number of CG evaluations is smaller for the “var” option. However, consider the most important measure which is z.zz and is the time of executing the SHAKE algorithm in microseconds per bond. The shortest times obtained are faster for the “M-up PCG-fix-it=7” option.

Method	DHFR	IHF	DOPC	Myosin
M-up PCG-fix-it=7	8(56*),2.69	8(56*)2.95	6(42*),1.73	7.1(49.6*),2.72
M-fix PCG-fix-it=7	13.6(95.5*),3.81	12.8(89.7*),3.89	10.7(75.1*),2.57	12.6(88.1*),3.77
M-up PCG-var-it=7	11.8(63.4),2.85	10.9(66.1),2.99	8.0(47.7), 1.75	10.3(61.3),2.77
M-fix PCG-var-it=7	17.0(82.9),3.99	16.0(81.5),3.91	13.7(66.5),2.59	16.1(79.0),3.83
M-up PCG-var-it=32	8.8(91.2),3.16	7.5(92.6),3.35	6.0(67.3),1.96	7.25(84.8),3.15
M-up GPU PCG-fix-it=7	7.8(55*),1.07	8(56 *),0.68	6(42.1)0.44	7.1(49.6*)0.53
M-fix GPU PCG-fix-it=7	13.7(95.9*),1.6	13(90.7 *),0.86	10.7(75.1*)0.61	12.6(88.1*)0.86
SHAKE-B (bond relaxation method)	61.4(n/a),4.36	66.6(n/a),4.38	21.62(n/a),1.44	61.1(n/a),4.0

\$watermark-text

\$watermark-text

\$watermark-text

Table 3

SHAKE on GPU vs multi-CPU SHAKE, system Myosin, ~100K atoms, ~100K non water bonds. The execution of the matrix SHAKE on the GPU + one CPU core to instruct/monitor the GPU allows concurrent execution of water matrix SHAKE on the remaining CPU cores (so at 4 CPU cores, 3 cores are shaking water molecules). Asynchronous computational gain for SHAKE per time step is illustrated. We also included the results for the bond relaxation approach, which is the most commonly used SHAKE variant. In all calculations water molecules were constrained by a specialized matrix SHAKE algorithm a calculation conducted solely on the CPU

	CPU cores	Time water	Time other bonds	Gain asynch	Total time SHK
Matrix SHAKE: CPU	4	5.75	5.70	none	11.45
Matrix SHAKE: GPU	1 + (3 water)	6.04	4.07	2.53	7.58
Bond-relx	1 + (3 water)	6.04	30.8	6.04	30.03

Table 4

Energy conservation in the DHFR system. The RESPA algorithm₂₈ was used with an inner time step of 2fs and outer time step of 4fs. Quadratic force interpolation for electrostatics with 256 knots in the lookup table is used as described in reference⁵. Under “Cut” we report the cutoff distance used to generate the non-bonded list (10.0Å) and the actual cutoff used in the calculation of the force (9.0 Å). The tolerance of error is defined as $|\Delta b/b|$ where b is the ideal distance and Δb the deviation from ideal value and is set to 10^{-10} .

Cut	Time step	SHAKE	DRIFT (% μ s)
10.0/9.0	2.0 fs	Light	0.4
10.0/9.0	2.0 fs	Matrix	0.1