

Enabling Simulation of High-Dimensional Micro-Macro Biophysical Models through Hybrid CPU and Multi-GPU Parallelism

Steven Cook

University of California Riverside
scook005@cs.ucr.edu

Tamar Shinar

University of California Riverside
shinar@cs.ucr.edu

ABSTRACT

Micro-macro models provide a powerful tool to study the relationship between microscale mechanisms and emergent macroscopic behavior. However, the detailed microscopic modeling may require tracking and evolving a high-dimensional configuration space at high computational cost. In this work, we present a parallel algorithm for simulation a high-dimensional micro-macro model of a gliding motility assay. We utilize a holistic approach aligning the data residency and simulation scales with the hybrid CPU and multi-GPU hardware. With a combination of algorithmic modifications, GPU optimizations, and scaling to multiple GPUs, we achieve speedup factors of up to 27 over our previous hybrid CPU-GPU implementation and up to 540 over our single-threaded implementation. This approach enables micro-macro simulations of higher complexity and resolution than would otherwise be feasible.

CCS CONCEPTS

• **Computing methodologies** → **Modeling methodologies; Multiscale systems; Vector / streaming algorithms; Massively parallel and high-performance simulations;** • **Applied computing** → **Biological networks; Systems biology;** • **Hardware** → **Hardware accelerators;**

KEYWORDS

Multiscale modeling and simulation, Micro-macro methods, Kinetic theory, GPGPU, MPI

ACM Reference format:

Steven Cook and Tamar Shinar. 2017. Enabling Simulation of High-Dimensional Micro-Macro Biophysical Models through Hybrid CPU and Multi-GPU Parallelism. In *Proceedings of 8th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems, Denver, CO USA, November 2017 (ScalA'17)*, 8 pages.
<https://doi.org/>

1 INTRODUCTION

Active gels exhibit macroscopic flow structures driven by the detailed microscopic interactions of constituent elements. Pronuclear centering and migration and cytoplasmic streaming are two such

examples, both being critical cellular processes driven by filament-motor mixtures. Reduced-component studies have found these systems to be highly sensitive to the microscopic interactions between motors and filaments; for instance, the detachment time of a motor protein at a filament end affects whether filaments form networks of asters or vortices [12–14]. Additionally, the tens-of-nanometers sized motor proteins bind, walk along, and detach from micrometer-length filaments on a faster timescale than the filament network evolution. Simulating even a millimeter-sized system with such disparate length and time scales and sensitivity to detailed interactions thus poses a challenging computational problem. Tracking interacting Lagrangian particles can become infeasible with large quantities of microstructural elements.

A promising approach lies in micro-macro methods, which couple a kinetic theory model of the microstructure (here, the configuration of the motors and filaments in the active gel) to the macroscale continuum mechanical representation of a viscoelastic fluid [11]. Kinetic theory models have been applied in the study of biological active matter [9], self-propelled particles [18], and networks of neurons [1]. They enable detailed microscale modeling that would otherwise be lost via closure approximations in macroscopic modeling approaches, and are particularly useful at scales where tracking individual particles and their interactions would be prohibitive. Compared to purely macroscopic methods, micro-macro methods are more computationally demanding, as they require evolving the microstructure density in a potentially high-dimensional configuration space.

[8] and [2] developed a micro-macro model for a gliding motility assay, consisting of immersed rigid filaments that glide along motor proteins anchored to the substrate of a chamber immersed in viscous fluid. This model includes hydrodynamic and steric interactions between the filaments. A high-dimensional kinetic theory describes the evolution of the filaments and motors. To make this model computationally feasible, parts of the microscale computation were ported to the GPU using Nvidia's CUDA C language [16]. In this work, we enable faster and significantly more detailed simulations through holistic restructuring of this algorithm, aligning the computation and data flow with the underlying heterogeneous computational resources. Moreover, these changes facilitate scaling to multiple GPUs across separate machines with MPI. We further utilize a variety of CPU and GPU optimizations. Our work expands the range of micro-macro models which can be simulated by direct solution of the kinetic theory and coupling equations to models with higher dimensional configuration spaces, at higher resolutions [11]. To our knowledge, [8] is the first GPU-accelerated micro-macro kinetic theory-based simulation. [8] achieved up to 20x speedups over a single-threaded CPU implementation, while the algorithm

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
ScalA'17, November 2017, Denver, CO USA
© 2017 Copyright held by the owner/author(s).
ACM ISBN .
<https://doi.org/>




Figure 1: Gliding motility assay. Motor proteins (black) anchored to the substrate bind to filaments (green), walk along them and exert forces, then detach.

presented here achieves a further 27x speedup over [8] and [2]. Key to our approach is moving the microscopic scale and related tasks, which are smaller scale in both space and time, to the GPU and limiting CPU-GPU communications to the longer timescale of the filaments and fluid. Such holistic approaches are recommended to achieve scalability in heterogeneous environments [15], [6]. We note that our method does not suffer from common GPU simulation challenges encountered in various other approaches such as building adjacency lists [10], reordering storage based on cell location [21], dynamic, irregular data accesses [22], thread divergence [5], or neighbor exchanges of halo regions.

The paper is organized as follows. A description of the model and implementation is presented in Section 2, algorithmic, data flow, GPU, and MPI modifications are discussed in Section 3, results are discussed in Section 4, and we conclude in Section 5.


2 MOTILITY ASSAY MODEL

Figure 1 illustrates a gliding motility assay. The geometry consists of top and bottom plates separated by a narrow vertical gap, which is filled with a viscous fluid containing ATP fuel. At the bottom plate, motor protein tails are anchored to a substrate. The motor protein heads diffuse in solution, tethered to their tail by a flexible stalk. When a filament enters the capture radius of a motor protein head, the head may bind to the filament. As the bound motor head walks toward the filament plus end, it exerts force, causing the filament to glide in the opposite direction, until the motor head detaches. When many filaments are present, the underlying microscopic mechanism coupled with hydrodynamic and steric interactions give rise to a variety of emergent macroscopic behaviors such as a lattice of vortices [19]. Through our modeling and simulation, we aim to better understand the relationship between the microscale interactions and the macroscopic phenomena. Related problems of emergent self-organization from simple interactions include gliding and swarming of birds, fish, and bacteria.

The model equations are presented in simplified, nondimensionalized form in Table 1. The filament density is parameterized by center-of-mass location \mathbf{x} , orientation \mathbf{p} , and time t as $\Psi(\mathbf{x}, \mathbf{p}, t)$.

The configuration space of bound motors is higher dimensional, as we need to track the center-of-mass position \mathbf{x} and orientation \mathbf{p} of the filament a bound motor with tail anchored at \mathbf{r}_0 is bound to, along with its arclength parameter along that filament s .

This yields the high-dimensional density $\mathcal{M}_b(\mathbf{r}_0, \mathbf{x}, \mathbf{p}, s, t)$ of bound motors per filament, illustrated in the figure to the right. A key observation is that motors with tail anchored at position \mathbf{r}_0 on the assay substrate can only bind to filament sections that are within the capture radius of the motor stalk r_c . This greatly reduces the feasible configurations a motor protein head may be bound in, and obviates the need to track configurations $|\mathbf{x} + \mathbf{sp} - \mathbf{r}_0| > r_c$. We denote by $B_{r_c}(\mathbf{r}_0)$ all feasible $\mathbf{x}, \mathbf{p}, s$ configurations such that $|\mathbf{x} + \mathbf{sp} - \mathbf{r}_0| \leq r_c$. We do not track unbound (free) motor heads, only their tail position \mathbf{r}_0 , so the density of free motors $\mathcal{M}_f(\mathbf{r}_0, t)$ is two-dimensional. We model the filament and bound motor protein densities as distributed by a smooth Dirac delta function in z about a plane a small distance z_0 above the bottom plate, i.e., $\Psi(\mathbf{x}, \dots) = \Psi_{z_0}(\mathbf{x}_2, \dots)\delta(z)$ and $\mathcal{M}_b(\mathbf{r}_0, \mathbf{x}, \dots) = \mathcal{M}_{b,z_0}(\mathbf{r}_0, \mathbf{x}_2, \dots)\delta(z)$. We thus evolve the lower-dimensional \mathcal{M}_{b,z_0} and Ψ_{z_0} in our simulation. We drop the z_0 from Ψ and \mathcal{M}_b in the remainder of the paper for brevity.



Filament equations

$$\partial_t \Psi + \nabla_2 \cdot (\dot{\mathbf{x}}_2 \Psi) + \partial_\theta (\dot{\theta} \Psi) = 0 \quad (1)$$

$$\dot{\mathbf{x}}_2 = -V_{\text{sp}} \mathbf{p}_2 + \mathbf{u}_2 + U_{t,\parallel}^0 \mathbf{p}_2 \mathbf{p}_2 : \nabla_2 \mathbf{D}_{2,z_0} - D_{t,\parallel} \nabla_2 \ln \Psi \quad (2)$$

$$\dot{\theta} = \nabla_2 \mathbf{u}_2 + U_r^0 \mathbf{D}_{2,z_0} : \mathbf{p}_2^\perp \mathbf{p}_2 - D_r \partial_\theta \ln \Psi \quad (3)$$

Motor equations

$$\partial_t \mathcal{M}_b + \partial_s \mathcal{M}_b = -k_{\text{off}} \mathcal{M}_b + k_{\text{on}} \mathcal{M}_f \mathbb{1}_{D_{r_c}} \quad (4)$$

$$\mathcal{M}_{b,\text{coarse}} = \iiint \mathcal{M}_b \Psi ds dx_2 d\theta \quad (5)$$

$$\mathcal{M}_f = \mathcal{M} - \mathcal{M}_{b,\text{coarse}} \quad (6)$$

Fluid equations

$$-\nabla_2^2 \mathbf{u}_2 - \frac{1}{\varepsilon^2} \partial_{zz} \mathbf{u}_2 + P_0 \nabla_2 q = \sigma_f \nabla_2 \cdot \boldsymbol{\sigma}^f - \sigma_t \nabla_2 \cdot \boldsymbol{\sigma}^t + \mathbf{F}_2 \quad (7)$$

$$-\nabla_2^2 w - \frac{1}{\varepsilon^2} \partial_{zz} w + P_0 \partial_z q = 0 \quad (8)$$

$$\nabla_2 \mathbf{u}_2 + \partial_z w = 0 \quad (9)$$

Motor force

$$\mathbf{F}_2 = F \iiint \mathbf{p}_2 \delta(y_2 + \frac{l}{L} s \mathbf{p}_2 - \mathbf{x}_2) \Psi \mathcal{M}_b ds d\mathbf{r}_0 dy_2 d\theta \quad (10)$$

Table 1: Summary of model equations for the filament and motor protein densities, the macroscopic fluid equations, and the motor force that couples them.

We represent the fluid velocity in three dimensions, with periodic boundary conditions in the x and y dimensions and no-slip conditions in the z dimension at the top and bottom plates. The system evolves on two timescales; the motors bind to, walk along, and unbind from the filaments on a faster timescale than the filaments and fluid evolve. Bound motor heads generate forces that are spread onto the fluid in an immersed boundary method fashion [17]. Together with stress terms arising from filament inextensibility and

steric interactions [4], the motor forces (Eq. (10)) couple the densities Ψ and \mathcal{M}_b to the fluid velocity (Eq. (7)). The x and y dimensions are discretized over a regular square grid, and the z dimension is discretized over an adaptive grid that is finely resolved around $z = z_0$ near the bottom plate and becomes coarser moving toward the upper plate. Allowable filament orientations are constrained to the (x, y) plane, so we can represent $\mathbf{p} = (\cos \theta, \sin \theta, 0)^T$. Orientation θ and arclength parameter s are discretized uniformly with the same resolution.

Algorithm 1 summarizes the process for evolving the filaments, motor proteins, and fluid velocity as in [8], [2]. First, we compute the adaptive time steps based on their stability conditions, with outer time step dt restricted by the advective fluxes in Eq. (1), and the inner time step dt^* restricted by the motor speed and binding/unbinding rates in Eq. (4). Next, $\Psi(t + dt)$ is solved using Crank-Nicolson for the diffusive terms and Adams-Bashforth 2 with upwinding for the advective terms (line 4). The bound motor density evolution routine in lines 7-9 performs the motor protein advection (bound motor heads walking along the filaments toward their plus-ends) and applies a Superbee flux limiter, as well as simulates the binding and unbinding of free and bound motor proteins respectively. After every configuration of $\mathcal{M}_b(\mathbf{r}_0\mathbf{x}, \theta, s, t)$ for a particular \mathbf{r}_0 has been updated, the integral $\mathcal{M}_{b,coarse}(\mathbf{r}_0, t) = \iiint \mathcal{M}_b(\mathbf{r}_0, \mathbf{x}, \theta, s, t)\Psi(\mathbf{x}, \theta, t)d\mathbf{x}d\theta ds$ is calculated at the same \mathbf{r}_0 to ensure that the number of bound motors does not exceed the total number of motors at \mathbf{r}_0 . If so, all \mathcal{M}_b configurations with that \mathbf{r}_0 are scaled down to conserve the total number of motors before \mathcal{M}_f is calculated. We next update the free motor density in line 9. The extra stress terms σ^f, σ^t arising from filament inextensibility and steric interactions [2] are computed as moments of Ψ in line 13. We perform a two-dimensional FFT and solve the transformed system of fluid equations for the three velocity components and pressure $\hat{u}, \hat{v}, \hat{w}, \hat{q}$ at every position on the z grid, then perform an inverse FFT to obtain the three-dimensional fluid velocity \mathbf{u} in line 13.

In a single-threaded implementation, the high dimensionality of \mathcal{M}_b makes the computations in lines 7-9 and line 11 prohibitively expensive for even moderate grid resolutions and experiment times. Fortunately, \mathcal{M}_b can be computed explicitly and easily parallelized over \mathbf{r}_0 . Thus in [2, 8], the \mathcal{M}_b and \mathbf{F} computations are ported to the GPU. On the other hand, the Ψ equation is stiff due to the diffusion terms and is computed semi-implicitly on the CPU. This decomposition of work is similar to several hybrid reactive flow solvers [15].

To perform the integral in line 11, [2, 8] use independent GPU threads to compute the integral at each \mathbf{r}_0 accumulating the partial results in thread-local storage, limiting the use of atomic operations to the final reduction over nearby \mathbf{r}_0 at each \mathbf{x} . Extra stress tensor calculation, forward and reverse fast Fourier transforms, and computation of the independent semispectral systems are all multithreaded on the CPU. In this work, we expand upon this hybrid computational approach as described below.

3 ACCELERATION METHODOLOGY

Our primary focus in this work is significant performance improvement through targeted algorithmic design enabling a multi-GPU

Algorithm 1: Evolution scheme for the coupled microtubule density, motor protein density, and fluid velocity equations.

```

1 Initialize  $\Psi$  and  $\mathcal{M}_b$ 
2 while  $t < t_{end}$  do
3   Compute adaptive time steps  $dt, dt^*$ 
4   Compute filament density  $\Psi(t + dt)$  (Eqs. (1)-(3))
5   set  $t_{end}^* = t + dt$ 
6   while  $t^* < t_{end}^*$  do
7     Compute bound motor density  $\mathcal{M}_b(t^* + dt^*)$  (Eq. (4))
8     Update coarsened density  $\mathcal{M}_{b,coarse}$  (Eq. (5))
9     Update free motor density  $\mathcal{M}_f$  (Eq. (6))
10  end
11  Calculate motor force  $\mathbf{F}_2$  (Eq. (10))
12  Calculate extra stresses  $\sigma^f, \sigma^t$  (Eq. (7))
13  Calculate fluid velocity  $\mathbf{u}$  (Eqs. (7)-(9))
14 end

```

decomposition, as well as single-GPU optimizations, described in detail below. Through these efforts we are able to scale to higher resolutions than previously possible and achieve up to 27x total simulation acceleration in a four GPU configuration. Use of additional GPUs is supported and should provide further speedup with similarly excellent scaling, although we did not test this in this work.

3.1 Holistic Algorithmic and Data Flow Restructuring

A primary goal of our approach is to restructure the algorithm so that the memory-intensive microscale motor protein data resides solely on the GPU, and only the smaller, coarsened data is transferred to/from main memory. A secondary goal is to support a multi-GPU decomposition. Additionally, we remove synchronization barriers and reduce GPU memory consumption by two-thirds. Figure 2 summarizes the changes and details follow.

Independent Time Steps. The original algorithm calculated a global dt^* and updated all \mathcal{M}_b configurations by this fixed time step to time $t + dt$, hindering performance in several ways. First, it artificially limits the inherently independent per-cell update operations, some of which may be able to complete in fewer steps as their local configuration and stability restrictions allow. Second, it requires an expensive reduction operation over the entire bound motor density microstructure. Third, if \mathcal{M}_b is distributed over multiple GPUs as desired, the reduction creates an unnecessary synchronization barrier. We instead compute a local dt^* for each \mathbf{r}_0 at the beginning of each inner time step and update \mathcal{M}_b at each \mathbf{r}_0 asynchronously. The most significant benefit of this change is enabling the multi-GPU implementation. Stability and accuracy were not adversely affected.

Numerical Integration Scheme. The algorithm in [2] used Adams-Bashforth 2 for time integration of the motor densities, which maintains the \mathcal{M}_b array at three distinct time points (t^{n+1}, t^n, t^{n-1}) . We instead use Runge-Kutta 2, which only requires the \mathcal{M}_b array at t^{n+1} and t^n . This change reduces GPU memory requirements by one-third while causing negligible impact on computation time.

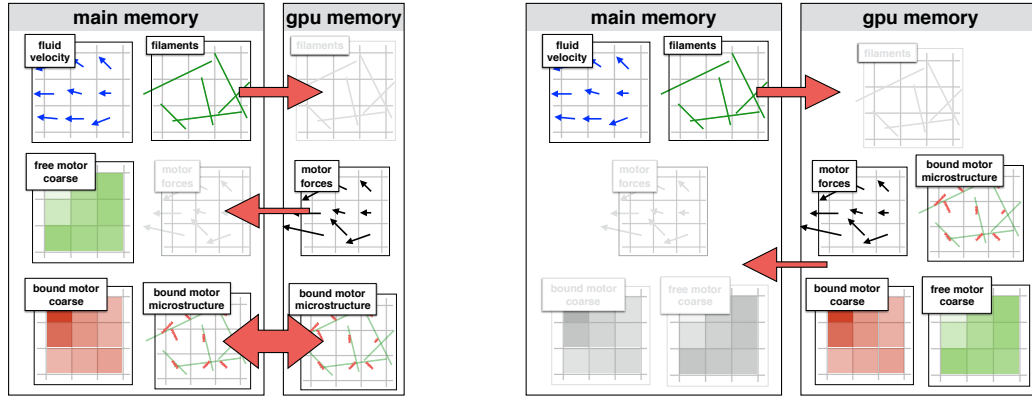


Figure 2: Data residency and computation before (left) and after (right) restructuring. Left, transfer of the microstructure creates significant communication overhead. Right, the bound motor microstructure now fully resides on the GPU eliminating expensive transfers, and the coarse bound and free motor densities are calculated on GPU and transferred back to main memory.

With these improvements, higher-resolution \mathcal{M}_b density representations may reside in scarce GPU memory.

Mixed Precision. We developed a mixed precision approach whereby we store and update \mathcal{M}_b and \mathcal{M}_f in single precision floating point while keeping the rest of the simulation as double precision. This saves space and improves performance without causing appreciable change in simulation behavior.

Data Residency. The algorithm in [2] updated \mathcal{M}_b one piece at a time due to GPU memory constraints, then transferred the complete updated \mathcal{M}_b to the GPU for the motor force calculation. With the new numerical integration scheme and the use of mixed precision, we have enough GPU memory to store the high-dimensional microstructure data \mathcal{M}_b solely on the GPU. This eliminates the overhead of transferring copies of \mathcal{M}_b before, during, and after the \mathcal{M}_b update. Since \mathcal{M}_b is required in order to calculate $\mathcal{M}_{b,coarse}$ and \mathbf{F} , we also do those calculations on the GPU, and transfer results to the CPU. $\mathcal{M}_{b,coarse}$ and \mathbf{F} are both macroscale data structures, and hence incur lower communication overhead. Finally, Ψ , which is also stored on the macroscale, is transferred as before. The updated data flow is shown in Figure 3.

Pseudocode describing the new GPU kernel is presented in Algorithm 2. The result is one large kernel that fully updates \mathcal{M}_b at each independent \mathbf{r}_0 value to $t + dt$ in as many steps as needed, using a local adaptive time step. The new memory access pattern is more amenable to caching as well, as each running block of threads on each GPU reads the same contiguous memory for all \mathcal{M}_b configurations at a fixed \mathbf{r}_0 location repeatedly until those configurations are fully updated before moving on.

We use CPU parallelism via OpenMP to further accelerate the simulation, specifically in the calculations of the fluxes and stress tensors for the fluid solves, the outer global time step dt calculation, and construction of the Ψ advection matrix. After moving the dt^* calculation into the motor force update kernel and multithreading the dt calculation, time step calculation becomes a negligible component of the total computation time.

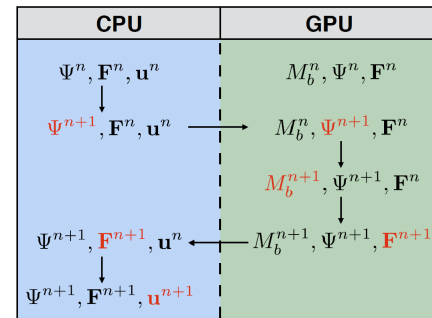


Figure 3: Residency and evolution of state from time step n to time step $n + 1$ on CPU and GPUs in new algorithm. Red indicates the quantity updated through computation or data transfer.

3.2 GPU Optimizations

This section describes various optimizations of the GPU kernel shown in Algorithm 2. Combined, these single-GPU optimizations yield an average improvement of 4.7-7.5x depending on resolution. The optimizations are described below and the individual effect of each is listed in Table 3.

Mixed precision. As previously detailed, switching \mathcal{M}_b , $\mathcal{M}_{b,coarse}$, and \mathcal{M}_f to single precision halves the GPU memory requirement. In addition, it provides a 4.3x to 5.8x speedup in our bound motor density evolution routine. This improvement will depend on the clock cycle ratio between single and double precision arithmetic for a given GPU family.

Fast Math. Compiling with CUDA's fast math library provides additional savings without noticeable change in simulation behavior. Accelerations of 1.35x were typical.

Launch Bounds. The `launch_bounds` macro in CUDA may be used to instruct the compiler to ensure a user-specified maximum number of concurrent threads and threads per block running on each GPU Streaming Multiprocessor (SM). Using a `launch_bounds`

configuration of 128 threads/block and 8 simultaneous blocks per SM gives the best performance of all configurations tested. Register spilling to global memory does occur at this configuration as each thread is limited to 32 registers. Newer architectures with more registers per SM will likely see immediate improvement by both reducing register spilling and enabling more threads per block. Accelerations of 1.2x were typical at the higher inner resolution and negligible at the lower inner resolution.

Dimension Mapping. CUDA threads are executed in simultaneous warps of 32 threads each, grouped first by their x-index then by their y-index. Since coalesced memory accesses are desirable for performance, the bound motor density evolution kernel was modified so that a thread's x-index maps to the s-index and the y-index maps to the θ -index. With this mapping threads executing in a warp will access \mathcal{M}_b storage in a coalesced fashion since sequential s-indices are contiguous as the innermost array indices. Accelerations of 1.16x to 1.57x were observed.

Reordering Storage. The Superbee flux limiter operates in the arclength s dimension. Since threads in a warp operate on subsequent arclength indices, and the flux limiter has a neighborhood access pattern of $(s - 2, s - 1, s, s + 1, s + 2)$, this gives coalesced memory accesses and pulls adjacent arclength data into the cache for subsequent iterations. The layout of memory in \mathcal{M}_b was modified to make s the innermost variable instead of θ in the storage of $\mathcal{M}_b(\mathbf{r}_0, \mathbf{x}, \theta, s)$, where \mathbf{x}, θ, s are represented as sequential flat four-dimensional arrays within a flat two-dimensional array over \mathbf{r}_0 . This prevents strides between subsequent s accesses. Another benefit to making s the innermost variable is that the value of $\Psi(\mathbf{x}, \theta)$ can be read after the θ loop instead of in the innermost loop. Reordering the loops in this fashion in the access-heavy motor force code resulted in a 1.54x acceleration. For the bound motor density update, accelerations of 1.1x were typical at the higher inner resolution and negligible at the lower inner resolution.

Unrolling Reductions. For the reduction step, [7] recommends manually unrolling a reduction when the number of remaining threads is less than the warp size (32 for our Tesla M2075), and performing part of a large reduction independently within each thread to reduce synchronization. We already follow the latter suggestion as each thread accumulates its contribution to $\mathcal{M}_{b,coarse}$ before storing this running sum in a shared memory array sized to the number of threads for the reduction step. We did not find meaningful performance improvements for the manual reduction.

Block Shaping. To update $\mathcal{M}_b(\mathbf{r}_0, \mathbf{x}, \theta, s)$ we assign one block of threads to each \mathbf{r}_0 position and map those threads to the \mathbf{x}, θ, s variables. For each \mathbf{x} , a two-dimensional block of threads is launched, with the threads' x and y indices corresponding to the innermost s and θ indices. Experimentation has shown 128 threads to be the optimal number in our implementation. The "block shaping" row of the optimizations table compares against running 256 threads in a 16x16 configuration. How these 128 threads are configured is important: $x = 8, y = 16$ runs faster than $x = 16, y = 8$. Both caching effects and memory coalescing play a role, and from our experience it is worthwhile to experiment with various configurations. Accelerations of 1.15x were typical at the higher inner resolution and accelerations of 1.65x were typical at the lower inner resolution.

Algorithm 2: Bound & free motor density update GPU kernel

```

1 Precondition:  $\Psi(t + dt)$  and  $dt$  are loaded into GPU memory.
2 Set  $t_{end}^* = t + dt$ 
3 for  $\mathbf{r}_0 \in \text{grid}$  do
4   while  $t^* < t_{end}^*$  do
5     Compute adaptive  $dt^*$ 
6     for  $\mathbf{x} \in B_{rc}(\mathbf{r}_0)$  do
7       for  $\theta \in B_{rc}(\mathbf{r}_0)$  do
8         for  $s \in B_{rc}(\mathbf{r}_0)$  do
9           Compute  $\mathcal{M}_b(\mathbf{r}_0, \mathbf{x}, \theta, s, t^* + \frac{dt^*}{2})$  (Eq. (4))
10          Compute  $\mathcal{M}_b(\mathbf{r}_0, \mathbf{x}, \theta, s, t^* + dt^*)$  (Eq. (4))
11        end
12      end
13    end
14    Compute  $\mathcal{M}_{b,coarse}(\mathbf{r}_0, t^* + dt^*)$  (Reduction) (Eq. (5))
15    if  $\mathcal{M}_{total}(\mathbf{r}_0) < \mathcal{M}_{b,coarse}(\mathbf{r}_0, t^* + dt^*)$  then
16      for  $\mathbf{x}, \theta, s \in B_{rc}(\mathbf{r}_0)$  do
17        Scale  $\mathcal{M}_b(\mathbf{r}_0, \mathbf{x}, \theta, s, t^* + dt^*)$ 
18        Set  $\mathcal{M}_{b,coarse}(\mathbf{r}_0, t^* + dt^*) = \mathcal{M}_{total}(\mathbf{r}_0)$ 
19      end
20    end
21    Compute  $\mathcal{M}_f(\mathbf{r}_0, t^* + dt^*)$  (Eq. (6))
22  end
23 end
```

To analyze performance of the \mathcal{M}_b and FGPU kernels, the Nvidia Visual Profiler v8.0 [3] was used. According to its output, arithmetic operations constitute the largest share of operations. No functional unit (load/store, arithmetic, control flow) is a bottleneck because of the balance of operations. We run the maximum possible number of simultaneous blocks per SM (8), but cannot run more threads per block without exhausting the available registers per SM. The result is a GPU occupancy of 66%, for which the profiler's heuristics report that increasing occupancy is unlikely to improve execution time. Our experience confirms this, as attempts to launch more threads per block to increase occupancy means decreasing registers per thread to keep the simultaneous blocks per SM maximized at 8, resulting in longer execution times. GPU occupancy is one of many factors that contributes to kernel performance, and it is possible to obtain high throughput at low occupancy levels [20].

3.3 Scaling to Multiple GPUs

As the spatial resolution of the \mathbf{r}_0 grid increases, two factors limit the performance of a single GPU. The first is that the number of blocks (each updating an independent \mathbf{r}_0) that can run concurrently on the GPU is limited by the number of SMs on the card, as we are running the maximum 8 simultaneous blocks per SM. Using two equivalent cards simultaneously doubles the throughput at which we can update the motor densities and calculate the motor force. The second factor is that once GPU memory is exhausted by the bound motor density (and scratch space for the intermediate values needed for numerical routines), additional large memory transfers to and from main memory become necessary every time step.

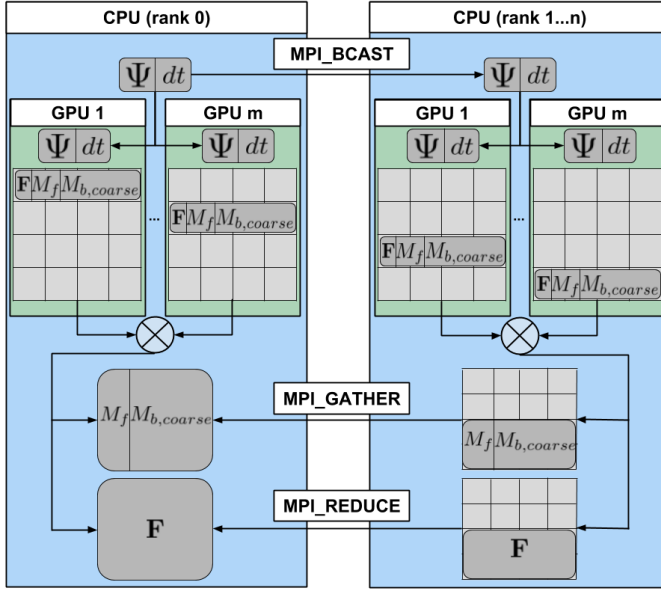


Figure 4: MPI control flow for multiple GPUs across multiple nodes. We achieve nearly ideal scaling of our GPU computation across multiple GPUs, indicating that the MPI overhead is negligible.

We expand our implementation to multiple GPUs using simultaneous CUDA streams and to multiple machines using MPI. After M_b is initialized at the beginning of the simulation, subsections of M_b are transferred to the memory of each GPU. The outer two-dimensional array of M_b (over r_0) is distributed among available GPUs by rows, which are contiguous in memory. Before invocation of the M_b and F kernels on the GPUs, the newly updated Ψ is broadcast to each with an MPI_BCAST from process rank 0. After the bound motor density update and motor force kernels complete, the M_f and $M_{b,coarse}$ values computed by each GPU are collected by process rank 0 using an MPI_GATHER operation. The motor force output F from each GPU contains overlapping force vectors that need to be summed together, so an MPI_REDUCE operation is used to combine them in process rank 0. From here on the fluid velocity update proceeds as normal. The process is summarized in Figure 4. OpenMPI 2.1 was used for this work.

4 RESULTS

We present results at different resolutions, scaling both the outer resolution of the x, y variables and the inner resolution of the θ, s variables. Increasing the outer resolution (x, y) affects the two-dimensional grid over which Ψ, M_f, M_b , and u are defined, and thus increases the workload across all steps of the simulation. Doubling the resolution of x and y increases by a factor of four the total workload of the M_b and F kernels. Increasing the inner resolution (θ, s) affects Ψ and M_b through their dependence on θ and M_b through its dependence on s . Doubling the resolution of θ and s increases by a factor of four the number of update tasks per thread in the motor density and force kernels, and adds two more iterations to the reduction step in the motor density kernel.

Simulations were run on one or more servers configured with 2 Tesla M2075 GPUs, 64 GB RAM, and dual AMD Opteron 6272 processors. Speedup factors for the algorithmic modifications and GPU optimizations vs. the original implementation in a single-node, single-GPU configuration are shown in Table 2. An average performance increase of between 5.75x and 9.98x per full simulation step is observed versus the original implementation. The bulk of the improvement comes in the bound motor update and motor force computations. Additionally, the dt computation is accelerated, the dt^* computation is moved onto the GPU, and the Ψ solve noticeably benefits from CPU acceleration.

We individually disable each GPU optimization and compare the running time for a single invocation of the M_b update kernel in Table 3. We see the largest performance improvement from switching from double precision to single precision, which affects both floating point arithmetic performance as well as cache and memory demand. The $256^2 \times 32^2$ resolution could not be tested with double precision on a single GPU as M_b exceeded GPU memory. The launch_bounds and storage reordering optimizations see their biggest impact when the inner variable resolution is increased. When this optimization was originally applied it showed a small improvement, but when it alone is removed from the final implementation, no discernable impact is observed.

Figure 5 shows that scaling the M_b and F kernels to multiple GPUs and across nodes is effective, with nearly ideal linear acceleration at the higher resolutions where acceleration is most needed. This demonstrates that the overhead to merge output between GPUs on the same machine plus the MPI overhead among multiple machines is small compared to execution time. As the inter-node communication consists of MPI broadcast, reduction, and gather operations, it is expected that scaling to 8 or more GPUs would likewise involve minimal overhead cost.

Table 4 summarizes overall performance of our optimized implementation using one, two, and four GPUs vs. the original single-GPU only implementation. We obtain higher accelerations for higher inner resolution sizes, which is desirable as we find the inner resolution of 16^2 too coarse at outer resolutions over 64^2 . Our maximum speedup factor over the original implementation was over 27x, obtained at the highest resolution. The simulation was previously limited to the $128^2 \times 32^2$ configuration given the running times involved. Reducing a day's worth of computation to less than one hour greatly facilitates the iterative exploration of the model's parameter space. Sample simulation results at the previously infeasible $256^2 \times 32^2$ resolution are shown in Figure 6.

5 CONCLUSIONS

Mapping the different scales of a simulation to different computational hardware, minimizing data transfers, and removing synchronization points like a global time step calculation allows us to explore the parameter space of our high-dimensional micro-macro simulation up to 540 times faster than a single-threaded implementation when using four GPUs simultaneously. This holistic approach significantly outperforms the commonly employed approach of accelerating individual functions in isolation [8]. Multi-node, multi-GPU overhead is minimal and the approach is expected to scale

outer \times inner resolution	$64^2 \times 16^2$	$64^2 \times 32^2$	$128^2 \times 16^2$	$128^2 \times 32^2$	$256^2 \times 16^2$	$256^2 \times 32^2$
total	332 (7.16)	854 (8.76)	2054 (6.37)	5026 (9.55)	8980 (7.36)	21408 (9.25)
Ψ solve	24 (6.17)	49 (5.37)	88 (6.44)	171 (7.88)	389 (7.28)	722 (9.19)
\mathcal{M}_b and F	193 (6.26)	682 (6.42)	1492 (6.03)	4366 (7.51)	7050 (6.65)	19125 (4.72)
fluid solve	107 (2.50)	112 (2.21)	447 (1.63)	456 (1.96)	1437 (2.00)	1438 (2.42)

Table 2: Effect of single-node, single-GPU optimizations. Average time per full outer time step (milliseconds) and speedup factors (bold) for original simulation vs. optimized simulation on a single machine with one GPU. Simulations with varied resolutions run to a fixed end time. The expensive dt^* calculation in the original simulation is now negligible, contributing to the increased overall speedup reported in the total time step row.

outer \times inner resolution	$64^2 \times 16^2$	$64^2 \times 32^2$	$128^2 \times 16^2$	$128^2 \times 32^2$	$256^2 \times 16^2$	$256^2 \times 32^2$
optimization removed:						
none	350	710	1490	3560	5960	15660
mixed precision	1520 (4.34)	4140 (5.83)	7390 (4.96)	19190 (5.39)	30170 (5.06)	X (X)
-fast-math	470 (1.34)	870 (1.23)	2010 (1.35)	4530 (1.27)	8050 (1.35)	19710 (1.26)
launch_bounds	340 (1.97)	860 (1.21)	1490 (1)	4270 (1.2)	5890 (99)	18200 (1.16)
dimension mapping	440 (1.26)	1120 (1.58)	1820 (1.22)	5210 (1.46)	8200 (1.38)	22420 (1.43)
reorder storage	350 (1)	820 (1.15)	1490 (1)	3960 (1.11)	6010 (1.01)	16950 (1.08)
unroll reduction	370 (1.06)	700 (0.99)	1510 (1.01)	3630 (1.02)	5920 (0.99)	15350 (0.98)
block shaping	550 (1.57)	820 (1.15)	2460 (1.65)	4330 (1.22)	9800 (1.64)	18070 (1.15)

Table 3: GPU optimizations. Time (milliseconds) and slowdown factors (bold) for the \mathcal{M}_b evolution kernel at different resolutions with various optimizations individually disabled.

outer \times inner resolution	$64^2 \times 16^2$	$64^2 \times 32^2$	$128^2 \times 16^2$	$128^2 \times 32^2$	$256^2 \times 16^2$	$256^2 \times 32^2$
1 node, 1 GPU	7.16	8.76	6.37	9.55	7.36	9.25
1 node, 2 GPU	9.70	14.04	9.73	16.50	11.98	16.60
2 node, 4 GPU	11.98	20.23	13.59	26.69	17.69	27.41

Table 4: Overall speedup resulting from our optimized, multi-GPU approach, as compared with original single-GPU simulation, as the number of nodes & GPUs is increased.




Figure 5: Multi-GPU scaling of the bound motor density and motor force computations is nearly ideal. The overhead of inter-node communication via MPI arising in the 4 GPU configuration does not have an appreciable effect.

well to a greater number of GPU accelerators. This approach capitalizes on the increasing prevalence of GPUs in high performance computing.

As the number of GPUs used increases the semispectral fluid solve and filament evolution update steps will become the next bottlenecks. Further adjustment of simulation flow to offload more of the fluid solve computations onto the otherwise idle CPU cores of non-root processes may then become cost effective.

It is our hope that our algorithmic design and breakdown of the various CPU and GPU optimizations will provide a useful reference for prioritizing optimizations in HPC software development and in porting of existing applications, where there is often an expectation that porting time should be recovered by faster runtimes. While the specific improvement will vary for different programs, quantifying the improvements corresponding to various optimizations contributes to the growing information in the literature regarding their efficacy [15].

ACKNOWLEDGMENTS

Removed for review.

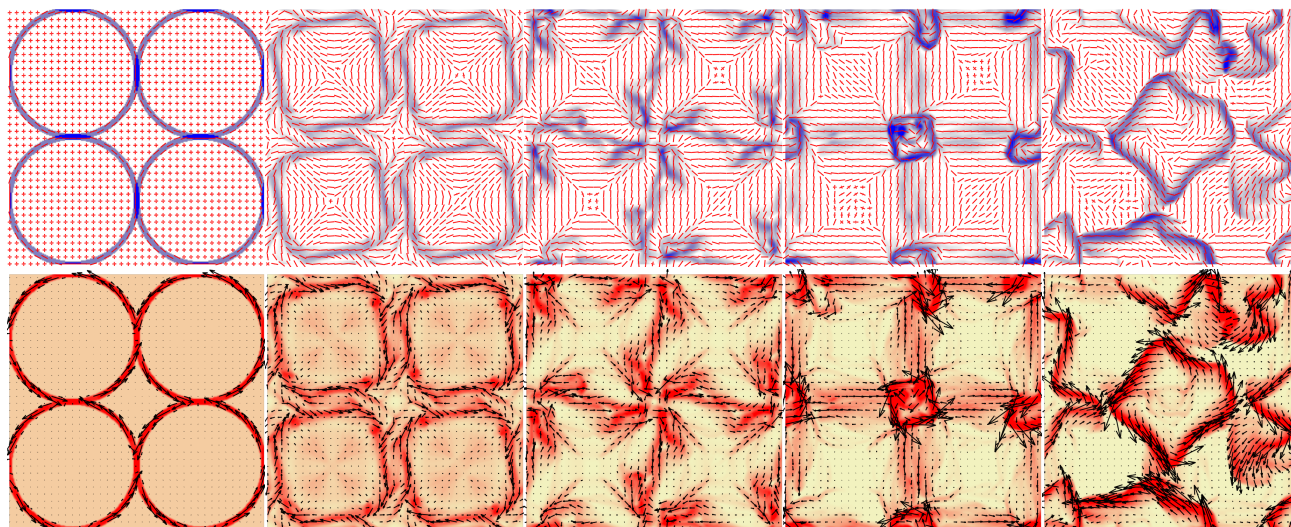


Figure 6: Sample simulation output at equally spaced timing intervals of the evolution of a lattice of overlapping filament rings from an overhead view of the assay, with periodic boundary conditions. Top row: filament density Ψ in blue with predominant orientation vector plotted every 8th cell in red. Bottom row: coarse bound motor density plotted in red with fluid velocity plotted as a black arrow every 8th cell.

REFERENCES

- [1] David Cai, Louis Tao, Michael Shelley, and David W McLaughlin. 2004. An effective kinetic representation of fluctuation-driven neuronal networks with application to simple and complex cells in visual cortex. *Proceedings of the National Academy of Sciences of the United States of America* 101, 20 (2004), 7757–7762.
- [2] S.C. Cook, C. Hohenegger, and T. Shinar. 2017. A Micro-Macro Framework for Analyzing Steric and Hydrodynamic Interactions in Gliding Assays. Submitted Jan 2017. *SIAM Journal on Multiscale Modeling and Simulation* (2017).
- [3] Nvidia Corporation. 2017. Nvidia Visual Profiler. <https://developer.nvidia.com/nvidia-visual-profiler>. (2017). <https://developer.nvidia.com/nvidia-visual-profiler>
- [4] Barath Ezhilan, Michael J Shelley, and David Saintillan. 2013. Instabilities and nonlinear dynamics of concentrated active suspensions. *Physics of Fluids* 25 (2013), 070607.
- [5] Aldo Frezzotti, Gian Pietro Ghiroldi, and Livio Gibelli. 2011. *Solving Kinetic Equations on GPU's*. Technical Report. DTIC Document.
- [6] Wei Ge, Ji Xu, Qingang Xiong, Xiaowei Wang, Feiguo Chen, Limin Wang, Chaofeng Hou, Ming Xu, and Jinghai Li. 2013. Multi-scale continuum-particle simulation on CPU-GPU hybrid supercomputer. In *GPU Solutions to Multi-scale Problems in Science and Engineering*. Springer, 143–161.
- [7] Mark Harris. 2007. Optimizing cuda. *SC07: High Performance Computing With CUDA* (2007).
- [8] Christel Hohenegger, Steve Cook, and Tamar Shinar. 2014. Dimensional Reduction of a Multiscale Continuum Model of Microtubule Gliding Assays. *SIAM J. Appl. Math.* 74, 5 (Jan. 2014), 1338–1353. <https://doi.org/10.1137/140961535>
- [9] Christel Hohenegger and Michael J. Shelley. 2011. Dynamics of complex biofluids. In *New Trends in the Physics and Mechanics of Biological Systems*, M. Ben Amar, A. Goriely, M. M. Müller, and L. F. Cugliandolo (Eds.). Oxford University Press, Chapter 3, 65–92.
- [10] Hanyu Jiang, Morisa Manzella, Luka Djapic, and Narayan Ganesan. 2016. Computational Framework for in-Silico Study of Virtual Cell Biology via Process Simulation and Multiscale Modeling. In *Proceedings of the 7th ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*. ACM, 384–393.
- [11] Roland Keunings. 2004. Micro-macro methods for the multiscale simulation of viscoelastic flow using molecular models of kinetic theory. *Rheology reviews* (2004), 67–98.
- [12] F. Nédélec. 2002. Computer simulations reveal motor properties generating stable antiparallel microtubule interactions. *The Journal of Cell Biology* 158, 6 (16 Sept. 2002), 1005–1015. <https://doi.org/10.1083/jcb.200202051>
- [13] François Nédélec and Dietrich Foethke. 2007. Collective Langevin dynamics of flexible cytoskeletal fibers. *New Journal of Physics* 9, 11 (2007), 427. <https://doi.org/10.1088/1367-2630/9/11/427>
- [14] François J. Nédélec and Thomas Surrey. 2001. Dynamics of microtubule aster formation by motor complexes. *Comptes-Rendus de l'Académie des Sciences Paris* 4, 2 (May 2001), 841–847.
- [15] Kyle E Niemeyer and Chih-Jen Sung. 2014. Recent progress and challenges in exploiting graphics processors in computational fluid dynamics. *The Journal of Supercomputing* 67, 2 (2014), 528–564.
- [16] Nvidia. 2011. Nvidia CUDA C programming guide. *Nvidia Corporation* (2011).
- [17] C. Peskin. 2003. The immersed boundary method. *Acta Numerica* 11 (July 2003), 479–517. <https://doi.org/10.1017/S0962492902000077>
- [18] D. Saintillan and M. J. Shelley. 2008. Instabilities, pattern formation, and mixing in active suspensions. *Physics of Fluids* 20, 12 (2008), 16. <https://doi.org/10.1063/1.3041776>
- [19] Yutaka Sumino, Ken H. Nagai, Yuji Shitaka, Dan Tanaka, Kenichi Yoshikawa, Hugues Chate, and Kazuhiro Oiwa. 2012. Large-scale vortex lattice emerging from collectively moving microtubules. *Nature* 483, 7390 (22 March 2012), 448–452. <https://doi.org/10.1038/nature10874>
- [20] Vasily Volkov. 2010. Better performance at lower occupancy. In *Proceedings of the GPU technology conference, GTC*, Vol. 10. San Jose, CA, 16.
- [21] E Westphal, SP Singh, C-C Huang, Gerhard Gompper, and Roland G Winkler. 2014. Multiparticle collision dynamics: GPU accelerated particle-based mesoscale hydrodynamic simulations. *Computer Physics Communications* 185, 2 (2014), 495–503.
- [22] Sergey Zabelok, Robert Arslanbekov, and Vladimir Kolobov. 2015. Adaptive kinetic-fluid solvers for heterogeneous computing architectures. *J. Comput. Phys.* 303 (2015), 455–469.