

Parallel Simulation for a Fish Schooling Model on a General-Purpose Graphics Processing Unit



Hong Li^{1,*}, Allison M. Kolpas^{2,†}, Linda Petzold^{1,‡}, Jeff Moehlis^{3,§}

¹ Department of Computer Science
University of California, Santa Barbara, CA 93106

² Department of Mathematics
University of California, Santa Barbara, CA 93106

³ Department of Mechanical Engineering
University of California, Santa Barbara, CA 93106

SUMMARY

We consider an individual-based model for fish schooling which incorporates a tendency for each fish to align its position and orientation with an appropriate average of its neighbors' positions and orientations, plus a tendency for each fish to avoid collisions. To accurately determine statistical properties of the collective motion of fish whose dynamics are described by such a model, many realizations are typically required. This carries a very high computational cost. The current generation of graphics processing units is well-suited to this task. We describe our implementation, and present computational experiments illustrating the power of this technology for this important and challenging class of problems.

KEY WORDS: GPU, Fish Schooling, Parallel Computing

1. INTRODUCTION

Many organisms display ordered collective motion [1], such as geese flying in a Chevron-shaped formation [7], wildebeests herding on the Serengeti plains of Africa [18], locusts swarming in sub-Saharan Africa [19], and fish schooling [16]. The number of individuals involved in such collective motion can be huge, from several hundred thousand wildebeests to millions of Atlantic cod off the Newfoundland coast. Despite these large numbers, the group can seem to move as a single

*E-mail: hongli@cs.ucsb.edu

†E-mail: allie@math.ucsb.edu

‡E-mail: petzold@enr.ucsb.edu

§E-mail: moehlis@enr.ucsb.edu



organism, with all individuals responding very rapidly to their neighbors to maintain the collective motion. Advantages of collective motion include a reduction in predation [12, 4], increased harvesting efficiency [10, 16], and improved aerodynamic efficiency [20, 11].

The mathematical study of biological collective motion has proceeded, broadly speaking, on two fronts. First, one can model each organism individually, with rules specifying its dynamics and interactions with other organisms. Such “individual-based models” can incorporate experimental observations of the behavior of the organisms, thereby giving biologically realistic models. On the other hand, one can define “continuum models” to describe the dynamics of the population; for example, one might study a model for the evolution of the density of the organisms. However, available continuum models are often qualitative caricatures that cannot capture all of the details which can be included in individual-based models, thereby compromising biological realism.

In this paper, we consider an individual-based model for fish schooling. This model incorporates a tendency for each fish to align its position and orientation with an appropriate average of its neighbors’ positions and orientations, plus a tendency for each fish to avoid collisions. Furthermore, randomness is included in the model to account for imperfections in the gathering of information and in acting on this information. To accurately determine statistical properties of the collective motion of fish whose dynamics are described by such an individual-based model, many realizations are typically required; see, e.g., [3]. This can entail a very large amount of computation. After describing the model in more detail, we will describe how a Graphics Processor Unit (GPU) can be used to very efficiently carry out parallel simulations of this model.

With the low cost and high performance processing capabilities of GPUs, computation on general purpose GPUs (GPGPU) [5] has become an active research field with a wide variety of scientific applications including fluid dynamics, cellular automata, particle systems, neural networks, and computational geometry [5, 15, 6]. Previous generation GPUs have required computations to be recast into a graphics application programming interface (API) such as OpenGL, which has made programming GPUs for non-graphics applications a significant challenge. NVIDIA’s Compute Unified Device Architecture (CUDA) [13] is a new technology which directly enables implementation of parallel programs in the C language using an API designed for general-purpose computation. In this paper, we will show how Single Instruction Multiple Data (SIMD) computation as well as parallel processing within a single realization can be implemented on a CUDA-enabled GPU to efficiently perform ensemble simulations of an individual-based fish schooling model.

2. FISH SCHOOLING MODEL

We consider a two-dimensional individual-based model for schooling with local behavioral interactions. This model is similar to that of [2], but without an informed leader, and with different weights of orientation and attraction response. Groups are composed of N individuals with positions $p_i(t)$, unit directions $\hat{v}_i(t)$, constant speed s , and maximum turning rate θ . At every time step of size τ , individuals simultaneously determine a new direction of travel by considering neighbors within two behavioral zones. The first zone, often referred to as the “zone of repulsion” [3], is represented by a circle of radius r_r about the individual. Individuals repel from neighbors that are within this zone, which typically has a radius of one body length. The second zone, a “zone of orientation and attraction”, is represented by an annulus of inner radius r_r and outer radius r_p about the individual. This zone also



includes a blind area, defined as a circular sector with central angle $(2\pi - \eta)$, for which neighbors within the zone are undetectable. Individuals orient with and are attracted towards neighbors within this zone.

These zones are used to define behavioral rules of motion. First, if individual i finds agents within its zone of repulsion, then it orients its direction away from the average relative directions of those within its zone of repulsion. Its desired direction of travel in the next time step is given by

$$v_i(t + \tau) = - \sum_{j \neq i} \frac{p_j(t) - p_i(t)}{|p_j(t) - p_i(t)|}. \quad (1)$$

This vector is normalized as $\hat{v}_i(t + \tau) = \frac{v_i(t + \tau)}{|v_i(t + \tau)|}$, assuming $v_i(t + \tau) \neq 0$. In the case that $v_i(t + \tau) = 0$, agent i maintains its previous direction of travel as its desired direction of travel, giving $\hat{v}_i(t + \tau) = \hat{v}_i(t)$.

If agents are not found within individual i 's zone of repulsion, then it will align with (by averaging the directions of travel of itself and its neighbors) and feel an attraction towards (by orienting itself towards the average relative directions of) agents within the zone of orientation and attraction. The desired direction of agent i is given by the weighted sum of these two terms:

$$v_i(t + \tau) = \omega_a \frac{\sum_{j \neq i} \frac{p_j(t) - p_i(t)}{|p_j(t) - p_i(t)|}}{\left| \sum_{j \neq i} \frac{p_j(t) - p_i(t)}{|p_j(t) - p_i(t)|} \right|} + \omega_o \frac{\sum_j \hat{v}_j(t)}{\left| \sum_j \hat{v}_j(t) \right|}, \quad (2)$$

where ω_a and ω_o are the weights of attraction and orientation, respectively. This vector is then normalized as $\hat{v}_i(t + \tau) = \frac{v_i(t + \tau)}{|v_i(t + \tau)|}$, assuming $v_i(t + \tau) \neq 0$. As before, if $v_i(t + \tau) = 0$, then agent i maintains its previous direction of travel.

We denote $r = \omega_o/\omega_a$ as the ratio of orientation and attraction tendencies. When $r = 0$ ($\omega_o = 0$), individuals have no desire to orient with their neighbors. As r approaches 1, individuals balance their orientation and attraction preferences. For $r > 1$, individuals are more interested in orientation with their neighbors than attraction towards them.

Stochastic effects are incorporated into the model by rotating agent i 's desired direction $\hat{v}_i(t + \tau)$ by an angle drawn from a circularly wrapped normal distribution with mean 0 and standard deviation σ . Also, since individuals can only turn $\theta\tau$ radians in one time step, if the angle between $\hat{v}_i(t)$ and $\hat{v}_i(t + \tau)$ is greater than $\theta\tau$, individuals do not achieve their desired direction, and instead rotate $\theta\tau$ towards it. Finally, each agent's position is updated simultaneously as

$$p_i(t + \tau) = p_i(t) + s\hat{v}_i(t + \tau)\tau. \quad (3)$$




Figure 1. CPU vs. GPU architecture. There are more transistors devoted to data processing rather than to data caching and flow control for the GPU, compared with the CPU.

3. THE GRAPHICS PROCESSOR UNIT - A NEW DATA PARALLEL COMPUTING DEVICE

3.1. Modern Graphics Processor Unit

The general-purpose Graphics Processing Unit (GPU) is becoming a viable option for many parallel programming applications. The GPU has a highly parallel structure with high memory bandwidth and more transistors devoted to data processing than to data caching and flow control, compared with a CPU architecture, as shown in Figure 1 [13]. Problems that can be implemented with stream processing and that use limited memory are well-suited to the GPU architecture. Single Instruction Multiple Data (SIMD) computation, which involves a large number of totally independent records being processed by the same sequence of operations simultaneously, is ideal for GPU application. We will show how simulation of the fish schooling model can be structured to fit within the constraints of the GPU architecture.

3.2. NVIDIA 8 Series GeForce-based GPU Architecture

The NVIDIA 8800 GTX chip, released at the end of 2006, has 768MB RAM and 681 million transistors on a 480mm^2 surface area. There are 128 stream processors on a GeForce 8800 GTX chip, divided into 16 clusters of multiprocessors with 8 streaming processors in each multiprocessor as shown in Figure 2 [13]. The 8 processors in each multiprocessor share 16 KB shared memory which brings data closer to the Arithmetic Logic Unit (ALU). The peak computation rate accessible for general-purpose application is $(16 \text{ multiprocessors} \times 8 \text{ processors / multiprocessor}) \times (2 \text{ flops / MAD}^*) \times (1 \text{ MAD / processor-cycle}) \times 1.35 \text{ GHz} = 345.6 \text{ GFLOP/s}$, since the processors are clocked at 1.35 GHz with dual processing of scalar operations supported. The maximum observed bandwidth between system and device memory is about 2GB/second.

*A MAD is a multiply-add.

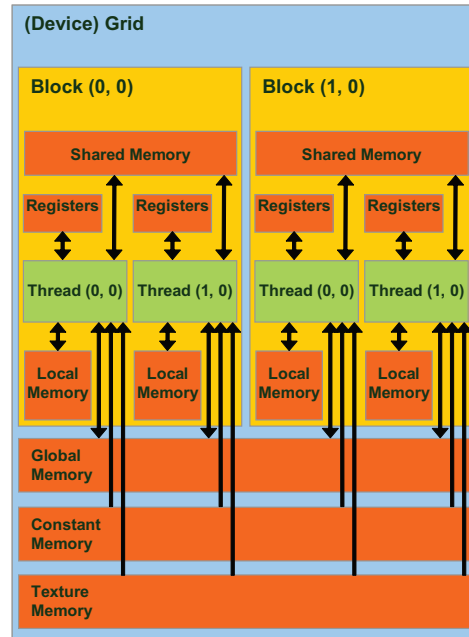


Figure 2. Hardware Model : a set of SIMD multiprocessors with on-chip shared memory.

Unfortunately, current GPU chips support only single-precision floating point, while current CPUs support 64-bit technology. Additionally, there is only 16K of fast read and write on-chip memory shared between the eight processors, on each multiprocessor.

3.3. CUDA: Compute Unified Device Architecture

The Compute Unified Device Architecture (CUDA) Software Development Kit (SDK) supplies general purpose functionality for non-graphics applications on the NVIDIA GPU. The CUDA provides an essential high-level development environment with standard C language, which results in a minimal learning curve for beginners to access the low-level hardware. For development flexibility, the CUDA provides both scatter and gather memory operations. It also supports a fast read and write shared memory [13].

Following the data-parallel model, the structure of CUDA computation allows each of the processors to execute the same instruction sequence on different sets of data in parallel. The data can be broken into a 1D or 2D grid of thread blocks. Each block can be specified as a 1D, 2D, or 3D array of threads which collaborate via the shared memory. Up to 512 threads can be active.

Currently, the CPU and GPU cannot run in parallel. Also, it is not possible to execute multiple kernels at the same time through the CUDA, or to download data and run a kernel in parallel. Users



can branch in the kernel based on the thread id to achieve multiple tasks on a single kernel, but the simulation will be slowed down with such branches. Individual GPU program launches can run at most 5 seconds on a GPU with a display attached.

4. PARALLEL SIMULATION

To accurately determine statistical properties of the collective motion of fish, many realizations of the fish schooling model are typically required for a given set of parameters. This can be very computationally intensive. There are essentially two ways to improve the performance: parallelize the simulation across the realizations, and parallelize the simulation within one realization. We do both.

We also tried simulating our fish model on a cluster. We got a linear performance improvement for parallelization across realizations, but for parallelization within one realization, the message passing overhead is a substantial burden. One might also consider using a shared memory computer to improve the performance of parallelization within one realization, but for multiple realizations where many threads must access the shared resources there are likely to be serious problems with contention. The best part of the GPU is that it is *ours*. It resides on our workstation. We don't have to share it, schedule it or find machine room space or a system administrator for it.

Migrating a C code from a workstation to the GPU is not a big job if you are familiar with programming multiple threads. Starting from a well-written sequential C code on our workstation, it took one of us about one week to port it to the GPU.

4.1. Random Number Generation

To generate the initial conditions and to add noise to our calculations at each time step, we need many uniformly distributed random numbers. So that our statistical results can be trusted, it is particularly important that these random numbers are independent. Generating independent sequences of random numbers is one of the most important issues of implementing simulations for ensembles of fish schooling models in parallel.

Originally we considered pre-generating a sequence of uniform random numbers in the CPU. Since the CPU and GPU can't run in parallel, we can pre-generate a huge number of random numbers and store them in the shared memory, and swap back to the CPU to generate more when they are used up. Alternatively, we could pre-generate a huge number of random numbers and store them in the global memory. These methods spend too much time on data access. Furthermore, the Scalable Parallel Random Number Generators Library (SPRNG) [8, 9], which might have been an ideal choice because of its excellent statistical properties, cannot be moved to the GPU efficiently due to its complicated data structure. The only solution appears to be to employ a simple random number generator on the GPU. Experts suggest using a mature random number generator instead of inventing a new one, since it requires great care and extensive testing to evaluate a random number generator [17]. Thus we chose the Mersenne Twister [14] in our application, which has passed many statistical randomness tests including the stringent Diehard tests.

In our implementation, we modified Eric Mills' multithreaded C implementation [14]. Since many random numbers are required by this model, we use the shared memory for random number generation to minimize the data launching and accessing time.



4.2. Parallelizing Across the Realizations

Parallelizing the independent realizations is an effective way to improve the performance. Ensembles of fish schooling runs are very well-suited for implementation on the GPU through the CUDA.

The whole fish schooling simulation can be put into a single kernel running in parallel on a large set of system state vectors $Px_i^j(t), Py_i^j(t), Vx_i^j(t), Vy_i^j(t)$, respectively representing the x and y components of the positions of the fish and the x and y components of the velocities of the fish, where i is id of the fish, j is the block id, and t represents time. The initial conditions $Px_i^j(0), Py_i^j(0), Vx_i^j(0), Vy_i^j(0)$ are randomly generated by different processors within each multiprocessor. The large set of final state vectors $Px_i^j(t_f), Py_i^j(t_f), Vx_i^j(t_f), Vy_i^j(t_f)$ will contain the desired results. We minimize the transfer between the host and the device by using an intermediate data structure on the device, and batch a few small transfers into a big transfer to reduce the overhead for each transfer.

4.3. Parallelizing Within One Realization

Before discussing parallelism within one realization, we outline the sequential algorithm. Individual fish are initialized in a bounded region with randomized positions and directions of travel. Then the simulation is executed for $t_f = 3000$ steps to reach a steady state. At each step, we calculate the influence on each fish from all the other fish. To calculate the influence on a fish, we compute the distance from this “goal fish” to all the other fish and record the influence coming from the different zones in a few variables. Next, we compute the net influences for this “goal fish”, including the noise term, based on the above data, and save them to an influence array. After computing all the influences for all the fish, we update the position and direction of each fish based on the influence array of this step.

For the GPU, there are some restrictions on the number of threads in each block and the total number of fish stored in the shared memory. The device has very limited shared memory in each multiprocessor but relatively large global memory. The global memory adjacent to the GPU chip has much higher latency and lower bandwidth than the on-chip shared memory: it takes about 400-600 clock cycles to read or write the global memory vs. 4 clock cycles to access the shared memory. To effectively use the GPU, our simulation makes as much use of on-chip shared memory as possible. Although it is better to have at least 128 threads for the best efficiency, we have to take the limited shared memory size into account.

The main goal of parallelizing within one realization is to decompose the problem domain into smaller subdomains. In our implementation, within each block j , the system state vector $Px_i^j, Py_i^j, Vx_i^j, Vy_i^j$ is divided among threads. Assume that we have N fish and we use n threads to do the simulation. Then we need order of $m = N/n$ time to calculate the influences on each of the fish. In the parallel implementation, we initialize all the fish with random positions and directions. At each step we load n fish by loading one fish in each thread. The system state vectors loaded to shared memory at this stage are the “goal fish”. Each thread holds only one “goal fish” at a time and calculates the influences on this “goal fish”. Thus during one step, each thread processes one m -element subvector. To calculate the distance from the “goal fish” to all the other fish, we need the position and direction of all the other fish. By using a temporary array of size n in shared memory, we load n fish at a time, with each thread loading one fish. The program will keep loading until all the desired data have

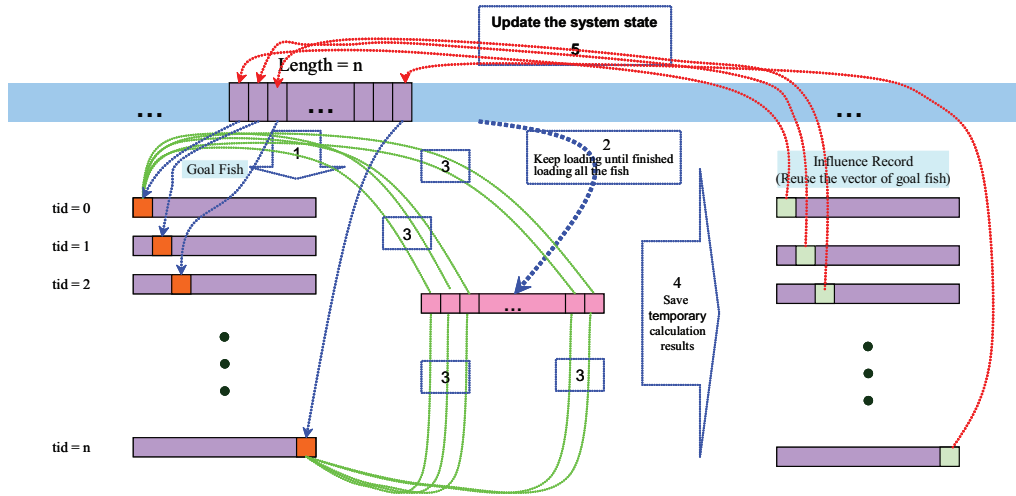


Figure 3. Memory layout of the fish schooling model. There are n threads to simulate one realization with thread id tid to identify the thread. In process 1, each thread loads one fish as its “goal fish” from the device memory to shared memory; In process 2, each thread loads fish to another array in shared memory; In process 3, each thread uses the full data from process 2 to compute the influences on its own “goal fish”; Processes 2 and 3 continue until all of the influences to the “goal fish” have been computed; In process 4, each thread saves its influence to the influence record array (this reuses the previous array to save shared memory); In process 5, each thread updates its “goal fish” to the system state vector on the device.

been loaded and used for calculation for the goal fish. Then, each thread will compute the influence information for the goal fish of this thread. At the end of the simulation step, each thread writes the results into its own influence subvector. After all the influence calculations are finished, each thread updates its system state subvector by its influence subvector. The memory layout of the whole process is shown in Figure 3. Our implementation does not make use of a concurrent update. During each calculation, each thread reads the full system state vector. The concurrent reads to the same data are supported by the hardware.

Within each realization, most calculations use the shared memory, for which the memory access latency is small. In addition to the parallelization within one realization, we have many blocks doing independent realizations. Whenever a thread is waiting to access the device memory, another thread will be running on the ALUs. Hence the memory access latency is not a problem.

5. RESULTS

Our simulations were run on the NVIDIA GeForce 8800GTX installed on a host workstation with Intel Pentium 3.00GHz CPU and 3.50GB of RAM with physical Address Extension.



To begin a simulation, individuals are placed in a bounded region (so that each agent initially interacts with at least one other agent), with random positions and directions of travel. The parameters were fixed to be $r_r = 1$, $r_p = 7$, $\eta = 350^\circ (\approx 6.1 \text{ radians})$, $s = 1$, $\tau = 0.2$, $\sigma = 0.01$, and $\theta = 115^\circ (\approx 2 \text{ radians})$. Simulations were run for approximately 3000 steps until they reached a steady state. Groups of size $N = 100$ were explored for a range of ratios r of attraction to orientation weightings.

Two observables were used to measure the structure of the schools: elongation and polarization [2]. Elongation measures the ratio of the length of the axis aligned with group motion to the axis perpendicular to group motion, and polarization

$$P(t) = \frac{1}{N} \left| \sum_{i=1}^N v_i(t) \right|,$$

measures the degree of group alignment. Both quantities take values between 0 and 1. To obtain statistics regarding the group structure for a given size N and ratio r , 1120 realizations (with different initial conditions) were run. The average group elongation and polarization were recorded, as well as the probability of group fragmentation.

We were able to nicely resolve the group statistics as a function of the ratio of attraction to alignment. We find that for r close to zero, groups exhibit swarm behavior, with small polarization and elongation near 1. As r is increased, groups become increasingly more aligned, forming elongated dynamically parallel groups and then highly parallel groups as r is further increased; see Figure 4. Thus, by changing the proportion of attraction to alignment tendencies, groups can shift between different collective behaviors. We also see that the probability of fragmentation first increases as the group elongation increases (from $1 < r < 2.5$), then decreases as elongation decreases (from $3 < r < 8$), and then increases again as schools become more highly aligned; see Figure 5.

The simulation performance in generating these results was extraordinary. We simulated 1120 realizations of the code with $N = 100$ individuals. The timing results for different parameters are shown in Table I. The simulation time of the GPU-based code includes the time of loading data to the GPU and copying the data back to the host. The parallel (GPU) simulation is about 230 – 240 times faster than the corresponding sequential simulation on the host workstation.

The code on the host workstation was carefully written to reduce memory accesses and combine loops, and was the starting point of the GPU code. SSE (Streaming SIMD Extensions) was not enabled on the workstation operating system, in the way that most practitioners (including us!) would normally use a workstation. The workstation computation was performed in single precision. Even so, one might reasonably ask why is the GPU making so much better use of its CPU resources? The answer is multithreading. In our simulation code, there are many branches and memory access operations, besides the floating point operations. On the GPU, the scheduler can arrange for other threads to run while the current threads are accessing the global memory.

The limitation of single-precision floating-point on the GPU seems to have no noticeable effect on the results, for this computation. Schools generated in double precision on the host workstation were found to be statistically indistinguishable from those obtained in single precision.

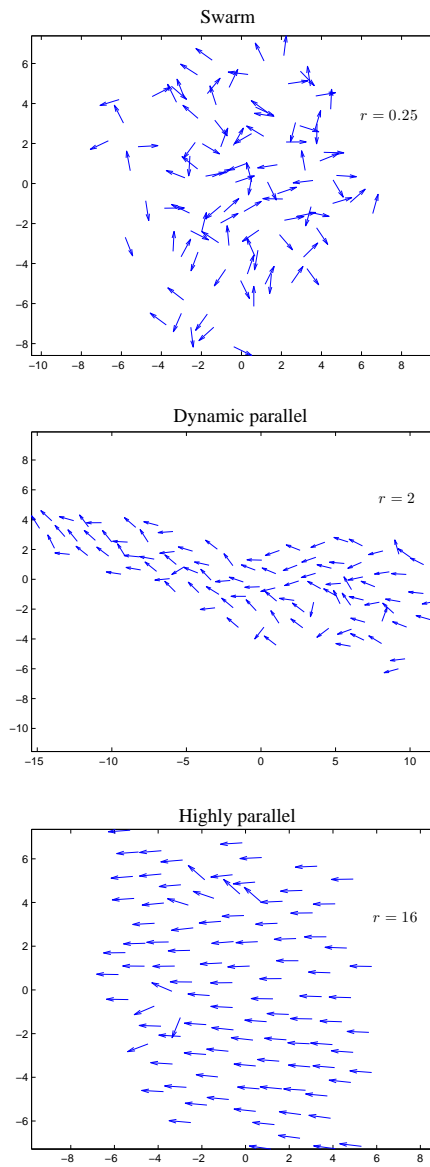


Figure 4. Swarm ($r = 0.25$), dynamic parallel ($r = 2$), and highly parallel ($r = 16$) collective motion for $N = 100$ member schools.




Figure 5. Statistics of group elongation (only including non-fragmented schools), group polarization (only including non-fragmented schools), and probability of group fragmentation, for $N = 100$ member schools. To obtain these statistics, for each value of r , 1120 simulations with different initial conditions were run for 3000 time steps, each using the GPU.

6. CONCLUSIONS

We considered an individual-based model for fish schooling which incorporates a tendency for each fish to align its position and orientation with an appropriate average of its neighbors' positions and orientations, plus a tendency for each fish to avoid collisions. Randomness is included in the model to account for imperfections in the gathering of information and in acting on this information. For the problem of calculating collective motion statistical measures from many realizations, we observed speedups of 230 – 240 times for our parallelized code running on a GPU over the corresponding sequential simulation on the host workstation. With this impressive performance improvement, in one day we can generate data which would require more than six months of computation with the sequential code.

In the future we hope to simulate larger fish schools using clusters of GPUs. Our computation should scale well for parallelization across the simulations. Effective parallelization within one realization



Table I. Performance Comparison for different parameters

Δr_o	Sequential Simulation Time (ms)	Parallel Simulation Time (ms)	Speedup
0.03125	7924067.2	33750.6	234.8
0.0625	7925786.6	33767.7	234.7
0.125	7908017.5	33750.6	234.3
0.25	7938803.4	33790.0	235.0
0.5	7934238.6	33774.8	235.0
1	7902735.7	33855.3	233.4
2	7799067.5	33755.0	231.0
4	8009521.9	33859.2	236.6
8	8091104.2	33681.4	240.2
16	8091787.8	33706.8	240.1

will depend, of course, on the communication speed of the cluster. The usefulness of individual-based models of biological systems, which can incorporate detailed experimental observations of the behavior of the organisms, is limited by the ability to simulate them in a reasonable amount of time. By exploiting the power of general-purpose GPUs as illustrated in this paper, we expect that it will be possible to simulate more detailed models with more agents for longer times and over more realizations. This GPU computational paradigm has the potential to revolutionize our ability to understand how collective behavior at the macroscopic, population-level arises from the interaction of individuals with each other and their environment.

7. ACKNOWLEDGMENTS

We thank Iain Couzin for helpful discussions related to the model, and Brent Oster for introducing us to the NVIDIA GPU.

This work was supported in part by the U.S. Department of Energy under DOE award No. DE-FG02-04ER25621, by NIH Grant EB007511, by the Institute for Collaborative Biotechnologies through grant DAAD19-03-D004 from the U.S. Army Research Office, and by National Science Foundation Grant NSF-0434328. J.M. also was supported by an Alfred P. Sloan Research Fellowship in Mathematics.

REFERENCES

1. S. Camazine, J. L. Deneubourg, N. R. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau. *Self-Organization in Biological Systems*. Princeton University Press, Princeton, 2003.
2. I. D. Couzin, J. Krause, N. R. Franks, and S. A. Levin. Effective leadership and decision making in animal groups on the move. *Nature*, 433:513–516, 2005.
3. I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and N. R. Franks. Collective memory and spatial sorting in animal groups. *J. theor. Biol.*, 218:1–11, 2002.



4. P. M. Driver and D. A. Humphries. *Protean Behavior: The Biology of Unpredictability*. Oxford University Press, Oxford, 1988.
5. GPGPU-Home. GPGPU homepage, 2007. <http://www.gpgpu.org/>.
6. H. Li and L. Petzold. Stochastic simulation of biochemical systems on the graphics processing unit. Technical report, Department of Computer Science, University of California, Santa Barbara, 2007. Submitted.
7. P. B. S. Lissaman and C. A. Shollenberger. Formation flight of birds. *Science*, 168:1003–1005, 1970.
8. M. Mascagni. SPRNG: A scalable library for pseudorandom number generation. In *Proceedings of the Ninth SIAM Conference on Parallel Processing for Scientific Computing*, San Antonio, Texas, 1999.
9. M. Mascagni and A. Srinivasan. SPRNG: A scalable library for pseudorandom number generation. In *ACM Transactions on Mathematical Software*, volume 26, pages 436–461, 2000.
10. D. R. Martinez and E. Klinghammer. The behavior of the whale *orcinus orca*: a review of the literature. *Z. Tierpsychol.*, 27:828–839, 1970.
11. R. M. May. Flight formations in geese and other birds. *Nature*, 282:778–780, 1979.
12. S. R. S. J. Neill and J. M. Cullen. Experiments on whether schooling by their prey affects the hunting behavior of cephalopods and fish predators. *J. Zoology*, 172:549–569, 1974.
13. NVIDIA Corporation. NVIDIA CUDA Compute Unified Device Architecture Programming Guide, 2007. <http://developer.download.nvidia.com>.
14. NVIDIA Forums members. NVIDIA forums, 2007. <http://forums.nvidia.com>.
15. John D. Owens, David Luebke, Naga Govindaraju, Mark Harris, Jens Krger, Aaron E. Lefohn, and Timothy J. Purcell. A survey of general-purpose computation on graphics hardware. In *Eurographics 2005, State of the Art Reports*, pages 21–51, August 2005.
16. B. L. Partridge. The structure and function of fish schools. *Sci. Am.*, 245:90–99, 1982.
17. Richard P. Brent. Uniform Random Number Generators for Vector and Parallel Computers. *Report TR-CS-92-02*, 1992.
18. A. R. E. Sinclair and M. Norton-Griffiths. *Serengeti: Dynamics of an Ecosystem*. University of Chicago, Chicago, 1979.
19. B. P. Uvarov. *Grasshoppers and Locusts*. Imperial Bureau of Entomology, London, 1928.
20. D. Wiehs. Hydrodynamics of fish schooling. *Nature*, 241:290–291, 1973.