

ECE 498AL

Lecture 21-22: Application Performance Case Studies: Molecular Visualization and Analysis

Guest Lecture by John Stone

Theoretical and Computational Biophysics Group



NIH Resource for Macromolecular Modeling and Bioinformatics

Beckman Institute for Advanced Science and Technology

<http://www.ks.uiuc.edu/Research/gpu/>


VMD – “Visual Molecular Dynamics”

- Visualization and analysis of molecular dynamics simulations, sequence data, volumetric data, quantum chemistry simulations, particle systems, ...
- User extensible with scripting and plugins
- <http://www.ks.uiuc.edu/Research/vmd/>



Experiences Integrating CUDA Kernels Into VMD

- VMD: molecular visualization and analysis
- State-of-the-art simulations require more viz/analysis power than ever before
- For some algorithms, CUDA can bring what was previously supercomputer class performance to an appropriately equipped desktop workstation




Ribosome: 260,790 atoms
before adding solvent/ions


Range of VMD Usage Scenarios

- Users run VMD on a diverse range of hardware: laptops, desktops, clusters, and supercomputers
- Typically used as a desktop science application, for interactive 3D molecular graphics and analysis
- Can also be run in pure text mode for numerically intensive analysis tasks, batch mode movie rendering, etc...
- GPU acceleration provides an opportunity to make some **slow, or batch** calculations capable of being run **interactively, or on-demand...**


CUDA Acceleration in VMD





Electrostatic field calculation, ion placement



Molecular orbital calculation and display



Imaging of gas migration pathways in proteins with implicit ligand sampling




Electrostatic Potential Maps

- Electrostatic potentials evaluated on 3-D lattice:

$$V_i = \sum_j \frac{q_j}{4\pi\epsilon_0|\mathbf{r}_j - \mathbf{r}_i|}$$

- Applications include:
 - Ion placement for structure building
 - Time-averaged potentials for simulation
 - Visualization and analysis



Overview of Direct Coulomb Summation (DCS) Algorithm


- One of several ways to compute the electrostatic potentials on a grid, ideally suited for the GPU
- Methods such as multilevel summation can achieve much higher performance at the cost of some numerical accuracy
- Begin with DCS for computing electrostatic maps:
 - conceptually simple algorithm well suited to the GPU
 - easy to fully explore
 - requires very little background knowledge, unlike other methods
- DCS: for each lattice point, sum potential contributions for all atoms in the simulated structure:
$$\text{potential}[j] += \text{atom}[i].\text{charge} / r_{ij}$$

Direct Coulomb Summation (DCS)

Algorithm Detail

- Each lattice point accumulates electrostatic potential contribution from all atoms:

$$\text{potential}[j] += \text{atom}[i].\text{charge} / r_{ij}$$



DCS Computational Considerations

- Attributes of DCS algorithm for computing electrostatic maps:
 - Highly data parallel
 - Starting point for more sophisticated algorithms
 - Single-precision FP arithmetic is adequate for intended uses
 - Numerical accuracy can be further improved by compensated summation, spatially ordered summation groupings, or with the use of double-precision accumulation
 - Interesting test case since potential maps are useful for various visualization and analysis tasks
- Forms a template for related spatially evaluated function summation algorithms in CUDA

Single Slice DCS: Simple (Slow) C Version

```
void cenergy(float *energygrid, dim3 grid, float gridspace, float z, const float *atoms,
            int numatoms) {
    int i,j,n;
    int atomarrdim = numatoms * 4;
    for (j=0; j<grid.y; j++) {
        float y = gridspace * (float) j;
        for (i=0; i<grid.x; i++) {
            float x = gridspace * (float) i;
            float energy = 0.0f;
            for (n=0; n<atomarrdim; n+=4) { // calculate potential contribution of each atom
                float dx = x - atoms[n ];
                float dy = y - atoms[n+1];
                float dz = z - atoms[n+2];
                energy += atoms[n+3] / sqrtf(dx*dx + dy*dy + dz*dz);
            }
            energygrid[grid.x*grid.y*k + grid.x*j + i] = energy;
        }
    }
}
```

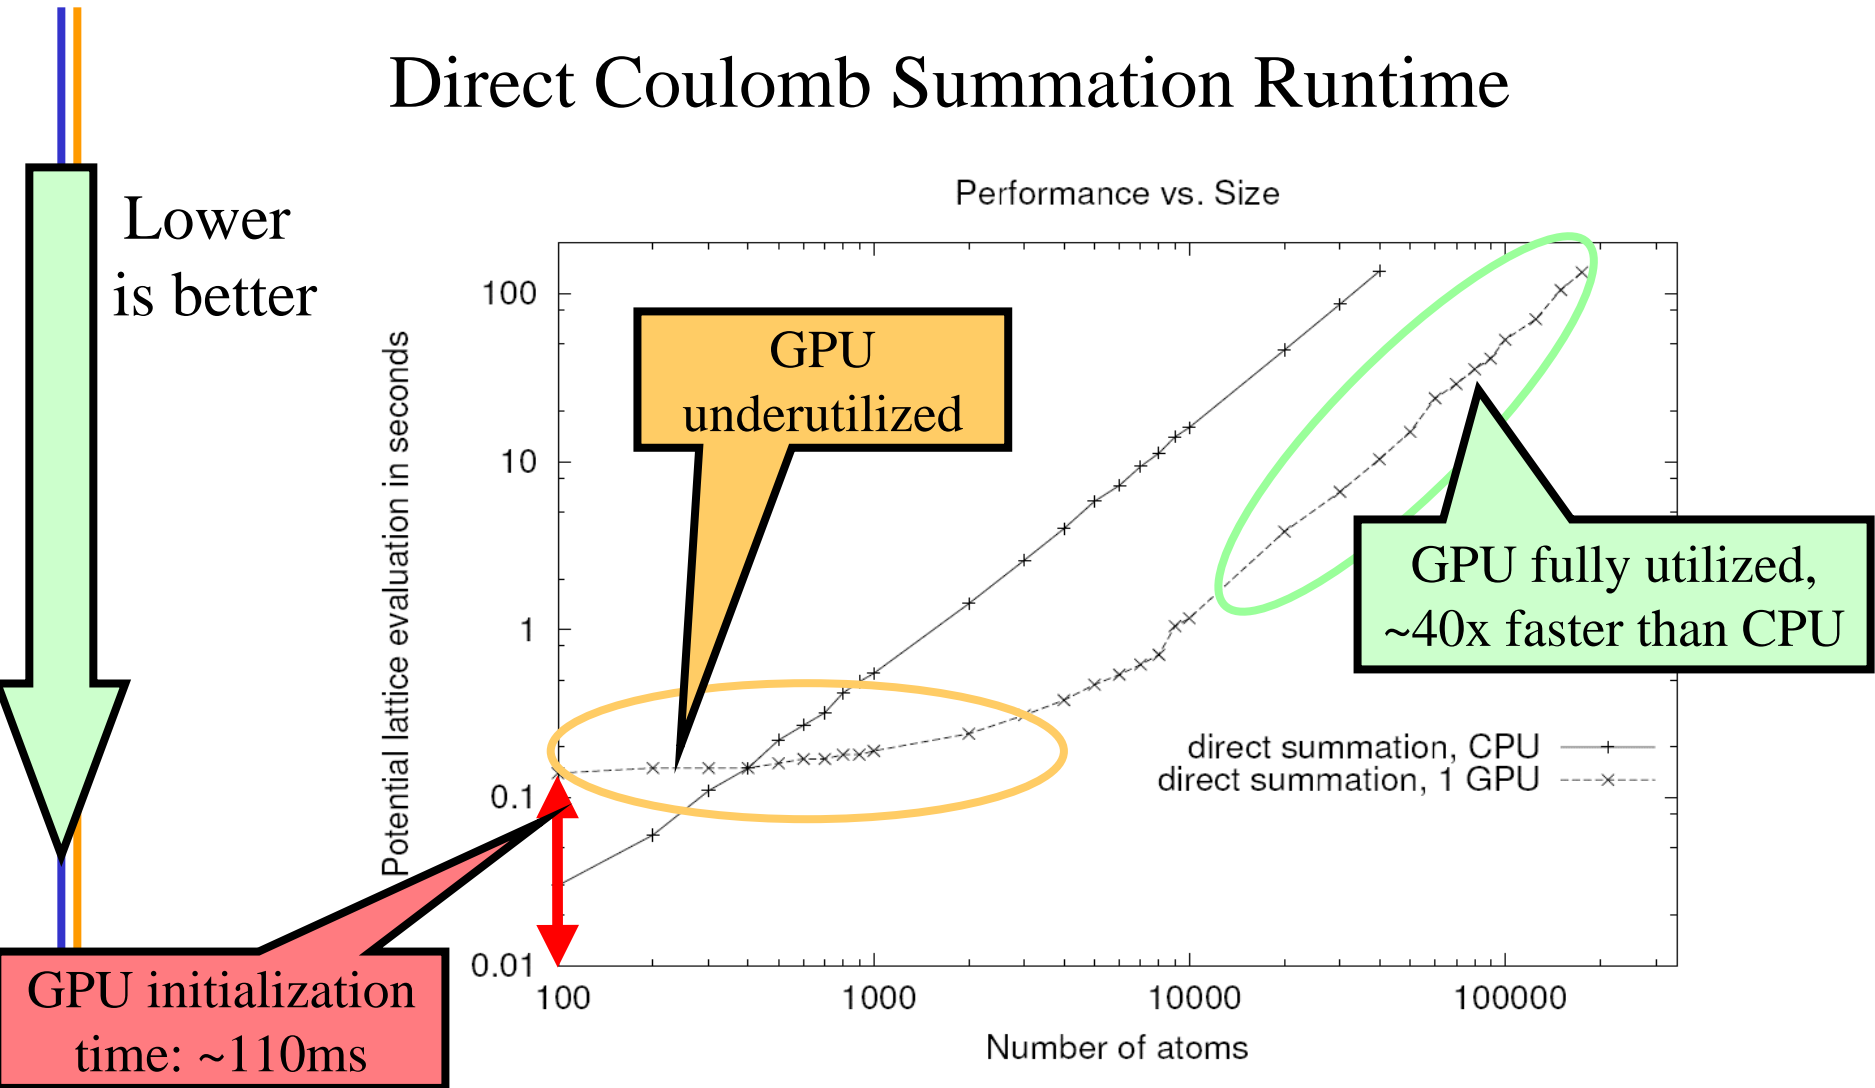
DCS Algorithm Design Observations

- Electrostatic maps used for ion placement require evaluation of ~20 potential lattice points per atom for a typical biological structure
- Atom list has the smallest memory footprint, best choice for the inner loop (both CPU and GPU)
- Lattice point coordinates are computed on-the-fly
- Atom coordinates are made relative to the origin of the potential map, eliminating redundant arithmetic
- Arithmetic can be significantly reduced by precalculating and reusing distance components, e.g. create a new array containing X , Q , and $dy^2 + dz^2$, updated on-the-fly for each row (CPU)
- Vectorized CPU versions benefit greatly from SSE instructions

An Approach to Writing CUDA Kernels

- Find an algorithm that can expose substantial parallelism, we'll ultimately need thousands of independent threads...
- Identify appropriate GPU memory or texture subsystems used to store data used by kernel
- Are there trade-offs that can be made to exchange computation for more parallelism?
 - Though counterintuitive, past successes resulted from this strategy
 - “Brute force” methods that expose significant parallelism do surprisingly well on current GPUs
- Analyze the real-world use case for the problem and select the kernel for the problem sizes that will be heavily used

Direct Coulomb Summation Runtime



Accelerating molecular modeling applications with graphics processors.
J. Stone, J. Phillips, P. Freddolino, D. Hardy, L. Trabuco, K. Schulten.
J. Comp. Chem., 28:2618-2640, 2007.

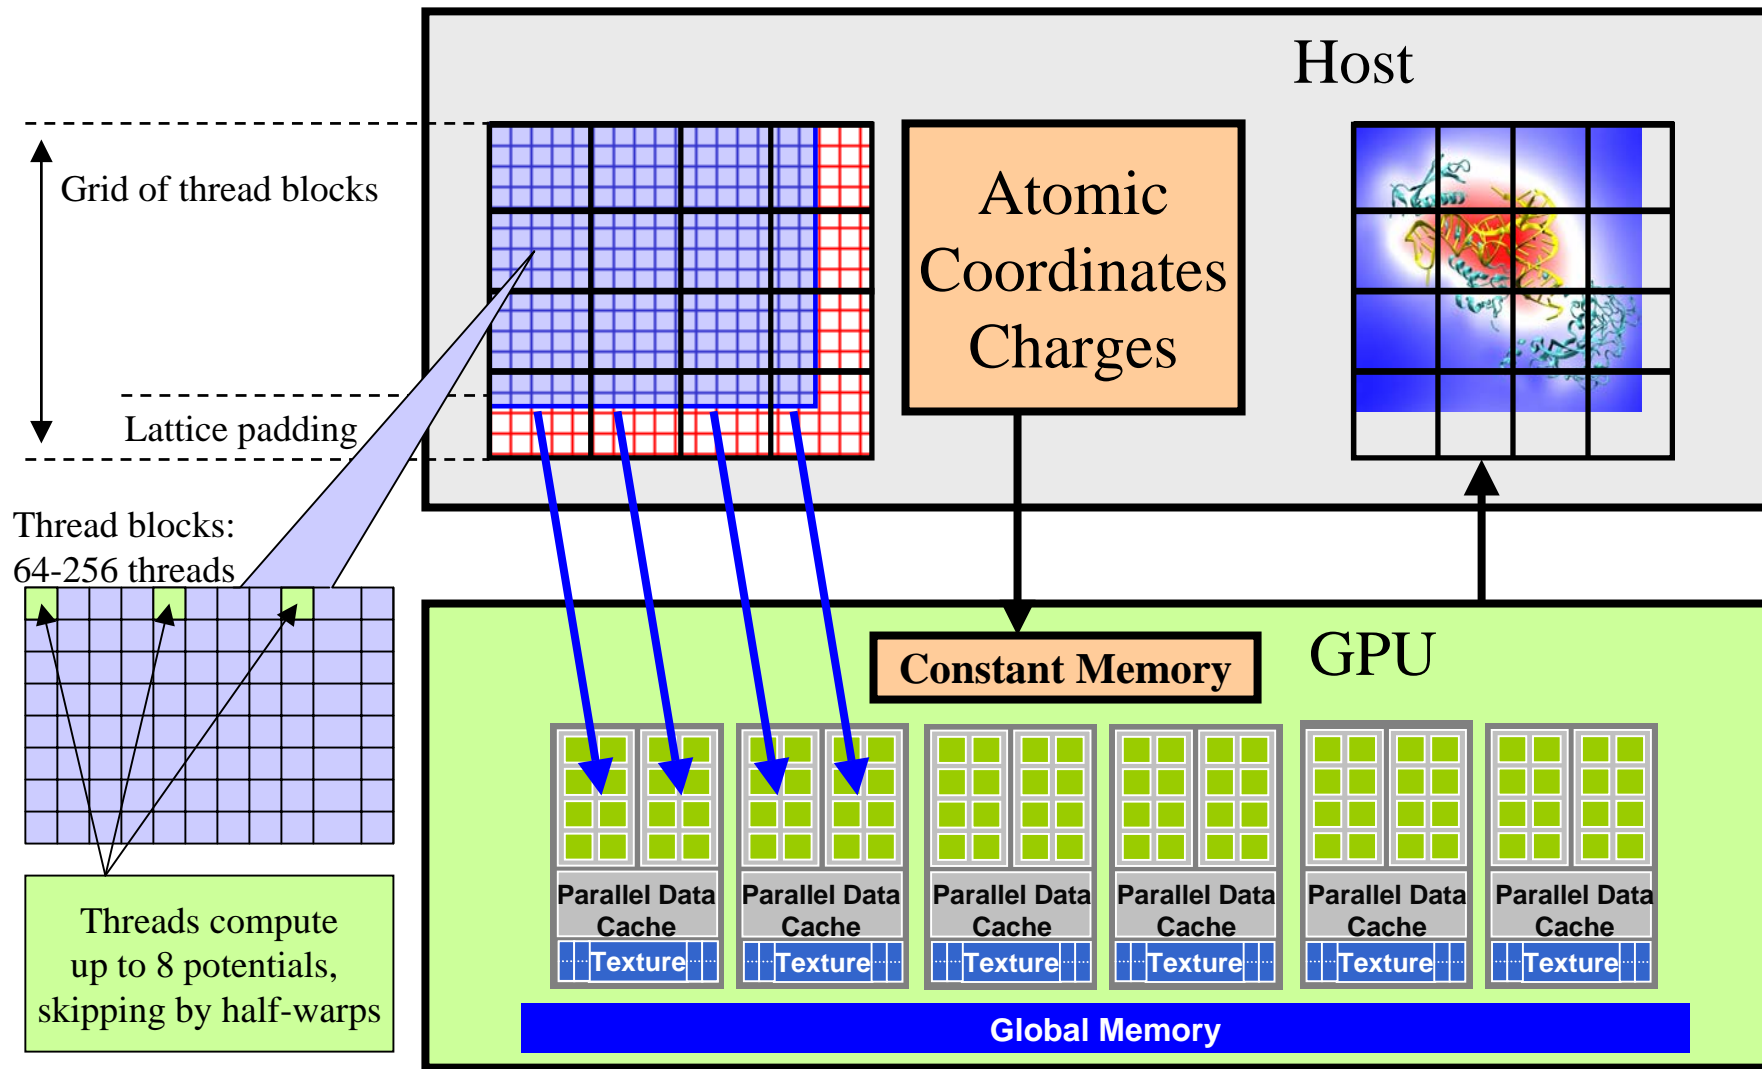
DCS Observations for GPU Implementation

- Naive implementation has a low ratio of FP arithmetic operations to memory transactions (at least for a GPU...)
- The innermost loop will consume operands VERY quickly
- Since atoms are read-only, they are ideal candidates for texture memory or constant memory
- GPU implementations must access constant memory efficiently, avoid shared memory bank conflicts, coalesce global memory accesses, and overlap arithmetic with global memory latency
- Map is padded out to a multiple of the thread block size:
 - Eliminates conditional handling at the edges, thus also eliminating the possibility of branch divergence
 - Assists with memory coalescing

CUDA DCS Implementation Overview

- Allocate and initialize potential map memory on host CPU
- Allocate potential map slice buffer on GPU
- Preprocess atom coordinates and charges
- Loop over slices:
 - Copy slice from host to GPU
 - Loop over groups of atoms until done:
 - Copy atom data to GPU
 - Run CUDA Kernel on atoms and slice resident on GPU accumulating new potential contributions into slice
 - Copy slice from GPU back to host
- Free resources

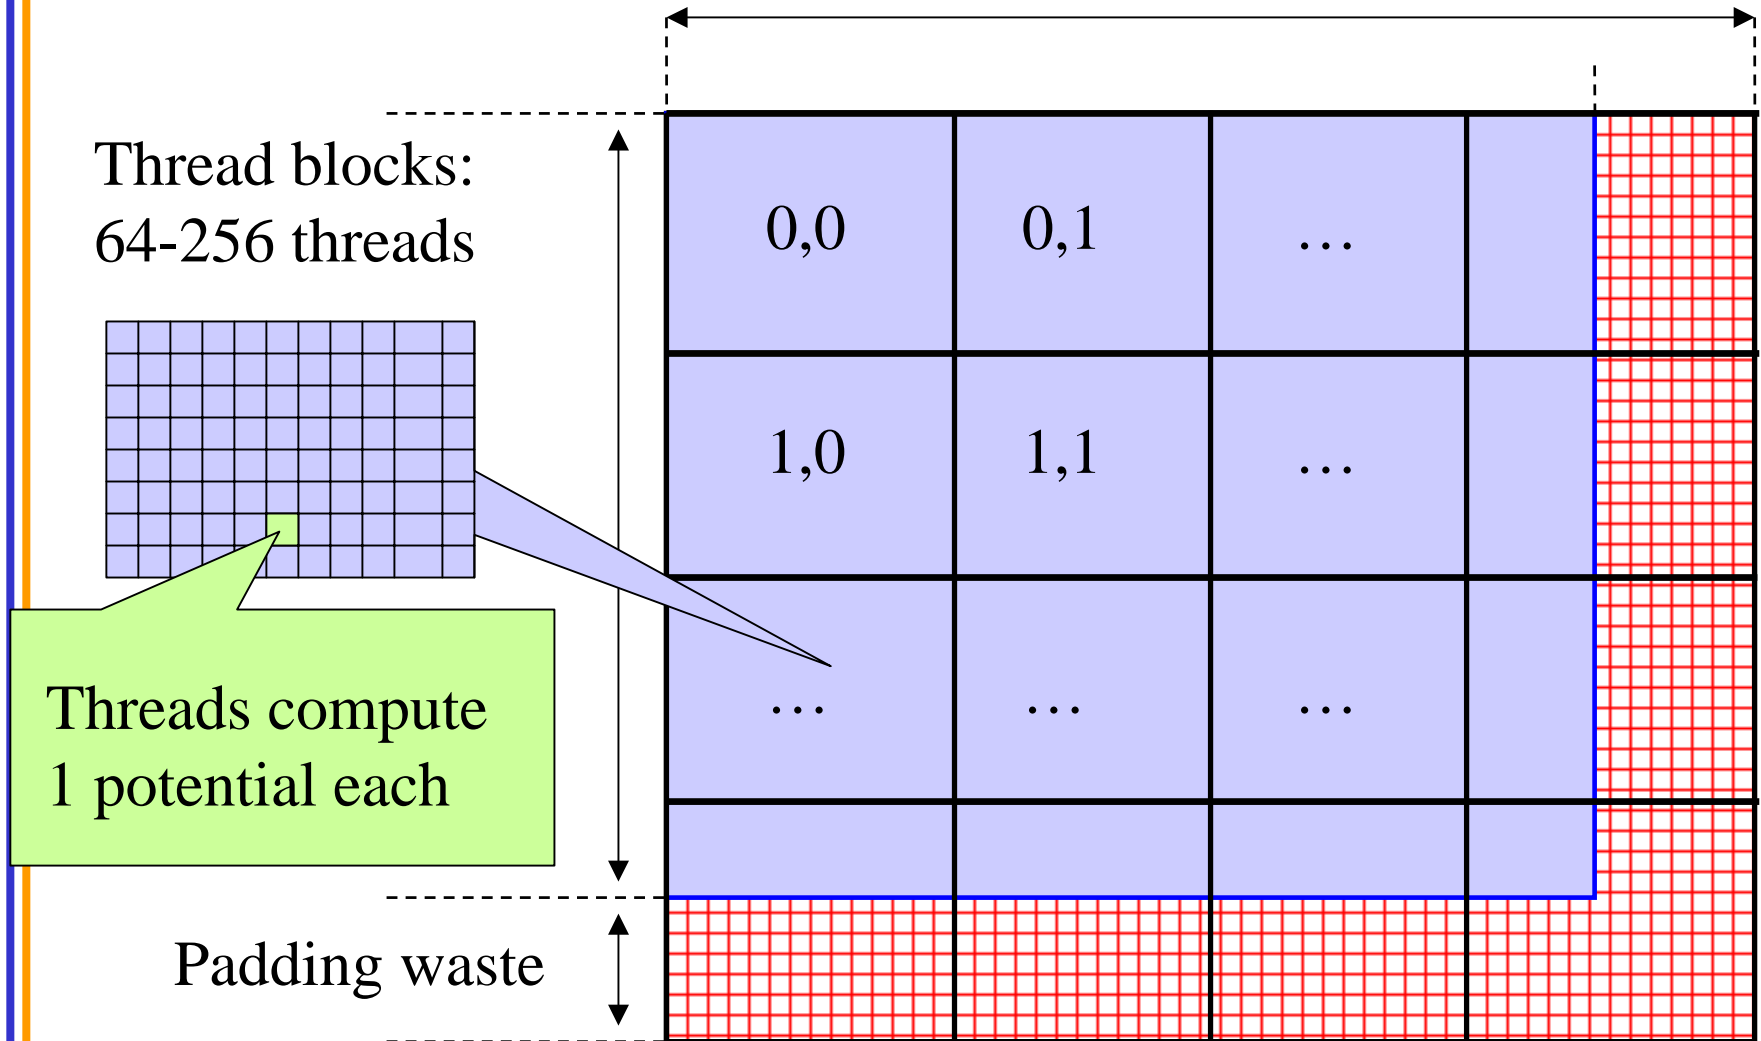
Direct Coulomb Summation on the GPU



DCS CUDA Block/Grid Decomposition

(non-unrolled)

Grid of thread blocks:



DCS CUDA Block/Grid

Decomposition (non-unrolled)

- 16x16 CUDA thread blocks are a nice starting size with a satisfactory number of threads
- Small enough that there's not much waste due to padding at the edges

Notes on Benchmarking CUDA Kernels: Initialization Overhead

- When a host thread initially binds to a CUDA context, there is a small (~100ms) delay during the first CUDA runtime call that touches state on the device
- The first time each CUDA kernel is executed, there's a small delay while the driver compiles the device-independent PTX intermediate code for the physical device associated with the current context
- In most real codes, these sources of one-time initialization overhead would occur at application startup and should not be a significant factor.
- The exception to this is that newly-created host threads incur overhead when they bind to their device, so it's best to re-use existing host threads than to generate them repeatedly

Notes on Benchmarking CUDA Kernels: Power Management, Async Operations

- Modern GPUs (and of course CPUs) incorporate power management hardware that reduces clock rates and/or powers down functional units when idle
- In order to benchmark peak performance of CUDA kernels, both the GPU(s) and CPU(s) must be awoken from their respective low-power modes
- In order to get accurate and repeatable timings, do a “warm up” pass prior to running benchmark timings on your kernel and any associated I/O
- Call `cudaThreadSynchronize()` prior to stopping timers to verify that any outstanding kernels and I/Os have completed

DCS Version 1: Const+Precalc

187 GFLOPS, 18.6 Billion Atom Evals/Sec

- Pros:
 - Pre-compute dz^2 for entire slice
 - Inner loop over read-only atoms, const memory ideal
 - If all threads read the same const data at the same time, performance is similar to reading a register
- Cons:
 - Const memory only holds ~ 4000 atom coordinates and charges
 - Potential summation must be done in multiple kernel invocations per slice, with const atom data updated for each invocation
 - Host must shuffle data in/out for each pass

DCS Version 1: Kernel Structure

...

```
float curenergy = energygrid[outaddr];
float coorx = gridspaceing * xindex;
float coory = gridspaceing * yindex;
int atomid;
float energyval=0.0f;
for (atomid=0; atomid<numatoms; atomid++) {
    float dx = coorx - atominfo[atomid].x;
    float dy = coory - atominfo[atomid].y;
    energyval += atominfo[atomid].w *
                rsqrtf(dx*dx + dy*dy + atominfo[atomid].z);
}
energygrid[outaddr] = curenergy + energyval;
```

Start global memory reads early. Kernel hides some of its own latency.

Only dependency on global memory read is at the end of the kernel...

DCS CUDA Block/Grid Decomposition (unrolled)

- Reuse atom data and partial distance components multiple times
- Use “unroll and jam” to unroll the outer loop into the inner loop
- Uses more registers, but increases arithmetic intensity significantly
- Kernels that unroll the inner loop calculate more than one lattice point per thread result in larger computational tiles:
 - Thread count per block must be decreased to reduce computational tile size as unrolling is increased
 - Otherwise, tile size gets bigger as threads do more than one lattice point evaluation, resulting on a significant increase in padding and wasted computations at edges

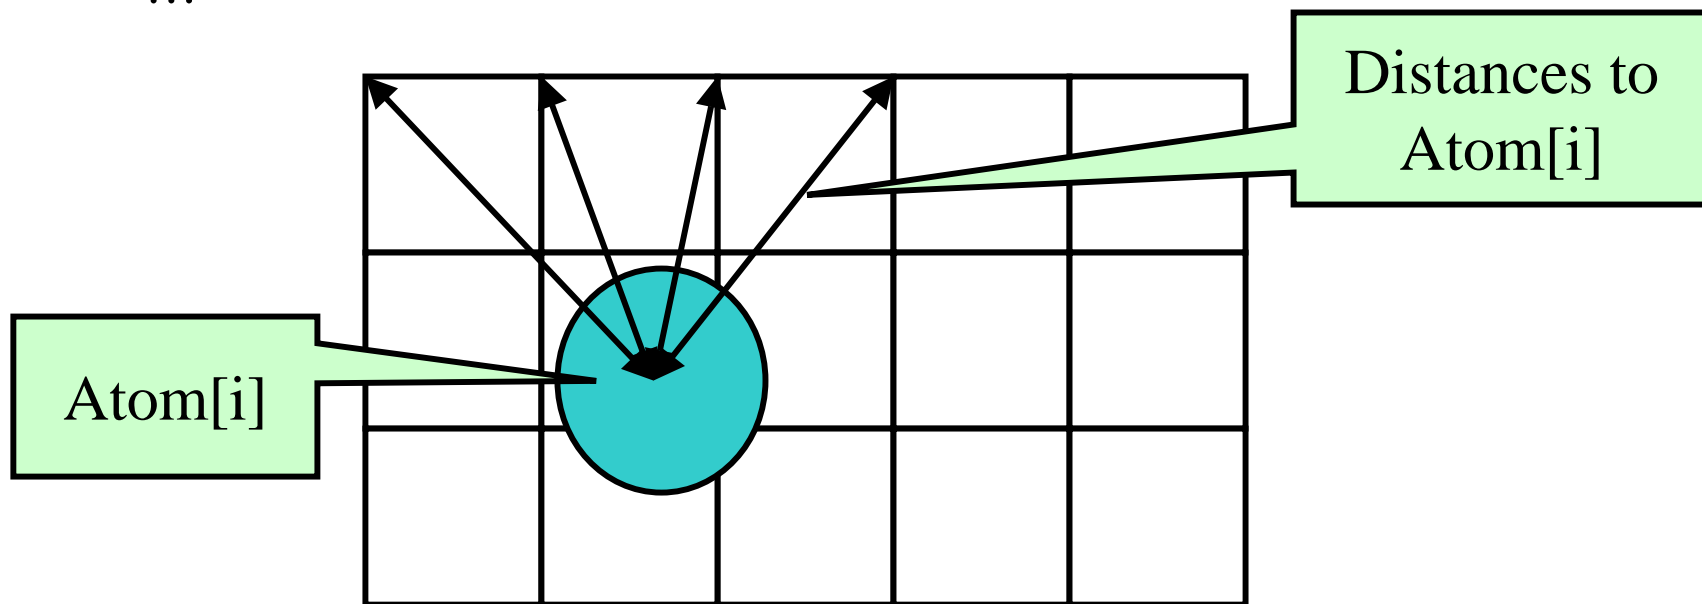
DCS CUDA Algorithm: Unrolling Loops

- Add each atom's contribution to several lattice points at a time, distances only differ in one component:

$$\text{potential}[j] \ += \ \text{atom}[i].\text{charge} / r_{ij}$$

$$\text{potential}[j+1] \ += \ \text{atom}[i].\text{charge} / r_{i(j+1)}$$

...




DCS CUDA Block/Grid Decomposition

(unrolled)

Grid of thread blocks:


Unrolling increases computational tile size

Thread blocks:
64-256 threads



Threads compute up to 8 potentials

Padding waste



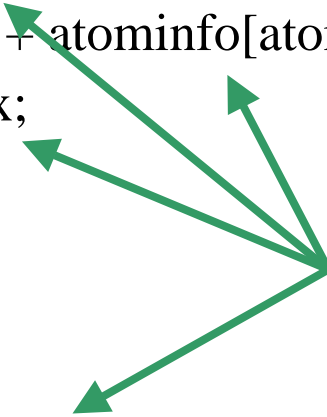
DCS Version 2: Const+Precalc+Loop Unrolling

259 GFLOPS, 33.4 Billion Atom Evals/Sec

- Pros:
 - Although const memory is very fast, loading values into registers costs instruction slots
 - We can reduce the number of loads by reusing atom coordinate values for multiple voxels, by storing in regs
 - By unrolling the X loop by 4, we can compute dy^2+dz^2 once and use it multiple times, much like the CPU version of the code does
- Cons:
 - Compiler won't do this type of unrolling for us (yet)
 - Uses more registers, one of several finite resources
 - Increases effective tile size, or decreases thread count in a block, though not a problem at this level

DCS Version 2: Inner Loop

```
...for (atomid=0; atomid<numatoms; atomid++) {  
    float dy = coory - atominfo[atomid].y;  
    float dysqpdzsq = (dy * dy) + atominfo[atomid].z;  
    float x = atominfo[atomid].x;  
    float dx1 = coorx1 - x;  
    float dx2 = coorx2 - x;  
    float dx3 = coorx3 - x;  
    float dx4 = coorx4 - x;  
    float charge = atominfo[atomid].w;  
    energyvalx1 += charge * rsqrtf(dx1*dx1 + dysqpdzsq);  
    energyvalx2 += charge * rsqrtf(dx2*dx2 + dysqpdzsq);  
    energyvalx3 += charge * rsqrtf(dx3*dx3 + dysqpdzsq);  
    energyvalx4 += charge * rsqrtf(dx4*dx4 + dysqpdzsq);  
}
```



Compared to non-unrolled kernel: memory loads are decreased by 4x, and FLOPS per evaluation are reduced, but register use is increased...

DCS Version 3:

Const+Shared+Loop Unrolling+Precalc
268 GFLOPS, 36.4 Billion Atom Evals/Sec

- Pros:
 - Loading prior potential values from global memory into shared memory frees up several registers, so we can afford to unroll by 8 instead of 4
 - Using fewer registers allows co-scheduling of more blocks, increasing GPU “occupancy”
- Cons:
 - Bumping against hardware limits (uses all const memory, most shared memory, and a largish number of registers)

DCS Version 3: Kernel Structure

- Loads 8 potential map lattice points from global memory at startup, and immediately stores them into shared memory before going into inner loop. We would otherwise consume too many registers and lose performance (on G80 at least...)
- Processes 8 lattice points at a time in the inner loop
- Additional performance gains are achievable by coalescing global memory reads at start/end

DCS Version 3: Inner Loop

```
...for (v=0; v<8; v++)
    curenergies[tid + nthr * v] = energygrid[outaddr + v];
float coorx = gridspacing * xindex;
float coory = gridspacing * yindex;
float energyvalx1=0.0f; [.....] float energyvalx8=0.0f;
for (atomid=0; atomid<numatoms; atomid++) {
    float dy = coory - atominfo[atomid].y;
    float dysqpdzsq = (dy * dy) + atominfo[atomid].z;
    float dx = coorx - atominfo[atomid].x;
    energyvalx1 += atominfo[atomid].w * rsqrtf(dx*dx + dysqpdzsq);
    dx += gridspacing;
[...]
```

```
    energyvalx8 += atominfo[atomid].w * rsqrtf(dx*dx + dysqpdzsq);
}
```

```
__syncthreads(); // guarantee that shared memory values are ready for reading by all threads
energygrid[outaddr    ] = energyvalx1 + curenergies[tid    ];
[...]
```

```
energygrid[outaddr + 7] = energyvalx2 + curenergies[tid + nthr * 7];
```

DCS Version 4:

Const+Loop Unrolling+Coalescing

291.5 GFLOPS, 39.5 Billion Atom Evals/Sec

- Pros:
 - Simplified structure compared to version 3, no use of shared memory, register pressure kept at bay by doing global memory operations only at the end of the kernel
 - Using fewer registers allows co-scheduling of more blocks, increasing GPU “occupancy”
 - Doesn’t have as strict of a thread block dimension requirement as version 3, computational tile size can be smaller
- Cons:
 - The computation tile size is still large, so small potential maps don’t perform as well as large ones

DCS Version 4: Kernel Structure

- Processes 8 lattice points at a time in the inner loop
- Subsequent lattice points computed by each thread are offset by a half-warp to guarantee coalesced memory accesses
- Loads and increments 8 potential map lattice points from global memory at completion of of the summation, avoiding register consumption
- Code is too long to show as a snippet
- Source code is available by request

DCS Version 4: Inner Loop

```
...float coory = gridspacing * yindex;
float coorx = gridspacing * xindex;
float gridspacing_coalesce = gridspacing * BLOCKSIZE;
int atomid;
for (atomid=0; atomid<numatoms; atomid++) {
    float dy = coory - atominfo[atomid].y;
    float dyz2 = (dy * dy) + atominfo[atomid].z;
    float dx1 = coorx - atominfo[atomid].x;
[...]
```

```
float dx8 = dx7 + gridspacing_coalesce;
energyvalx1 += atominfo[atomid].w * rsqrtf(dx1*dx1 + dyz2);
[...]
```

```
energyvalx8 += atominfo[atomid].w * rsqrtf(dx8*dx8 + dyz2);
}
energygrid[outaddr          ] += energyvalx1;
[...]
```

```
energygrid[outaddr+7*BLOCKSIZE] += energyvalx7;
```

Points spaced for
memory coalescing

Reuse partial distance
components $dy^2 + dz^2$

Global memory ops
occur only at the end
of the kernel,
decreases register use


DCS CUDA Block/Grid Decomposition

(unrolled, coalesced)

Grid of thread blocks:


Unrolling increases computational tile size

Thread blocks:
64-256 threads




Threads compute up to 8 potentials, skipping by half-warps

Padding waste



Direct Coulomb Summation Performance

Number of thread blocks modulo number of SMs results in significant performance variation for small workloads



CUDA-Unroll8clx:
fastest GPU kernel,
44x faster than CPU,
291 GFLOPS on
GeForce 8800GTX

CUDA-Simple:
14.8x faster,
33% of fastest
GPU kernel

GPU computing. J. Owens, M. Houston, D. Luebke, S. Green, J. Stone, J. Phillips. *Proceedings of the IEEE*, 96:879-899, 2008.

Multi-GPU DCS Potential Map Calculation

- Both CPU and GPU versions of the code are easily parallelized by decomposing the 3-D potential map into slices, and computing them concurrently
- Potential maps often have 50-500 slices in the Z direction, so plenty of coarse grain parallelism is still available via the DCS algorithm

Multi-GPU DCS Algorithm:

- One host thread is created for each CUDA GPU, attached according to host thread ID:
 - First CUDA call binds that thread's CUDA context to that GPU for life
 - Map slices are decomposed cyclically onto the available GPUs
 - Handling error conditions within child threads is dependent on the thread library and, makes dealing with any CUDA errors somewhat tricky. Easiest way to deal with this is with a shared exception queue/stack for all of the worker threads.
- Map slices are usually larger than the host memory page size, so false sharing and related effects are not a problem for this application

Multi-GPU Direct Coulomb Summation


- Effective memory bandwidth scales with the number of GPUs utilized
- PCIe bus bandwidth not a bottleneck for this algorithm
- 117 billion evals/sec
- 863 GFLOPS
- 131x speedup vs. CPU core
- Power: 700 watts during benchmark



Quad-core Intel QX6700

Three NVIDIA GeForce 8800GTX

Multi-GPU Direct Coulomb Summation



NCSA GPU Cluster


<http://www.ncsa.uiuc.edu/Projects/GPUcluster/>

	Evals/sec	TFLOPS	Speedup*
4-GPU (2 Quadroplex) Opteron node at NCSA	157 billion	1.16	176
4-GPU GTX 280 (GT200) In new NCSA Lincoln cluster	241 billion	1.78	271

*Speedups relative to Intel QX6700 CPU core w/ SSE

Multi-GPU DCS Performance: Initial Ion Placement Lattice Calculation

- Original virus DCS ion placement ran for 110 CPU-hours on SGI Altix Itanium2
- Same calculation now takes 1.35 GPU-hours
- 27 minutes (wall clock) if three GPUs are used concurrently
- CUDA Initial ion placement lattice calculation performance:
 - 82 times faster for virus (STMV) structure
 - 110 times faster for ribosome
- Three GPUs give performance equivalent to ~330 Altix CPUs for the ribosome case



Satellite Tobacco Mosaic Virus (STMV)
Ion Placement

Multi-Level Summation Method for Coulomb Potential

Infinite vs. Cutoff Potentials

- Infinite range potential:
 - All atoms contribute to all lattice points
 - Summation algorithm has quadratic complexity
- Cutoff (range-limited) potential:
 - Atoms contribute within cutoff distance to lattice points
 - Summation algorithm has linear time complexity
 - Has many applications in molecular modeling:
 - Replace electrostatic potential with shifted form
 - Short-range part for fast methods of approximating full electrostatics
 - Used for fast decaying interactions (e.g. Lennard-Jones, Buckingham)


Short-range Cutoff Summation

- Each lattice point accumulates electrostatic potential contribution from atoms within cutoff distance:

if ($r_{ij} < \text{cutoff}$)

potential[j] += (charge[i] / r_{ij}) * $s(r_{ij})$

- Smoothing function $s(r)$ is algorithm dependent



Cutoff Summation on the GPU

Atoms are spatially hashed into fixed-size bins

CPU handles overflowed bins (GPU kernel can be very aggressive)

GPU thread block calculates corresponding region of potential map,


Bin/region neighbor checks costly; solved with universal table look-up

Each thread block cooperatively loads atom bins from surrounding neighborhood into shared memory for evaluation


Shared memory

atom bin

Global memory



Potential map regions

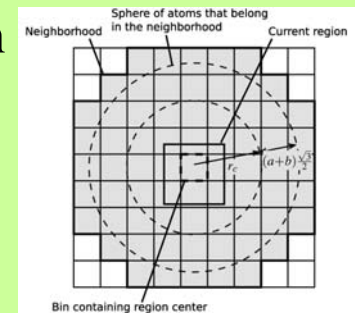


Bins of atoms

Constant memory

Offsets for bin neighborhood


Look-up table encodes "logic" of spatial geometry



Using the CPU to Improve GPU Performance

- GPU performs best when the work evenly divides into the number of threads/processing units
- Optimization strategy:
 - Use the CPU to “regularize” the GPU workload
 - Use fixed size bin data structures, with “empty” slots skipped or producing zeroed out results
 - Handle exceptional or irregular work units on the CPU while the GPU processes the bulk of the work
 - On average, the GPU is kept highly occupied, attaining a much higher fraction of peak performance

Cutoff Summation Runtime



GPU cutoff with
CPU overlap:
17x-21x faster than
CPU core

If asynchronous
stream blocks due
to queue filling,
performance will
degrade from
peak...

GPU acceleration of cutoff pair potentials for molecular modeling applications.
C. Rodrigues, D. Hardy, J. Stone, K. Schulten, W. Hwu. *Proceedings of the 2008 Conference On Computing Frontiers*, pp. 273-282, 2008.

Cutoff Summation Observations

- Use of CPU to handle overflowed bins is very effective, overlaps completely with GPU work
- One caveat when using streaming API is to avoid overfilling the stream queue with work, as doing so can trigger blocking behavior (greatly improved in current drivers)
- The use of compensated summation (all GPUs) or double-precision (GT200 only) for potential accumulation resulted in only a ~10% performance penalty vs. pure single-precision arithmetic, while reducing the effects of floating point truncation

Multilevel Summation


- Approximates full electrostatic potential
- Calculates sum of smoothed pairwise potentials interpolated from a hierarchy of lattices
- Advantages over PME and/or FMM:
 - Algorithm has linear time complexity
 - Permits non-periodic and periodic boundaries
 - Produces continuous forces for dynamics (advantage over FMM)
 - Avoids 3-D FFTs for better parallel scaling (advantage over PME)
 - Spatial separation allows use of multiple time steps
 - Can be extended to other pairwise interactions
- Skeel, Tezcan, Hardy, *J Comp Chem*, 2002 — Computing forces for molecular dynamics
- Hardy, Stone, Schulten, *J Paral Comp*, 2009 — GPU-acceleration of potential map calculation

Multilevel Summation Main Ideas

Split the $1/r$ potential

Interpolate the smoothed potentials

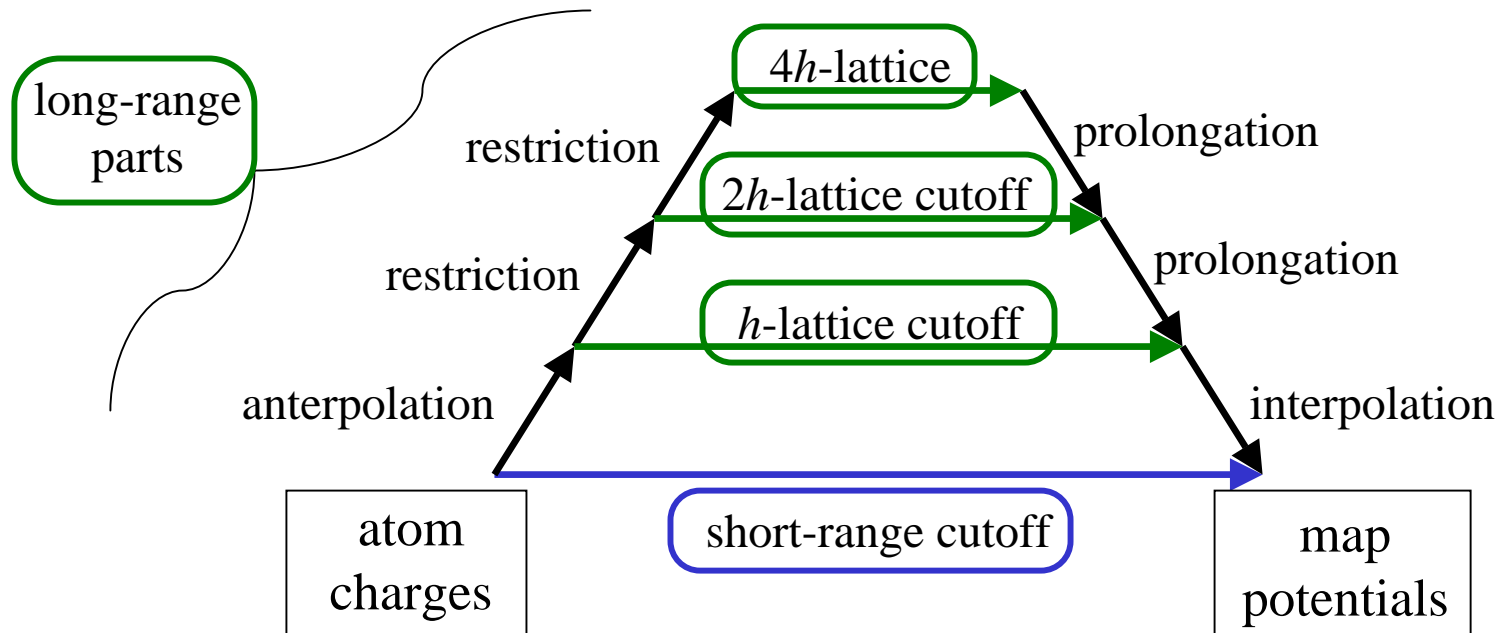
- Split the $1/r$ potential into a short-range cutoff part plus smoothed parts that are successively more slowly varying. All but the top level potential are cut off.
- The smoothed potentials are interpolated from successively coarser lattices.
- The lattice spacing is doubled at each successive level. The cutoff distance is also doubled.



Multilevel Summation Calculation

$$\text{map potential} = \text{exact short-range interactions} + \text{interpolated long-range interactions}$$

Computational Steps




Multilevel Summation on the GPU

Accelerate **short-range cutoff** and **lattice cutoff** parts

Performance profile for 0.5 Å map of potential for 1.5 M atoms.
Hardware platform is Intel QX6700 CPU and NVIDIA GTX 280.

Computational steps	CPU (s)	w/ GPU (s)	Speedup
Short-range cutoff	480.07	14.87	32.3
Long-range anterpolation	0.18		
restriction	0.16		
lattice cutoff	49.47	1.36	36.4
prolongation	0.17		
interpolation	3.47		
Total	533.52	20.21	26.4



Lessons Learned

- GPU algorithms need fine-grained parallelism and sufficient work to fully utilize the hardware
- Much of per-thread GPU algorithm optimization revolves around efficient use of multiple memory systems and latency hiding
- Concurrency can often be traded for per-thread performance, in combination with increased use of registers or shared memory
- Fine-grained GPU work decompositions often compose well with the comparatively coarse-grained decompositions used for multicore or distributed memory programming

Lessons Learned (2)

- The host CPU can potentially be used to “regularize” the computation for the GPU, yielding better overall performance
- Overlapping CPU work with GPU can hide some communication and unaccelerated computation
- Targeted use of double-precision floating point arithmetic, or compensated summation can reduce the effects of floating point truncation at low cost to performance

Summary

- GPUs are not a magic bullet, but they can perform amazingly well when used effectively
- There are many good strategies for extracting high performance from individual subsystems on the GPU
- It is wise to begin with a well designed application and a thorough understanding of its performance characteristics on the CPU before beginning work on the GPU
- By making effective use of multiple GPU subsystems at once, tremendous performance levels can potentially be attained

Acknowledgements

- Additional Information and References:
 - <http://www.ks.uiuc.edu/Research/gpu/>
 - <http://www.ks.uiuc.edu/Research/vmd/>
- Questions, source code requests:
 - John Stone: johns@ks.uiuc.edu
- Acknowledgements:
 - D. Hardy, J. Saam, J. Phillips, P. Freddolino, L. Trabuco, J. Cohen, K. Schulten (UIUC TCB Group)
 - Prof. Wen-mei Hwu, Christopher Rodrigues (UIUC IMPACT Group)
 - David Kirk and the CUDA team at NVIDIA
 - NIH support: P41-RR05969
 - UIUC NVIDIA Center of Excellence

Publications

<http://www.ks.uiuc.edu/Research/gpu/>

- Long time-scale simulations of in vivo diffusion using GPU hardware. E. Roberts, J. Stone, L. Sepulveda, W. Hwu, Z. Luthey-Schulten. In *Eighth IEEE International Workshop on High Performance Computational Biology*, 2009. In press.
- High Performance Computation and Interactive Display of Molecular Orbitals on GPUs and Multi-core CPUs. J. Stone, J. Saam, D. Hardy, K. Vandivort, W. Hwu, K. Schulten, *2nd Workshop on General-Purpose Computation on Graphics Processing Units (GPGPU-2)*, *ACM International Conference Proceeding Series*, volume 383, pp. 9-18, 2009.
- Multilevel summation of electrostatic potentials using graphics processing units. D. Hardy, J. Stone, K. Schulten. *J. Parallel Computing*, 35:164-177, 2009.
- Adapting a message-driven parallel application to GPU-accelerated clusters. J. Phillips, J. Stone, K. Schulten. *Proceedings of the 2008 ACM/IEEE Conference on Supercomputing*, IEEE Press, 2008.
- GPU acceleration of cutoff pair potentials for molecular modeling applications. C. Rodrigues, D. Hardy, J. Stone, K. Schulten, and W. Hwu. *Proceedings of the 2008 Conference On Computing Frontiers*, pp. 273-282, 2008.
- GPU computing. J. Owens, M. Houston, D. Luebke, S. Green, J. Stone, J. Phillips. *Proceedings of the IEEE*, 96:879-899, 2008.
- Accelerating molecular modeling applications with graphics processors. J. Stone, J. Phillips, P. Freddolino, D. Hardy, L. Trabuco, K. Schulten. *J. Comp. Chem.*, 28:2618-2640, 2007.
- Continuous fluorescence microphotolysis and correlation spectroscopy. A. Arkhipov, J. Hüve, M. Kahms, R. Peters, K. Schulten. *Biophysical Journal*, 93:4006-4017, 2007.