Multiplying the speed-ups: GPU-accelerated, fast multipole BEM, for applications in protein electrostatics.

Lorena A Barba[1], Rio Yokota[1], Jaydeep P Bardhan[2], Matthew G Knepley[3]

[1] Boston University, [2] Rush University Medical Center, [3] University of Chicago

BOSTON UNIVERSITY

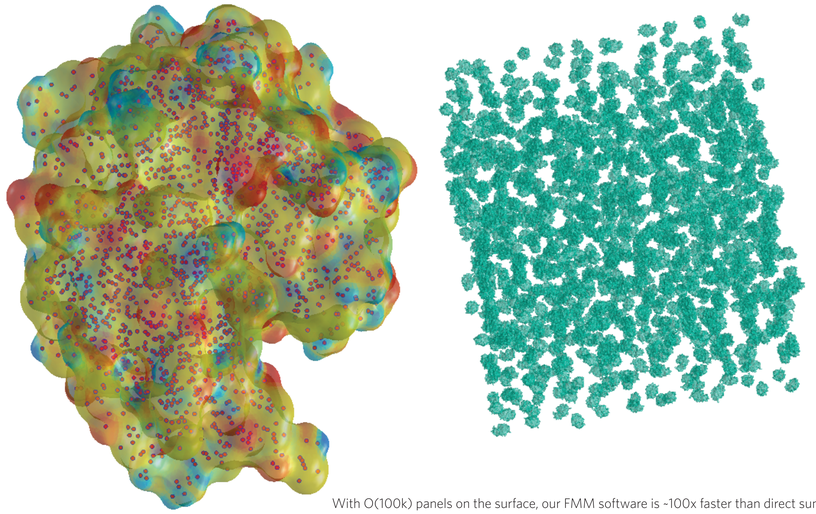

Far left: A lysozyme molecule surface, shown in transparency, with the atomic locations inside. **The surface has 100k points**. *Left*: An arrangement of 1000 such molecules, randomly placed inside a cubic volume.

Lysozyme—abundant in secretions, such as tears and saliva—is part of the immune system and a natural form of protection from pathogens such as E.coli.

With O(100k) panels on the surface, our FMM software is ~100x faster than direct summation on the GPU. On one GPU, the FMM-accelerated boundary element method (BEM) solver is ~10x faster than on one CPU—this is application speedup (not an inner kernel!).

A quote. It expresses our motivation

> *"... the fundamental law of computer science [is]: the faster the computer, the greater the importance of speed of algorithms."*
>
> Trefethen & Bau, Numerical Linear Algebra, SIAM (1997)

What's new? Fast algorithms on the GPU

The first wave of successful applications of GPUs to scientific computing was crowded with highly parallel methods. The paradigmatic example is molecular dynamics and other *N*-body simulations, where the embarrassingly parallel problem of calculating the all-pairs interactions exploits the fine-grained parallelism of the hardware very well. **But the easy pickings are running out.**

Many important applications involve intricate algorithms that require *"going back to the algorithmic drawing board"* (to quote an Intel blog) for a successful implementation on the GPU. This is the case of fast, O(*N*) algorithms like the fast multipole method (FMM). The FMM accelerates *N*-body problems by representing clusters of bodies with series expansions, and using a hierarchical tree structure to organize the bodies in space. There are various operations needed in a tree or FMM algorithm. We have ported all of them to the GPU.

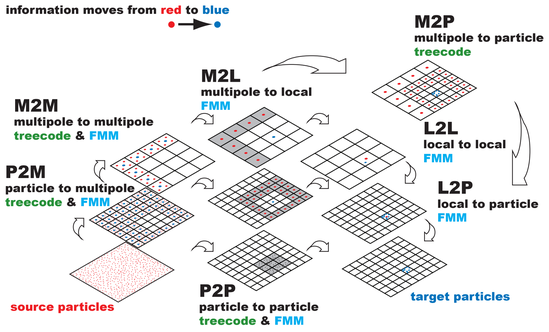

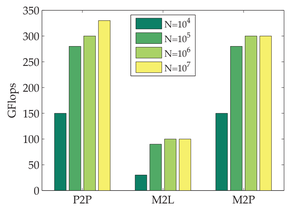

The impact of GPU acceleration is greater for large problems. Also, the cross-over point between a direct evaluation and the FMM occurs at higher *N* on the GPU than CPU.

For $N > 4 \times 10^4$, the FMM on the GPU is faster than the direct all-pairs evaluation.

Actual performance on the GPU of three core kernels, for four different values of N.

The model. Continuum electrostatics

Molecular dynamics is very detailed and accurate, but spends too much time computing the surrounding water molecules. An alternative model considers the molecule and the water as continuum dielectric media. Point charges are placed at the location of atoms inside the molecule.
This results in a **mixed-dielectric Poisson problem**.

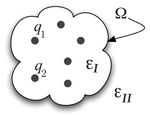

The problem can be written as a boundary-integral equation for charge density on the molecular surface:

$$\left(1 - \frac{\epsilon_I}{\epsilon_{II}}\right)\left(\frac{\partial}{\partial n(r)}\sum_{i=1}^{n_c}\frac{q_i}{4\pi||r-r_i||} + \frac{\partial}{\partial n(r)}\int_\Omega \frac{\sigma(r')}{4\pi||r-r'||}dA'\right) = \sigma(r)$$

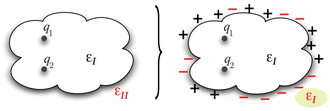

The surface charge density reproduces the potential of the orginal problem, but in a *homogeneous* dielectric space.

Method. Fast multipole BEM

The boundary element method (BEM) solution of the integral equation problem results in a linear system with *N* unknows, with a dense matrix. Solving it with iterative methods would require $O(N^2)$ calculations for the matrix-vector products.

We have developed a fast multipole (BEM) for biomolecular electrostatics. With GPU acceleration of the FMM, there is a **multiplicative speed-up** resulting from the fast O(*N*) algorithm and GPU hardware. We can obtain converged results for multi-million atom system in less than an hour, using multi-GPU clusters.

Now the bottleneck is generating the surface mesh.

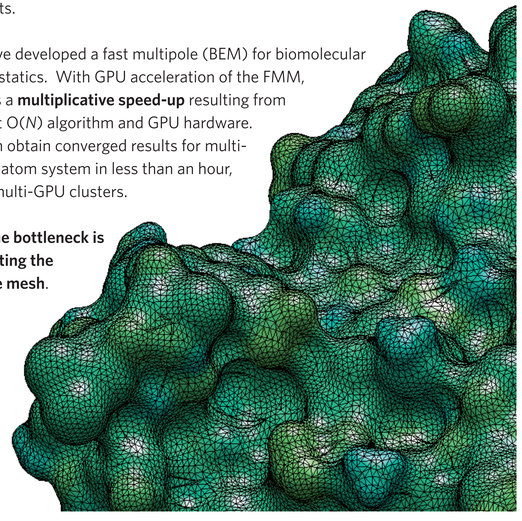

Results. Multi-GPU performance

To demonstrate the capability on large problems, we use a large collection of randomly oriented lysozyme molecules, arranged inside a cubic volume. One such collection is shown in the figure above. This setup is meant to mimic Brownian dynamics of a crowded molecule environment.

The largest calculation we conducted consists of

- 10,648 lysozyme molecules
- each surface discretized into 102,486 elements
- more than 20 million atoms
- **over 1 billion unknowns**

This calculation required only **~1 min per iteration on 512 GPUs**, using the cluster of the Nagasaki Advanced Computing Center, which was inexpensively built with 144 host nodes and 288 GTX 295 cards (PI: Prof. T. Hamada).
The strong scaling of the FMM on multi-GPUs is shown below, up to 128 GPUs.

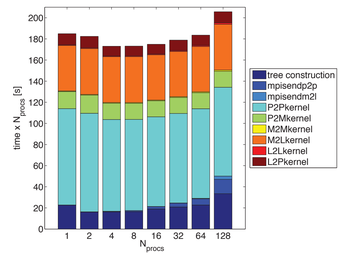

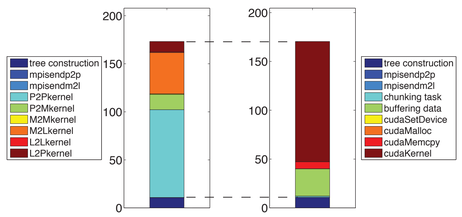

Who cares? Bioelectrostatics is important

Electrostatic interactions play a crucial role in the function of biological molecules. Functional properties that are ruled by electrostatics include:

1) electron transfer reactions
   - involved in key transduction processes, *e.g.*, photosynthesis
2) ligand binding to proteins
   - involved in structure-based drug desing
3) enzyme catalysis
   - involved in all chemistry of life!
   - behind much of biotechnology, *e.g.*, biofuels
4) protein folding and stability
   - involved in diseases like Alzheimer's

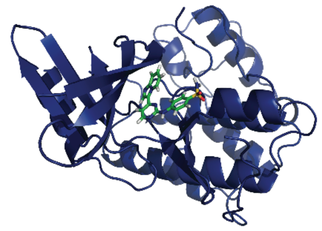

Want more? Papers and software are online

All the codes developed in our group are free (like free beer) and open source. To download them, follow the links from our group website:

http://barbagroup.bu.edu

Also on the website are up-to-date bibliographic references, and papers for download. ***Please visit!***

GPU Technology Conference, September 2010