Game-changing triad: fast algorithms, many-core hardware, open software

A case study with biomolecular electrostatics

Lorena A Barba, Boston University

Scientific computing has entered a new era. There are many examples of computational capability leading to major discoveries, thanks to the power to tackle more complex and accurate models than ever before. *E.g.*, a recent feature in *Science* reports unprecedented simulations of the entire Earth mantle, revealing new dynamics—an example of how the right methods can provide leaps in capability much larger than what Moore's law would achieve in a given period. And *Physiome*, the international effort to quantitatively cross-link all scales of life's organization (proteins to cells, tissues, organs and organism), epitomizes the complexity being tackled by computing—in this case, an example of how open science and open source can enable huge computational projects.

What does this new era look like? Without doubt, the rebirth of vector-type hardware such as GPUs is a protagonist. It is not a fad; the future *is* heterogeneous computing. Naturally, the first wave of successes with GPUs was crowded with highly parallel algorithms that adapted well to the hardware. But history (and Trefethen's quote) tells us that as machines grow faster, we will rely more on approximate but fast algorithms. The $O(N)$ fast multipole method (FMM) is a case in point. And a third ingredient that is vital for game-changing computational science is the open model for data and software.

> ***"... 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

We see at work this triad of extreme-scale computing in our application to biomolecular electrostatics. In the implicit-solvent model, a vast acceleration is possible via fast algorithms for long-range interactions, in particular the FMM. With an open-source, GPU-enabled solver based on boundary element methods, we aim at unprecedented capability. Our recent results solve 1 billion unknowns in ~1min, using 512 GPUs. On a desktop with one GPU, treating millions of unknows can be routine. We have also initiated a collaboration with NRL to apply these methods to investigate interactions of biomolecules with nanostructures used for biosensing.

right: The surface mesh of a biomolecule, for electrostatic modeling via BEM.
far right: Prof Barba, with her postdoc Rio Yokota (left) and collaborator Tsuyoshi Hamada (right), at the Nagasaki Advanced Computing Center.

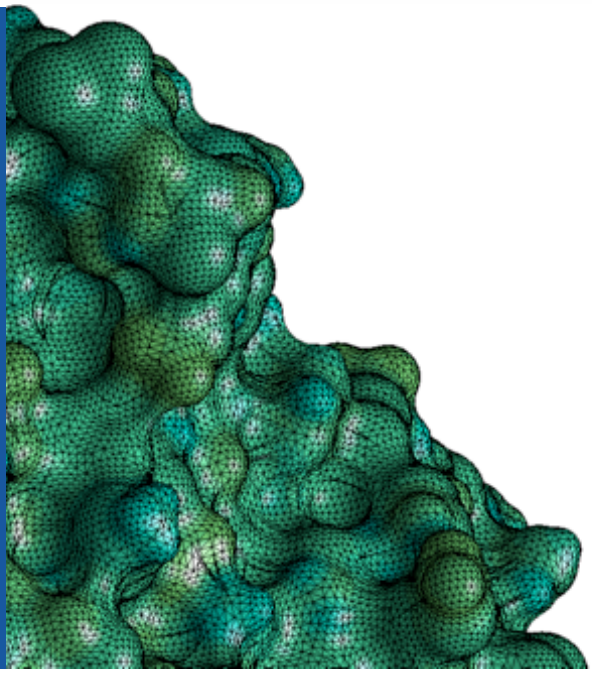

Seminar time and location: **2:45–3:45 pm, 210 Mudd (APAM conference room)**