Boundary-Integral Methods in Molecular Science and Engineering

Lecture 2: There's More Than One Way to Skin a Cat

Jaydeep P. Bardhan

Rush University Medical Center

jaydeep_bardhan@rush.edu

Outline for Lectures

✓ Biology is awesome. If you can solve Poisson, you can join in the fun!

- There's more than one way to skin a cat. Sometimes PDEs can be advantageously reframed as *integral equations*.

- Numerical solution of integral equations presents different challenges than do PDEs.

- A diversity of unusual computational challenges will continue to drive biological simulation.

Deriving a Boundary Integral Equation

- Key Concept: Unknowns are on *boundaries* between regions

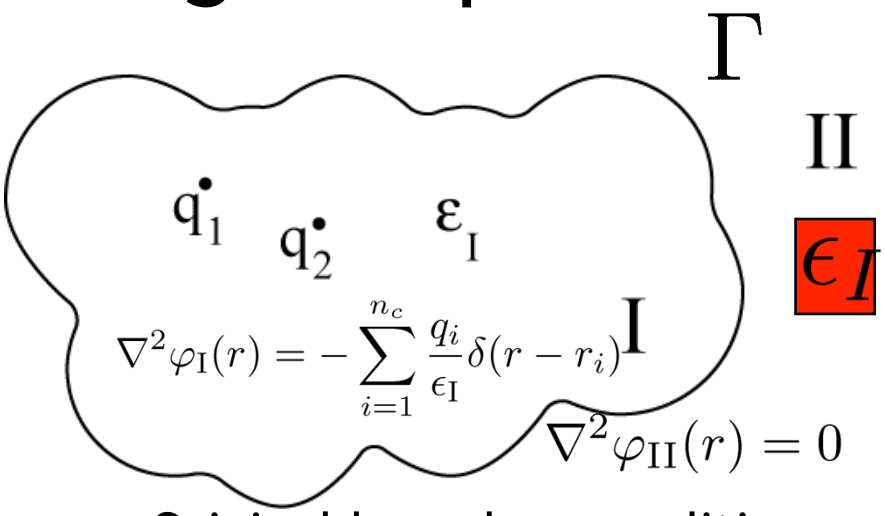

In homogeneous dielectric:

$$\hat{\varphi}(r) = \sum_i \frac{q_i}{4\pi\epsilon_I |r - r_i|} + \int_\Omega \frac{\sigma_p(r')}{4\pi\epsilon_I |r - r'|} dA'$$

$$\frac{\partial \hat{\varphi}_{\mathrm{II}}}{\partial n^+} - \frac{\partial \hat{\varphi}_{\mathrm{I}}}{\partial n^-} = \sigma_p(r)/\epsilon_{\mathrm{I}}$$

Original boundary conditions

$$\varphi_{\mathrm{I}}(r_\Omega) = \varphi_{\mathrm{II}}(r_\Omega)$$

$$\epsilon_{\mathrm{I}} \frac{\partial \varphi_{\mathrm{I}}}{\partial n}(r_\Omega) = \epsilon_{\mathrm{II}} \frac{\partial \varphi_{\mathrm{II}}}{\partial n}(r_\Omega)$$

$$\sigma_p(\mathbf{s}) + \frac{\Delta\epsilon(\mathbf{s})}{4\pi\bar{\epsilon}(\mathbf{s})}\mathbf{n}(\mathbf{s}) \cdot \int_\Omega \frac{\mathbf{s} - \mathbf{s}'}{|\mathbf{s} - \mathbf{s}'|^3} \sigma_p(\mathbf{s}') d\mathbf{s}' = -\frac{\Delta\epsilon(\mathbf{s})}{4\pi\bar{\epsilon}(\mathbf{s})}\mathbf{n}(\mathbf{s}) \cdot \sum_k \frac{q_k}{\epsilon(\mathbf{r}_k)} \frac{\mathbf{s} - \mathbf{r}_k}{|\mathbf{s} - \mathbf{r}_k|^3}$$

This can be derived from any of several paths: variational principles, Gauss's law, or Green's theorem

Why Bother With Integral Equations?

Easy problem:

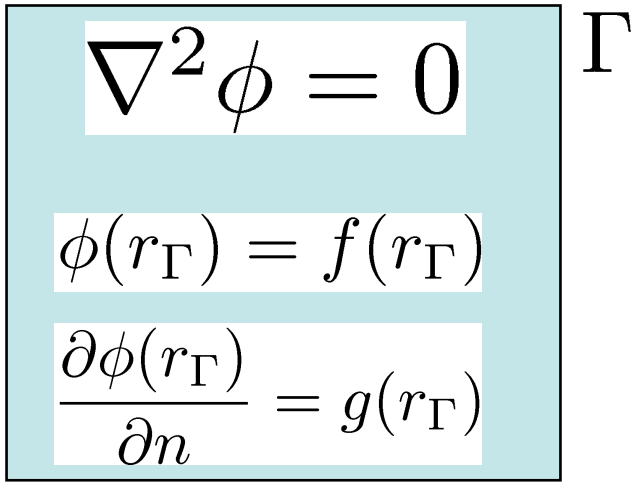

Medium problem:

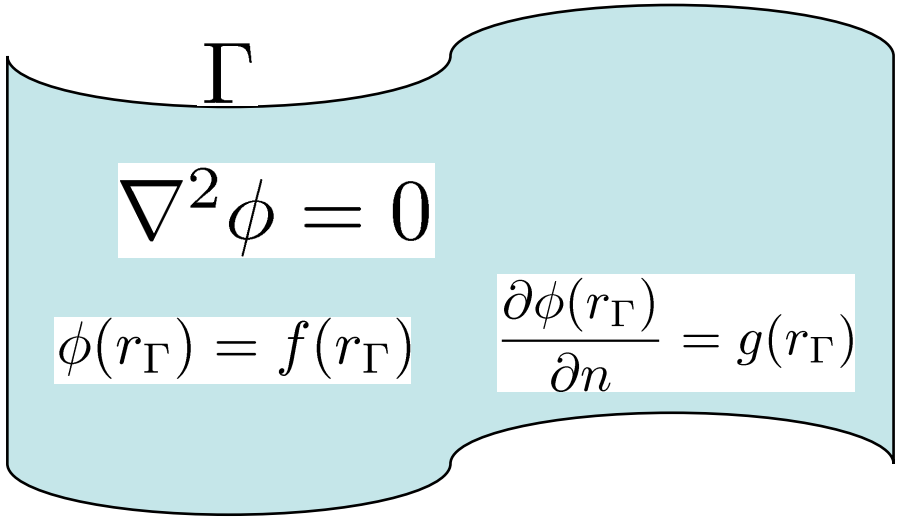

Exterior problems?

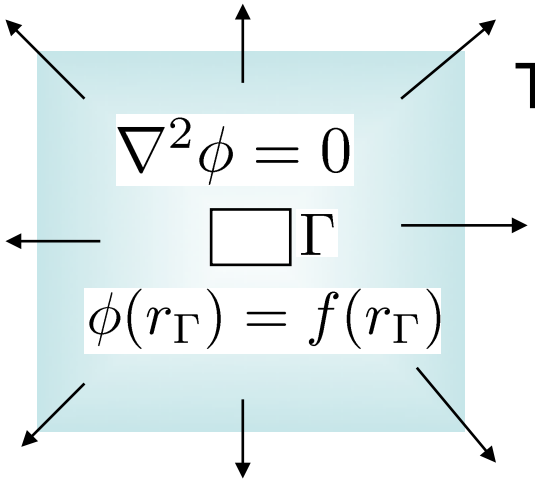

To infinity

Problems with mostly empty, uninteresting space

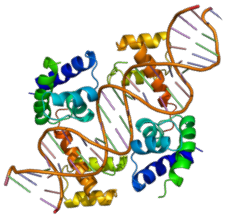

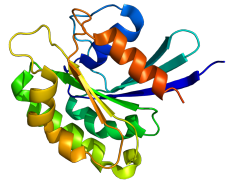

The Advantages of PDE Solvers

1. More general
   - nonlinear problems
   - continuously varying material properties
2. Easier to parallelize
3. Sometimes easier to write down

The Capacitance Problem

- Charge accumulates on surface of conductor when it is raised to a potential relative to ground:

$$Q = CV$$

$$V = \int_\Gamma \sigma(r') \frac{1}{|r - r'|} dA'$$

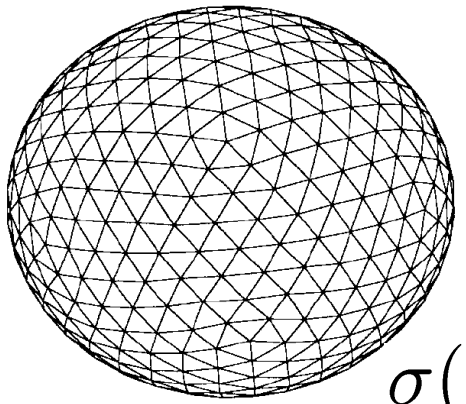

What about the singularity?

The singularity is *integrable!*

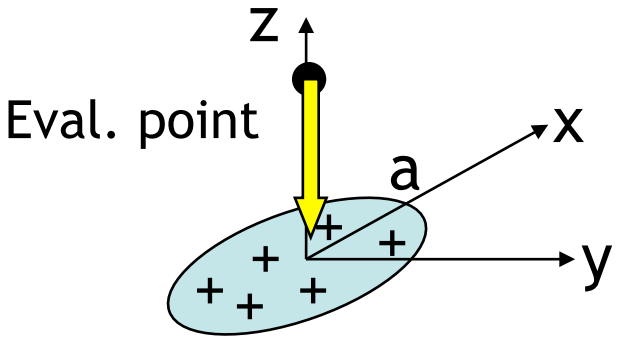

$$\phi(r) = \int \frac{1}{|r - r'|} dA'$$

$$= \int_0^a \frac{1}{|r'|} [2\pi r' dr']$$

$$= 2\pi a$$

Similarity Between FEM and BEM

- Both **weighted residual** methods:

FEM

$$\mathcal{L}^{\text{diff}} u = f$$

$$\tilde{u}(r) = \sum_i u_i \chi_i(r)$$

$$\mathcal{L}^{\text{diff}} \tilde{u} = f = e(r)$$

Enforce

$$< e, \chi_i >= 0$$

(Galerkin method)

BEM

$$\mathcal{L}^{\text{int}} u = f$$

$$\tilde{u}(r) = \sum_i u_i \chi_i(r)$$

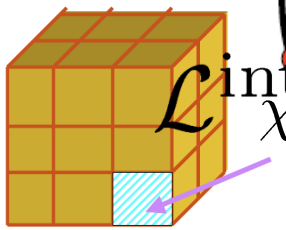

$$\mathcal{L}^{\text{int}} \chi_i(r) = f = e(r)$$

1 on panel i
0 elsewhere

Enforce

$$< e, w_i >= 0$$

Galerkin: $w_i(r) = \chi_i(r)$

Differences Between BEM and FEM

1. Extra freedom in choosing test functions

Collocation: test = delta functions

$$< e, \delta(r - r_i) >= 0$$

$$\int_\Omega \delta(r - r_i) \cdot \left(\mathcal{L}^{\text{int}} \tilde{u} - f\right) dA = 0$$

$$e(r_i) = \left(\mathcal{L}^{\text{int}} \tilde{u} - f\right)(r_i) = 0$$

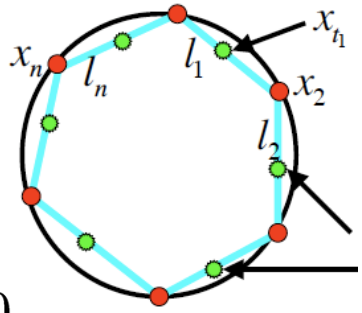

Centroids of elements

2. Matrix elements are harder to compute

Galerkin FEM:

$$A_{ij} = \int_\Omega \nabla\chi_i(r) \cdot \nabla\chi_j(r) d\Omega$$

Smooth integrand:
Easily computed with quadrature!

Galerkin BEM:

$$<\chi_i, \mathcal{L}^{\text{int}} \tilde{u} > = < \chi_i, f >$$

$$\int_\Omega \chi_i(r) \int_\Omega G(r, r') \mathcal{L}^{\text{int}} \left[\sum_j u_j \chi_j(r')\right] dA' dA d\Omega$$

Double integral of a singular function!!

Using Quadrature to Compute Panel Integrals

- In 1D, N-point Gauss quadrature is exact for polynomials up to order 2N-1

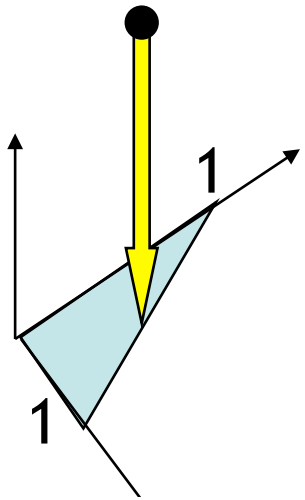

- What happens as the field point approaches the panel? (Here, the middle of the hypotenuse)

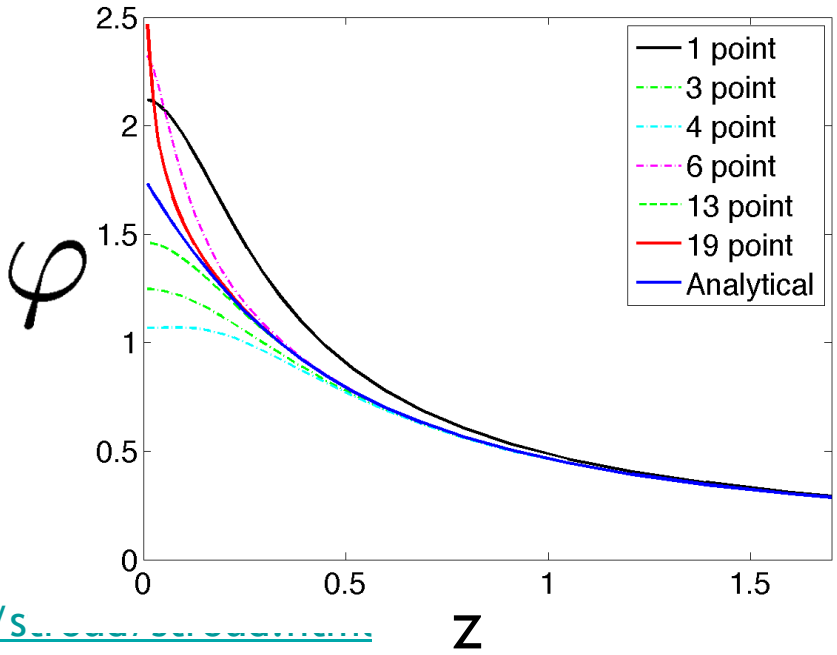

- Using Stroud's rules

- http://people.scs.fsu.edu/~burkardt/m_src/s...

Good News: Analytical Laplace Integrals

- For a planar element with polynomial charge distribution, we can analytically compute

$$\varphi(r) = \int_{\text{element}} \frac{\sigma(r')}{|r - r'|} dA'$$

Potential due to monopole distribution

$$\frac{\partial \varphi(r)}{\partial n(r)} = n(r) \cdot \nabla \int_{\text{element}} \frac{\sigma(r')}{|r - r'|} dA'$$

Normal field due to monopole

$$\varphi(r) = \int_{\text{element}} n(r') \cdot \nabla \frac{\mu(r')}{|r - r'|} dA'$$

Potential due to *dipole* distribution

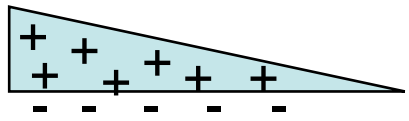

- Many people fear BIE/BEM due to panel integrals when they don't have to!

Hess+Smith, 1964; Newman, 1986

Green's Representation Formula

- Well-known fact: a function harmonic in a region *D* is completely specified by its boundary values

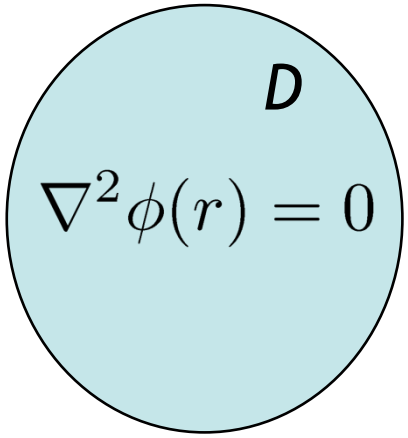

Dirichlet: given $\phi(r_D) = f(r_D)$

Neumann: given $\frac{\partial\phi}{\partial n}(r_D) = g(r_D)$

- Not so well known: if you know both, the potential anywhere in *D* is given by

$$\phi(r) = \int_S \phi(r') \frac{\partial}{\partial n(r')} \left(\frac{1}{|r - r'|} \right) dA' - \int_S \frac{\partial\phi(r')}{\partial n(r')} \left(\frac{1}{|r - r'|} \right) dA'$$

Thus you can solve Laplace by finding the other boundary condition!

Another Formulation

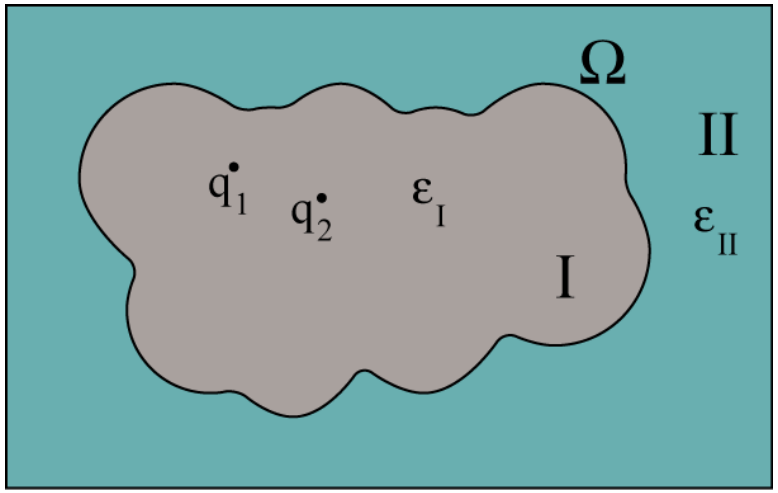

Region I:

$$\nabla^2\varphi(r) = -\sum_{i=1}^{n_c} q_i\delta(r - r_i)$$

Region II:

$$\nabla^2\varphi(r) = \kappa^2\varphi(r)$$

$$\begin{bmatrix} \frac{1}{2}I + D^a_{I,a} & -S^a_{I,a} \\ \frac{1}{2}I - D^a_{II,a} & \epsilon_{I,II}S^a_{II,a} \end{bmatrix} \begin{bmatrix} \phi_a \\ \frac{\partial\phi_a}{\partial n} \end{bmatrix} = \begin{bmatrix} \sum_i \frac{q_i}{\epsilon_I} G^a_{I,i} \\ 0 \end{bmatrix}$$

- Derivable using Green's theorem in the interior and exterior regions
- Unknowns are potential and its normal derivative!

Fast Solvers For Integral Equations

- Consider the physical meaning of

$$\phi(r) = \int_{\Omega} \sigma(r') \frac{1}{|r - r'|} dA'$$

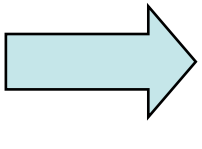

Computing the field at a number of surface points due to a distribution of sources!

- Adopt fast-summation methods like fast multipole, etc, with *preconditioned Krylov methods* such as GMRES

$$x^k \in \{b, Ab, A^2b, \ldots, A^{k-1}b\}$$

The Fast-Multipole Method

- Optimal (linear-scaling) algorithm for the N-body problem

$$f(\vec{x}_j) = \sum_{i=1}^{N} c_i K(\vec{x}_i - \vec{x}_j) \qquad j \in [1 \ldots N]$$

Direct $O(N^2)$ → FMM $O(N)$

Prohibitive!

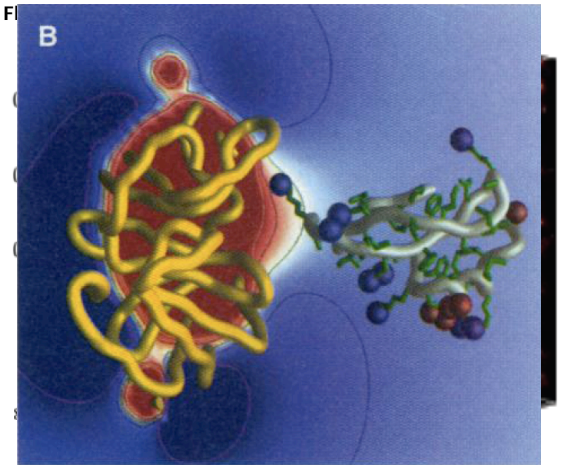

Image: T. A. Cruz
Honig+Nicholls (1995)

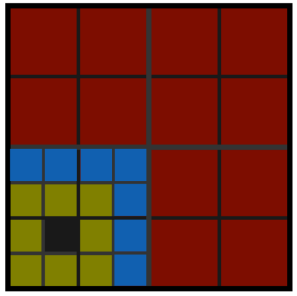

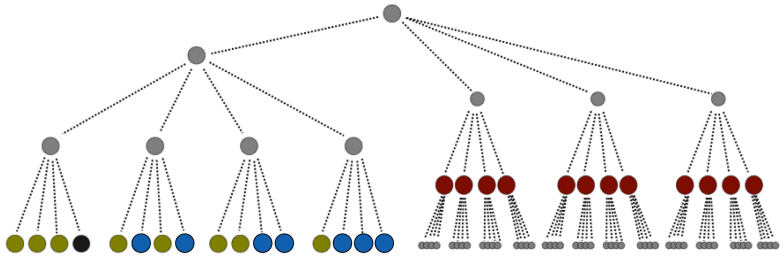

Hierarchical spatial decomposition gives well -separated clusters of charges

- Multipole expansions approximate source distributions
- Local expansions approximate resulting potential fields

Greengard+Rokhlin (1987)

The FMM in One Dimension

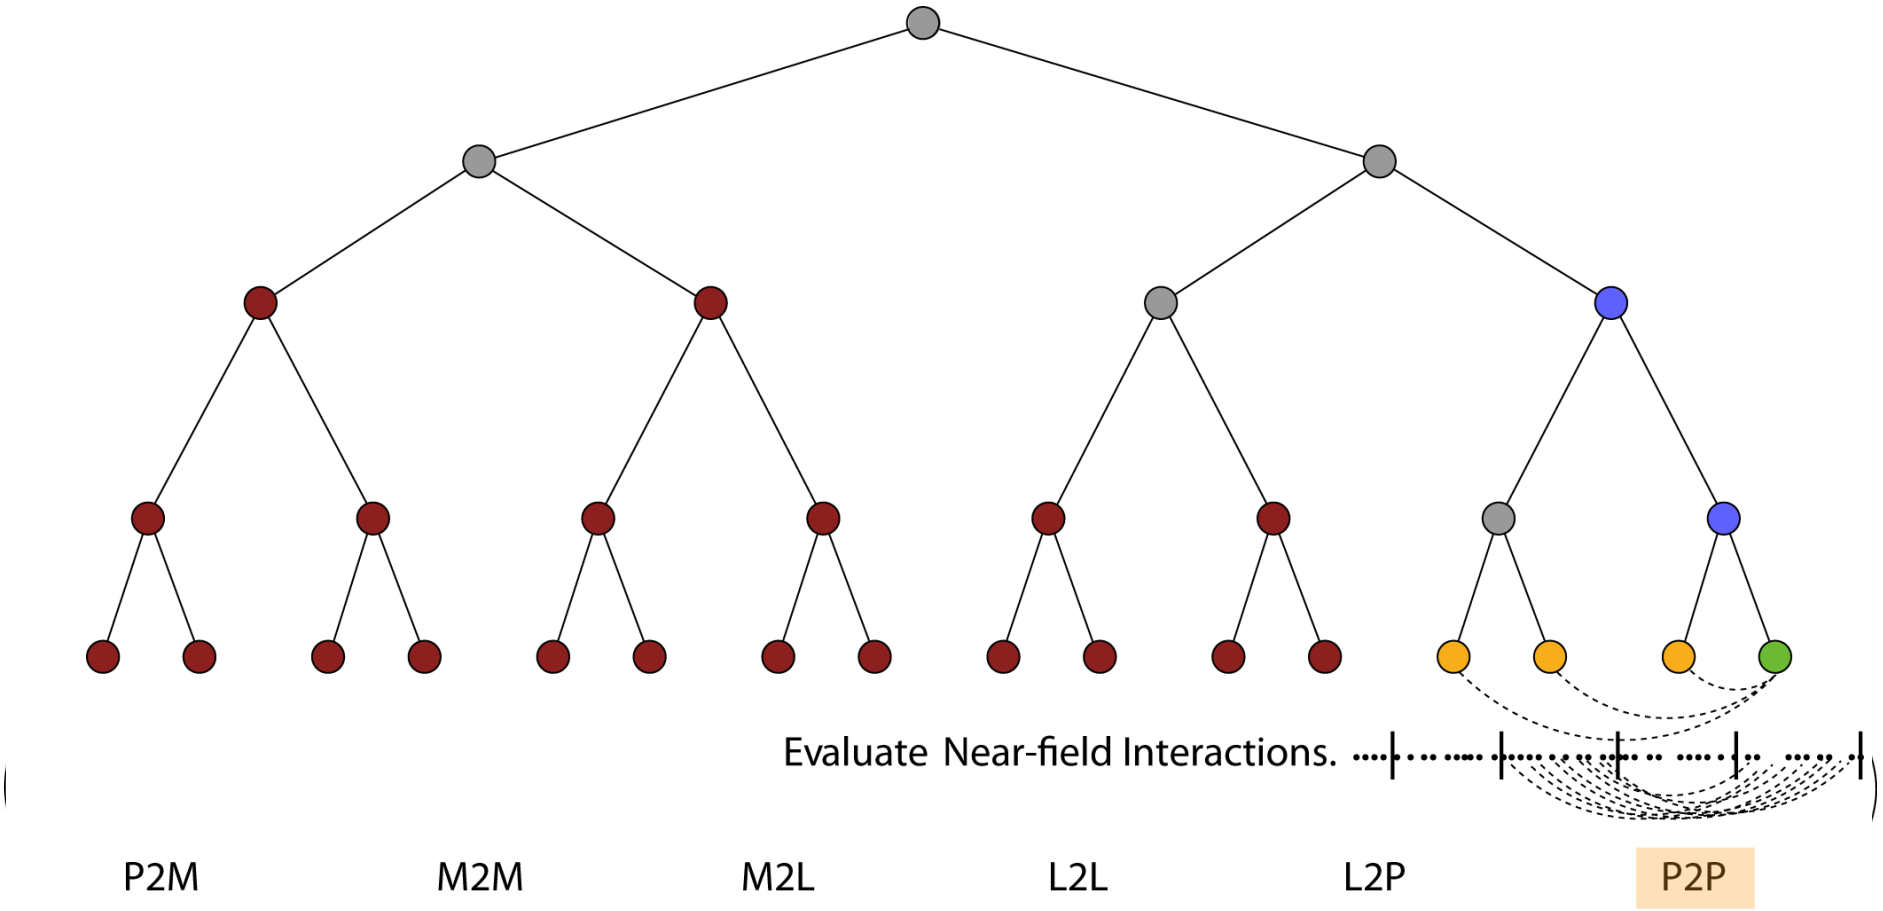

- "Natural" to combine computation with traversal

Greengard+Rokhlin (1987)

Modifications for BEM

- Fast multipole method, etc., are generally built around interacting “point” sources:

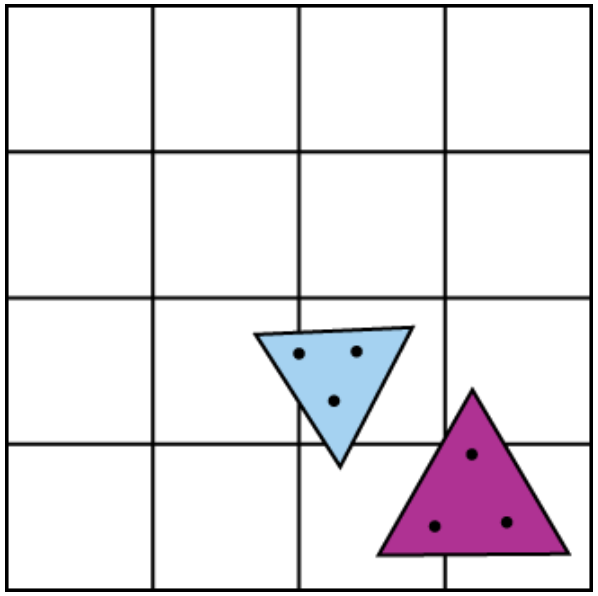

Potentials at all points are accurately computed…

BUT as we have seen, the representation of source distributions as point charges is NOT accurate!

- Different ways to address this:
  - Easy: add a sparse “local correction” matrix
  - Pretty, but difficult: compute multipole coefficients directly from basis functions

PetFMM: Open-source GPU FMM

- **Separate** tree traversal and computation

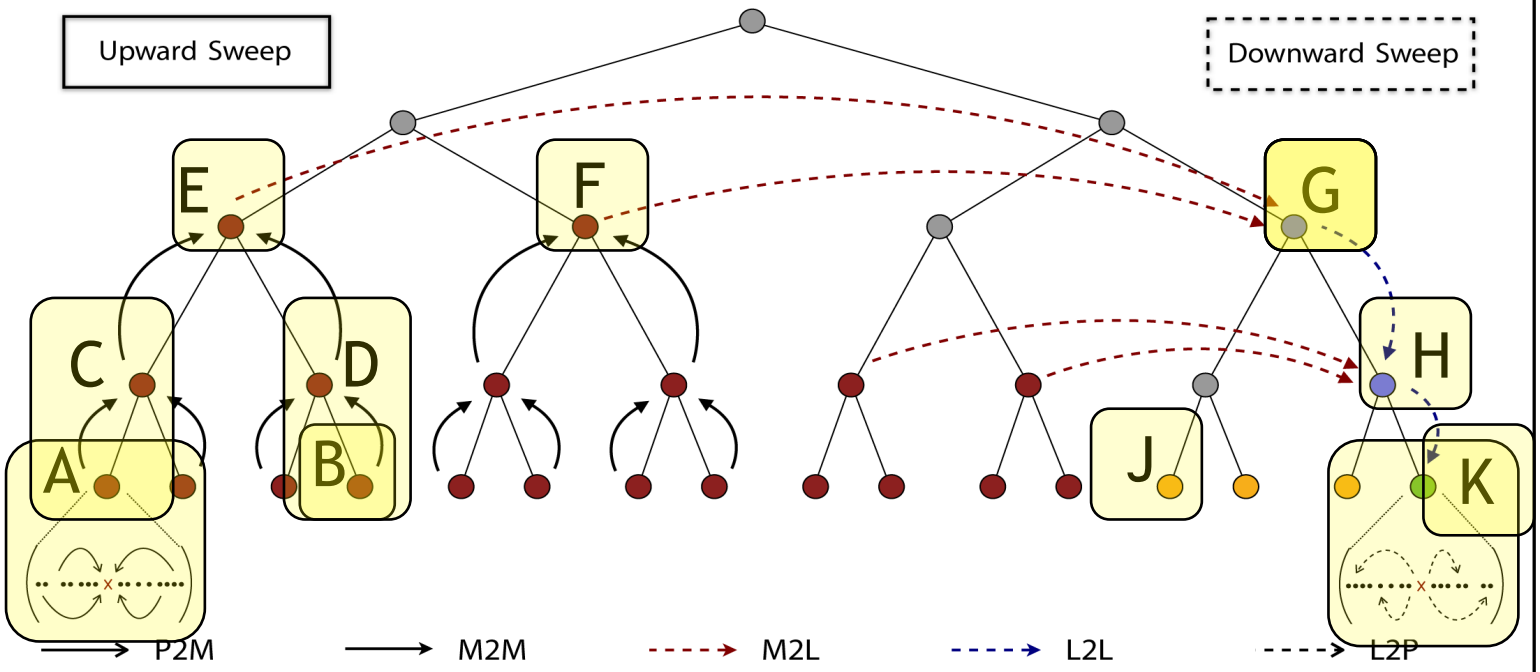

Queue

- P2M (A)
- P2M (B)
- M2M (A,C)
- M2M (B,D)
- M2L (E,G)
- M2L (F,G)
- L2L (G,H)
- L2P (K)
- P2P(J,K)

- Use a data-aware queuing system
  - Can handle dependencies if necessary
  - Non-FMM tasks can also be queued

http://barbagroup.bu.edu/Barba_group/PetFMM.html

PetFMM: Open-source GPU FMM

- **Queuing** improves memory access efficiency

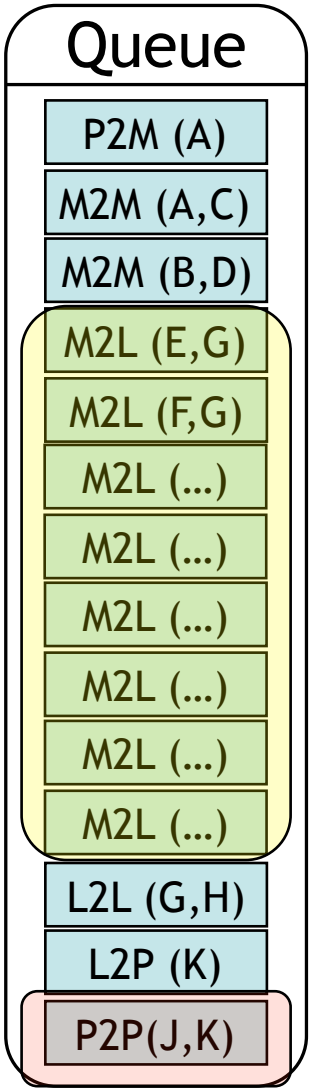

Bottleneck: small, separated data transfers for M2Ls

Solution: Batch all M2Ls associated with a given "target"

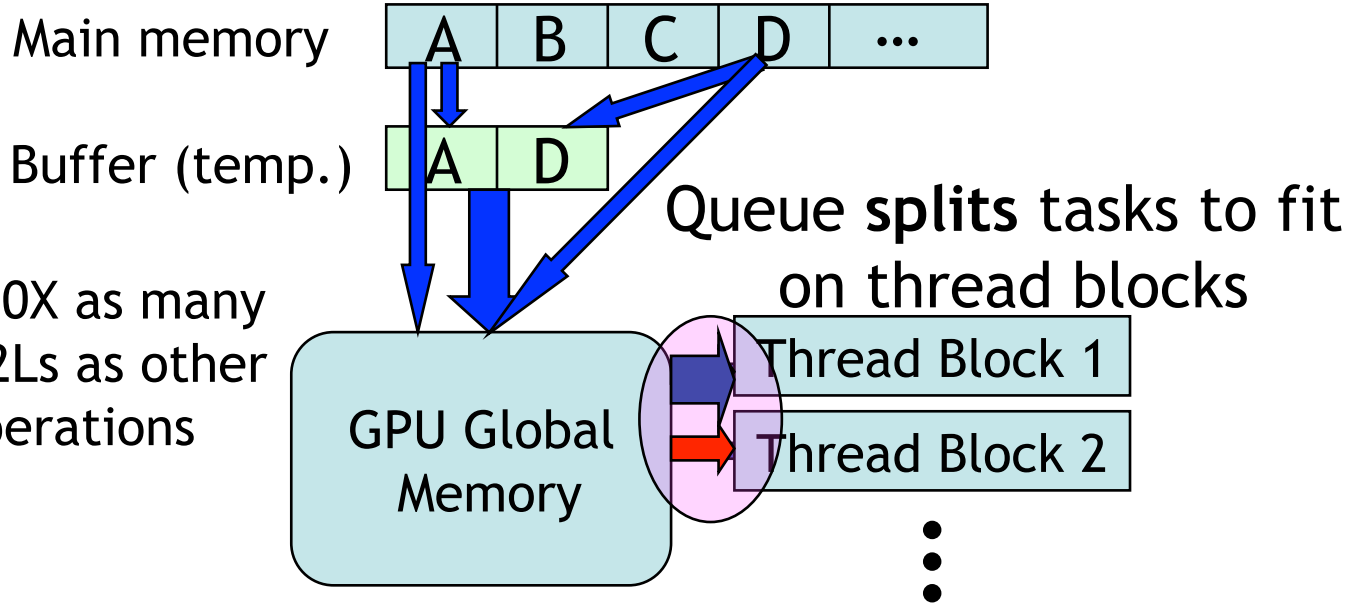

GPU: Can give 10X speedup on downward pass!
CPU: Less important but still advantageous

http://barbagroup.bu.edu/Barba_group/PetFMM.html

PetFMM Performance: CPU and GPU

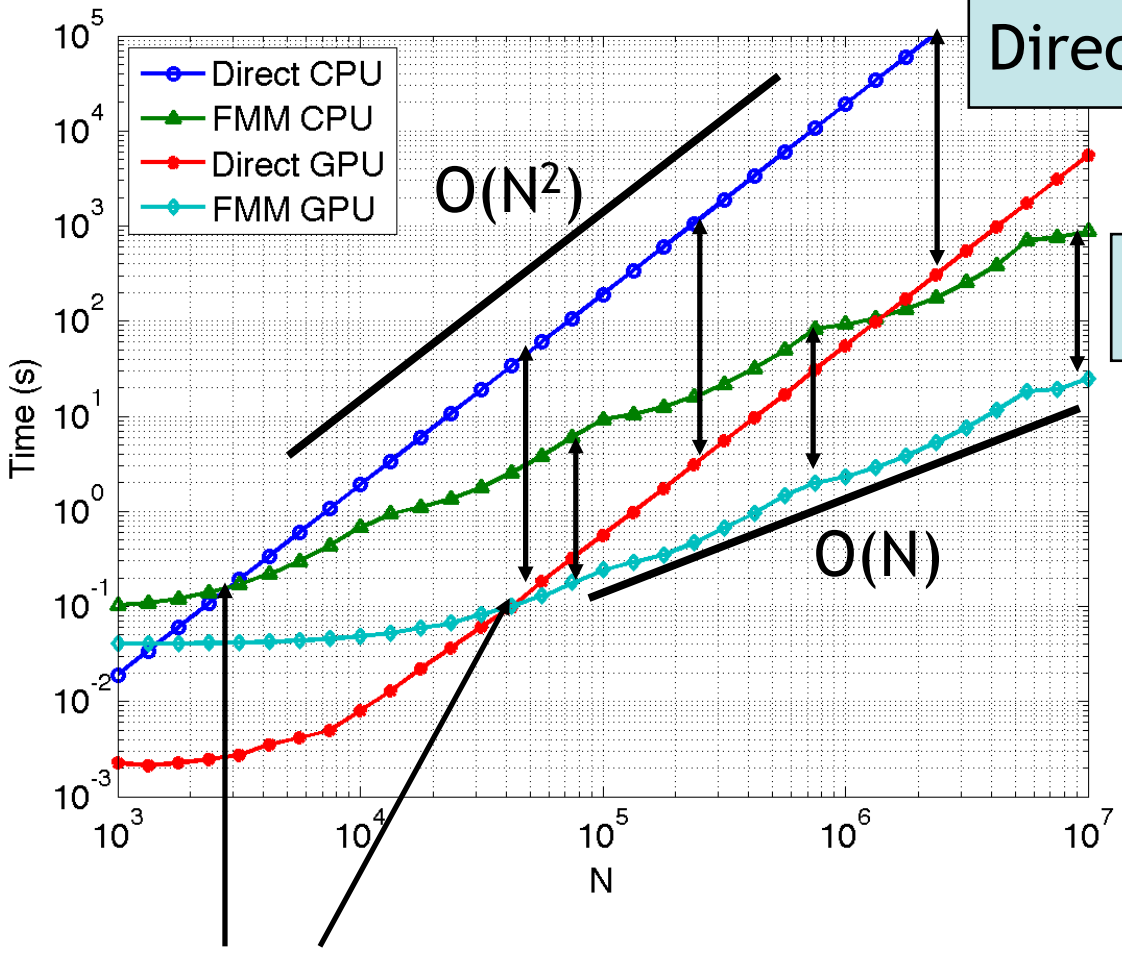

Direct: $T_{CPU} > 200\ T_{GPU}$

FMM: $T_{CPU} \sim 20\text{-}30\ T_{GPU}$

Cross-over for GPU ~ 10X cross-over for CPU

Parallelization: Graph Partitioning

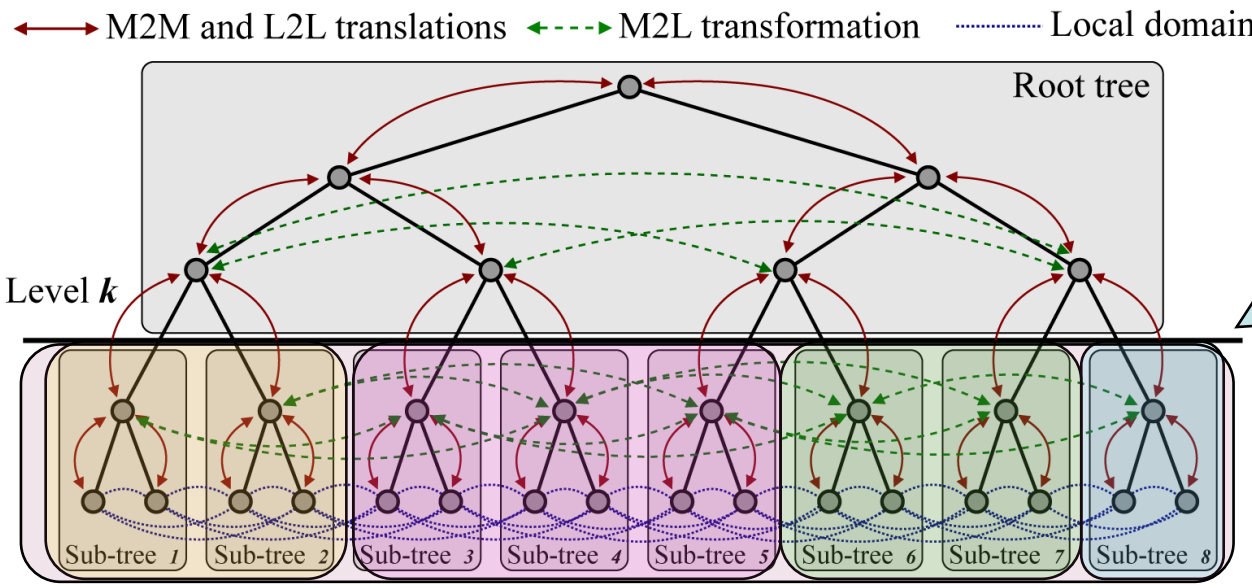

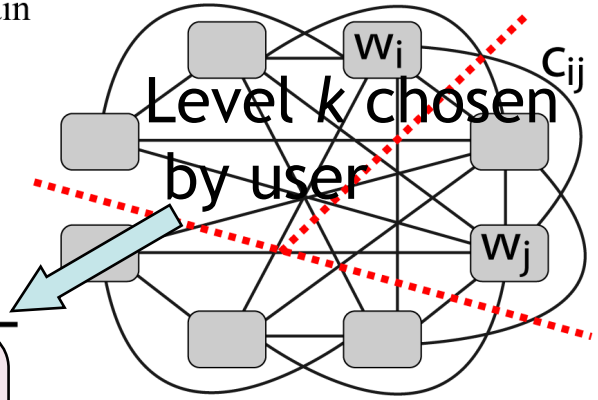

C_{ij}: communication cost
w_i: computation cost

- Provably good partitioning by ParMETIS
- Partitioning is fast: < 1% of total cost on CPU
- *Complete* re-use of serial code!

Cruz, Knepley, Barba arXiv:0905.2637

PetFMM Scales to Hundreds of GPUs

- 760-node GPU cluster (T. Hamada, Nagasaki)

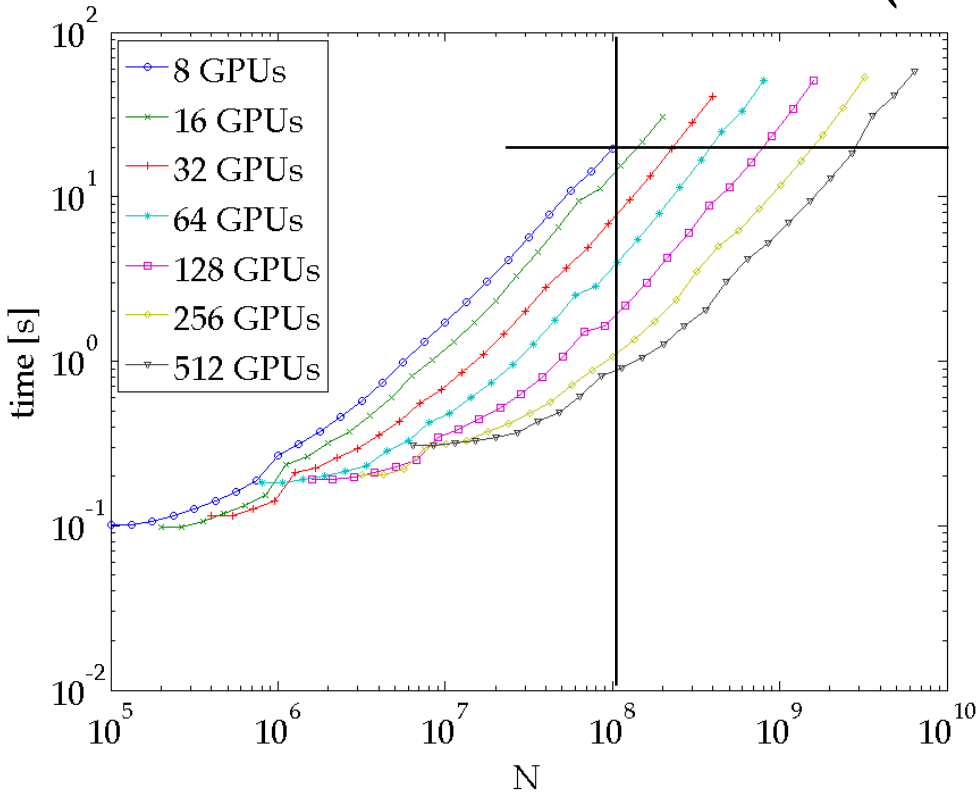

N=10^8: $\frac{\mathrm{T}(512\mathrm{GPU})}{\mathrm{T}(8\mathrm{GPU})} \approx 20$

In 20 sec:

$$8\ \mathrm{GPU} \rightarrow N \approx 1 \times 10^8$$
$$512\ \mathrm{GPU} \rightarrow N \approx 3 \times 10^9$$

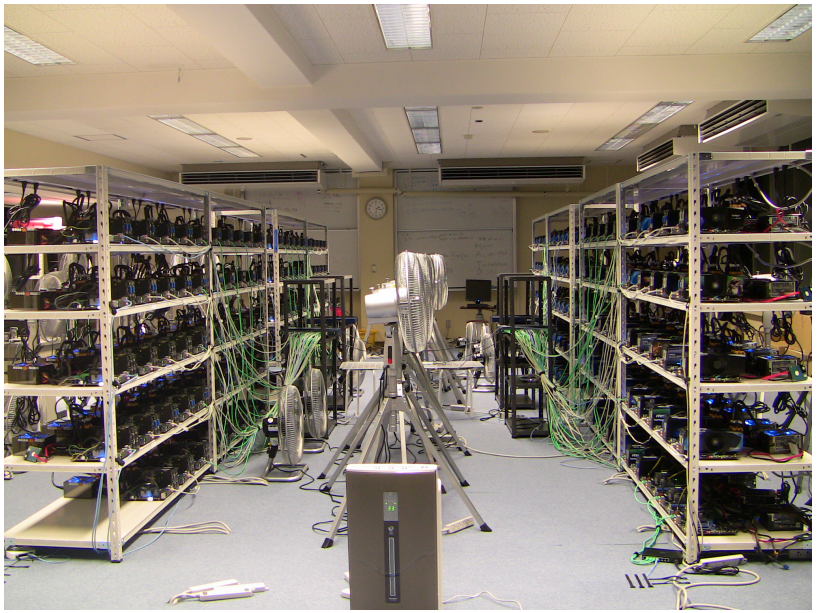

Cost of cluster: ~ US $420,000

Sustained: 34.6 Tflops

Performance/price: 80 Mflops/$

PetFMM Scaling on Modest GPU Clusters

- 100 million point charges, uniformly distributed

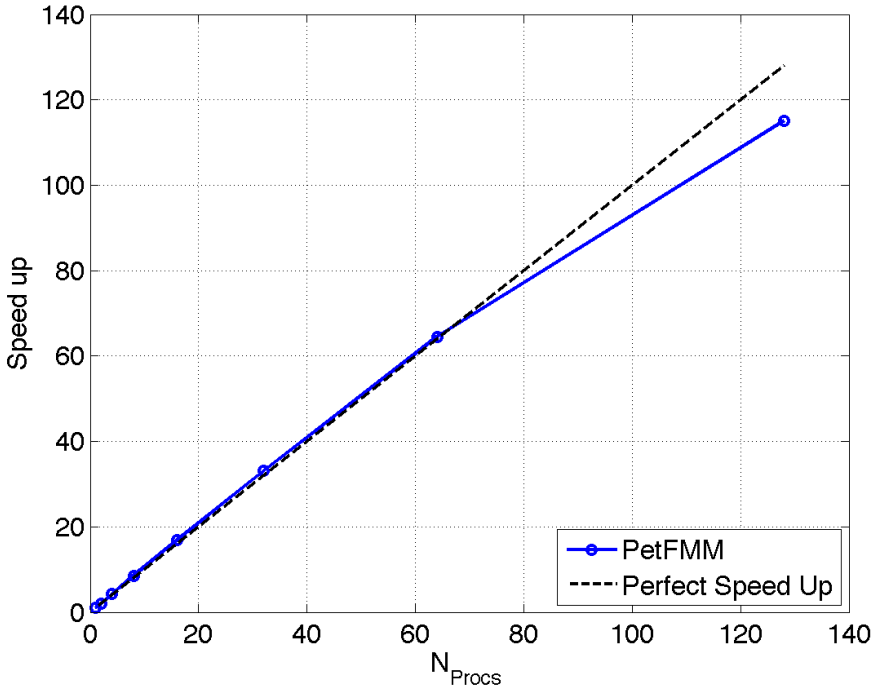

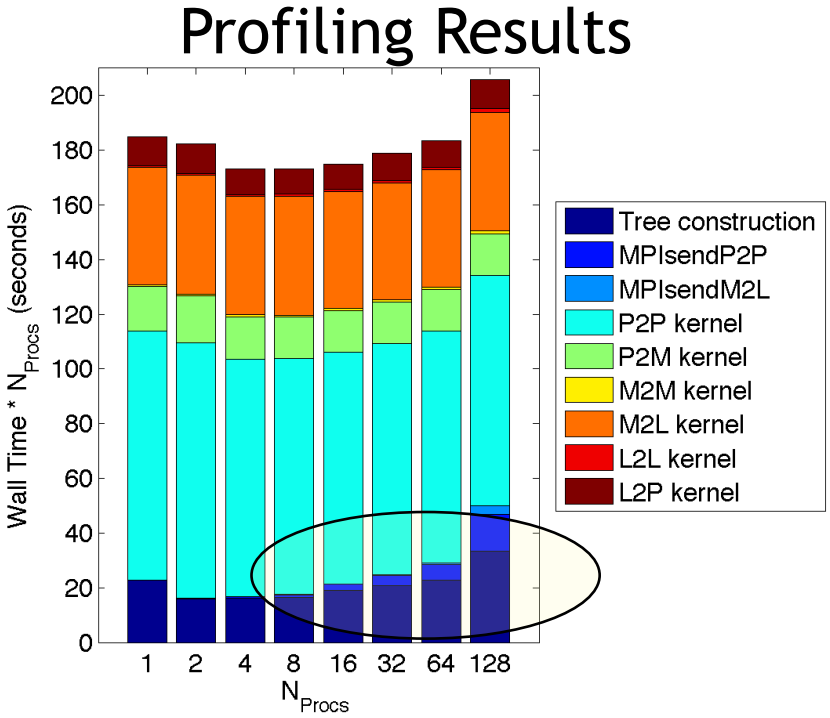

- Near perfect scaling up to 64 GPUs!
- Further optimizations are in progress

Pre-corrected FFT Algorithm

- Inspired by particle-particle particle-mesh (P^3M)
- O(N log N) but competitive in speed with fast multipole
- Algorithm is ***KERNEL INDEPENDENT***
  Laplace, Helmholtz, others…

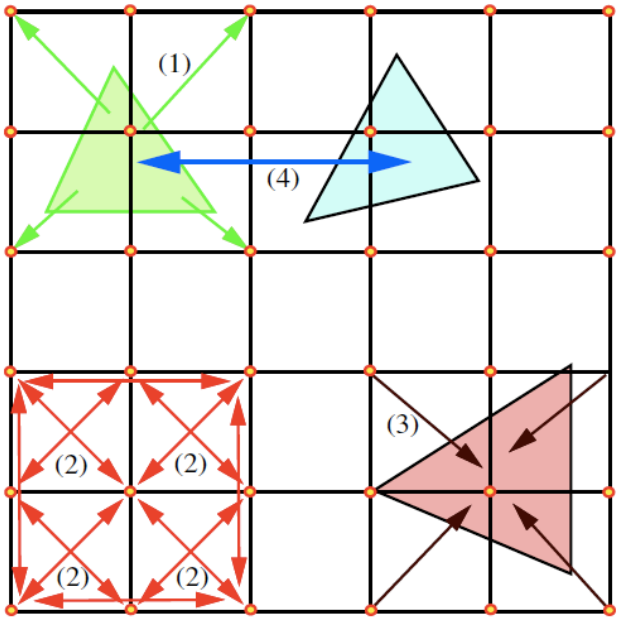

1. Project charges to grid
2. FFT convolution of kernel with grid sources
3. Interpolate grid potentials
4. “Pre-correct” so that local interactions are accurate

Phillips and White (1997)

Applications of pFFT

Proteins

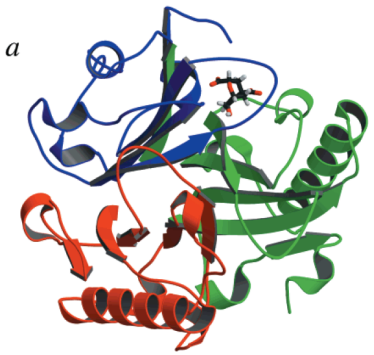

FastAero: coupled pFFT/tree code

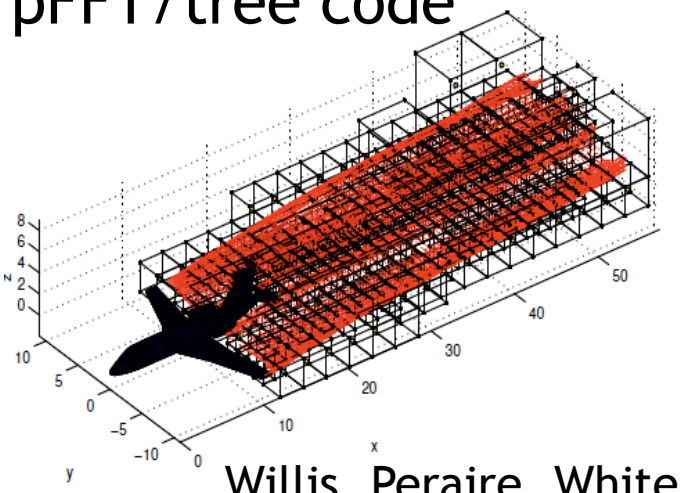

Willis, Peraire, White

Circuit Simulation

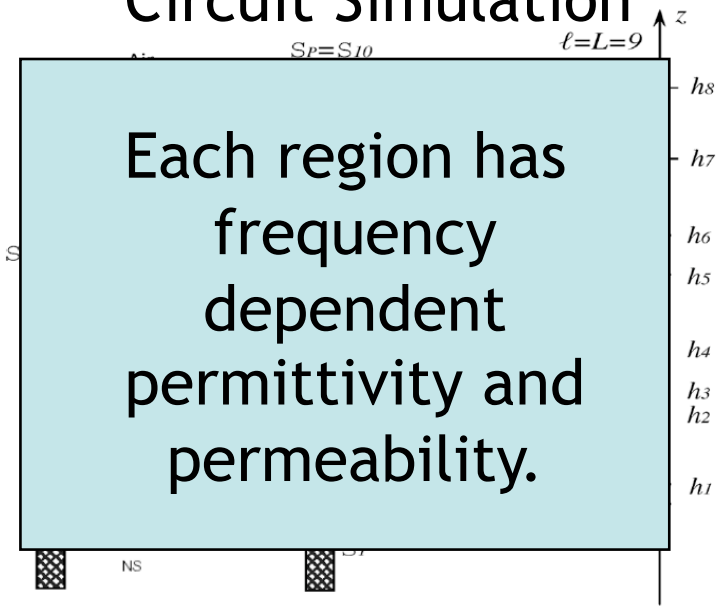

Each region has frequency dependent permittivity and permeability.

Fig. 1. Cross section of a nine-layer technology process ($\ell = 1, \ldots, L, L = 9$) featuring six metal layers ("NS", "Rx", "M1", "MT", "LY", and "M4") and four via layers ("CA", "V1", "VY", "AV"). The total surface enclosing the interconnect volume at each metal/via layer is denoted $\mathbf{s}_p, p = 1, \ldots, P; P = 10$.

Cadence Design Systems

EEG and MEG modeling

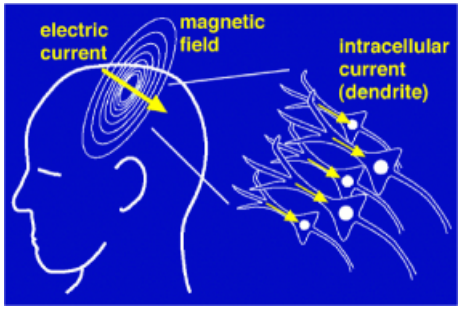

Summary:

- Many problems in potential theory can be rewritten from PDE form to BIE form

- BIEs tend to be most advantageous for *exterior* problems or problems with highly irregular boundaries

- Boundary-integral operators map *surface source distributions* to *surface potentials and fields*

References:

1. Atkinson, K. E. “The Numerical Solution of Boundary-Integral Equations of the Second Kind.” The authoritative book.

2. Atkinson and Han. “Theoretical Numerical Analysis.” Covers BEM in much less detail but it also introduces the theoretical background in good detail and also analyzes FDM and FEM.

3. McLean, W. “Strongly Elliptic Systems and Boundary Integral Equations.” First chapter has a nice historical account of integral equations and offers a more general framework than Atkinson.

4. Hsiao, G. C. and Wendland, W. L. “Boundary Element Methods: Foundations and Error Analysis.” Encyclopedia of Computational Mechanics, pp 1-54, 2004.

5. Hildebrandt, A. PhD thesis.
   http://scidok.sulb.uni-saarland.de/volltexte/2007/1400/