Fast Methods in Scientific Computation — Review

Techniques covered — quadrature:

We discussed methods for numerically evaluating integrals such as:

$$I = \int_\Omega f(x)\, dx$$

where f is a given function, and Ω is a subset of $\mathbb{R}^d$.

When f is smooth and $\Omega = [a, b]$, *very* efficient techniques exist (e.g. Gaussian quadrature). These work when we are free to pick the quadrature nodes.

Complications not covered in class.

- High dimensions.
- Irregular domains.
- Integrand f that is not smooth (or worse, singular).
- You may not be free to pick the quadrature points.

Very well-studied subject!
For any particular application, just do some reading.

Techniques covered — Fourier methods:

Thanks to the FFT, Fourier transforms can be evaluated very rapidly computationally.

Applications:

- Constant coefficient linear PDEs "without" boundary conditions:
  - Free space problems.
  - Rectangle with periodic boundary conditions.
- Signal and image processing.
- Spectral analysis.
- Multiplication of polynomials and large numbers.
- ...

Techniques covered — Fourier methods:

Thanks to the FFT, Fourier transforms can be evaluated very rapidly computationally.

Topics not covered in class (or covered very briefly):

- Sampling and connection to continuous Fourier transform.
- *Non-uniform FFT* — points in physical and/or Fourier space are not placed on a regular grid.
- FFT-like algorithm for other unitary transforms:
  - Wavelets.
  - Hadamard transform.
  - Expansions in special functions (Legendre polynomials, spherical harmonics, etc).
- Spectral methods for solving PDEs on a variety of domains.

Techniques covered — fast methods for N-body problems:

Problem formulation:

We are given points $\{x_i\}_{i=1}^N$ in $\mathbb{R}^d$.

We are given charges $\{q_i\}_{i=1}^N$.

We are given a kernel $k(x, y)$.

We seek to evaluate the potentials

$$u_i = \sum_{j=1}^{N} k(x_i, x_j)\, q_j \qquad i = 1,\, 2,\, 3,\, \ldots,\, N.$$

The *Fast Multipole Method* by Rokhlin and Greengard executes this task in $O(N)$ operations for a broad class of problems. ($d = 1,\, 2,\, 3$)

Also methods based on expansions in Gaussians $\rightarrow$ Beylkin lectures.

Techniques covered — fast methods for N-body problems:

$$u_i = \sum_{j=1}^{N} k(x_i, x_j)\, q_j \qquad i = 1,\, 2,\, 3,\, \ldots,\, N.$$

- $k(x, y) = \log|x - y|$ — Laplace equation in $\mathbb{R}^2$. *The case we studied.*
- $k(x, y) = 1/|x - y|$ — Laplace equation in $\mathbb{R}^3$.
- $k(x, y) = e^{-k|x-y|}/|x - y|$ — "screened" Laplace potential.
- $k(x, y) = e^{ik|x-y|}/|x - y|$ — Helmholtz equation in $\mathbb{R}^3$.
- $k(x, y) = e^{-|x-y|^2}$ — Gauss kernel.
- Fundamental solutions for Stokes, elasticity, time-harmonic Maxwell, *etc.*

Not covered in class: different kernels, geometries, hardware ... *But:*

You can now solve these problems.

You can (and should) almost always use "black box" codes in practice!

Techniques covered — fast methods for N-body problems:

When combined with quadrature techniques, the FMM can be used to solve free space problems for elliptic PDEs. For example, consider

$$-\Delta u(x) = f(x), \qquad x \in \mathbb{R}^3,$$

where f is a function supported on some domain and we impose appropriate decay conditions at infinity. The exact solution is:

$$u(x) = \int_{\mathbb{R}^3} \frac{1}{4\pi|x-y|} f(y)\, dy.$$

The integral can be evaluated via the FMM.

The technique described works for most constant-coefficient elliptic PDEs.
It can be modified to work with periodic boundary conditions on a rectangle, and certain other "simple" boundary conditions.

Techniques covered — superposition principle for linear problems:

We seek to solve

$$A\,u = g \qquad \text{on } \Omega, \tag{1}$$

$$u = f \qquad \text{on } \Gamma, \tag{2}$$

for some complication domain Ω with boundary Γ.

Suppose that we can solve

$$A\,v = g \qquad \text{on } \Psi,$$

for some simple domain Ψ such that $\Omega \subseteq \Psi$. (If Ψ is a rectangle, the FMM and the FFT work; if $\Psi = \mathbb{R}^d$, the FMM works.) Then set

$$w = u - v$$

and observe that w satisfies

$$A\,w = g \qquad \text{on } \Omega, \tag{3}$$

$$w = f - v \qquad \text{on } \Gamma. \tag{4}$$

You can handle the body load and the boundary condition separately.

Techniques covered — Boundary Integral Equations:

Suppose that we are interested in solving the Boundary Value Problem

$$A\,u = 0 \qquad \text{on } \Omega, \tag{5}$$

$$u = f \qquad \text{on } \Gamma, \tag{6}$$

where A is a linear elliptic partial differential operator with constant coefficients, where Ω is a domain in $\mathbb{R}^d$ with boundary Γ, and where f is a pre-scribed function on Γ. It is in many cases possible to reformulate a BVP such as (5), (6) as a boundary integral equation of the form

$$v(x) + \int_\Gamma k(x,y)\, v(y)\, dl(y) = g(x), \tag{7}$$

where k is a *kernel function* derived from the fundamental solution associated with the operator A, and g is a data function derived from the boundary data f.

Techniques covered — Boundary Integral Equations:

Example: We consider the BVP

$$-\Delta u = 0 \qquad \text{on } \Omega, \tag{8}$$

$$u = f \qquad \text{on } \Gamma, \tag{9}$$

where Ω is a simply connected open set in $\mathbb{R}^2$ with smooth boundary Γ, and where f is a prescribed function on Γ. We seek to rewrite (8, 9) as an equation defined on Γ only. To this end, we make the Ansatz

$$u(x) = \int_\Gamma \frac{n(x') \cdot (x - x')}{2\pi |x - x'|^2} \, \sigma(x') \, dl(x'), \qquad \text{for } x \in \Omega, \tag{10}$$

where $n(x')$ is the unit outwards pointing normal of Γ at x', and where σ is an unknown function on Γ. Note that for any σ, the function u defined by (10) satisfies (8).[a] Equation (9) is satisfied if and only if the function u defined by (10) satisfies

$$\lim_{y \to x,\, y \in \Omega} u(y) = f(x), \qquad \forall\, x \in \Gamma. \tag{11}$$

A simple evaluation of the limit (11) yields

$$-\frac{1}{2}\sigma(x) + \int_\Gamma \frac{n(x') \cdot (x - x')}{2\pi |x - x'|^2} \, \sigma(x') \, dl(x') = f(x), \qquad \text{for } x \in \Gamma. \tag{12}$$

In a strong sense, the BVP (8, 9) is equivalent to the BIE (12).

[a] Conversely, it is known from classical potential theory that any solution of (8, 9) can be written in the form (10) for some function σ.

Techniques covered — Boundary Integral Equations:

Benefits:

1. The *condition number* of the numerical problem is typically similar to the condition number of the actual "physics" of the problem. This is in contrast to discrete systems arising from discretizations of differential equations, which typically have very large condition numbers.

2. The *computational complexity is asymptotically optimal* in the sense that the number of floating point operations required scales linearly with the true complexity of the problem. (The data provided is defined on a set of dimension $d-1$ so we should not have to discretize a set of dimension d.)

3. The construction of discretization schemes that are both stable and of *high order* is somewhat easier in the integral equation environment than it is in the differential equation environment.

4. Handling *exterior problems* where the domain Ω is the region outside some given contour or surface Γ is much easier in the BIE environment than in the PDE setting.

Techniques covered — Boundary Integral Equations:

Drawbacks:

1. You get dense linear systems. This complication can be handled (e.g. via the FMM) but it does take effort.

2. Less versatile technique — multiphysics, non-standard equations, unusual boundary conditions, non-linearity.

3. Integral equations *can* be used for problems with variable coefficients, but are generally less competitive in this environment as they no longer gain from dimension reduction. (Integral equations can still be a good idea in specialized environments.)

4. Much less infrastructure! Codes are not available to the same extent.

Techniques covered — Boundary Integral Equations:

In class, we covered the case of Laplace's equation in $\mathbb{R}^2$.

Conceptually, a broad class of equations can be solved using analogous techniques:

- Equations of elasticity.
- Stokes equations.
- Helmholtz.
- Time-harmonic Maxwell.
- Yukawa equations.
- ...

Techniques covered — Boundary Integral Equations:

Caveats:

- Problems in $\mathbb{R}^3$ are *much* harder:
  - Representing the geometry is hard.
  - Constructing high-order quadratures on surfaces is hard.
  - Corners in *curves* are slightly difficult to manage.
    Corners and edges in *surfaces* are very hard (if you want high order).
  - 3D FMMs are significantly slower than 2D versions.

  (Of course, 3D is challenging to all methods so the competition is weaker.)

- Constructing a good BIE formulation requires specialized knowledge.
  For standard problems, you can often find the right formulation in the literature, but you need to dig a bit.

- Problems involving scattering cause substantial difficulties (ill-conditioning due to resonances, trickiness in constructing FMMs, lots of points needed, etc). However, scattering is sometimes almost impossible to solve without integral equations.

Techniques covered — Krylov methods:

Problem formulation:
Suppose that A is a given square matrix and that you either want to solve

$$Ax = b$$

or compute a partial spectral decomposition of A. We assume that:

- A is very large.
- A can rapidly be applied to vectors (A is sparse, "FMM-able," etc).

The idea of a *Krylov* method is to study the vector space

$$\mathcal{K}_k(A, b) = \text{Span}\{b,\ A\,b,\ A^2\,b,\ A^3\,b,\ \ldots,\ A^{k-1}\,b\}.$$

- For solving $x = A^{-1}b$ we seek the best possible solution within $K_k(A, b)$:
  - $x_k = \text{argmin}\{||A^{-1}b - z|| :\ z \in \mathcal{K}_k(A, b)\}$. (GMRES)
  - $x_k = \text{argmin}\{||A^{-1}b - z||_A :\ z \in \mathcal{K}_k(A, b)\}$. (CG)
  - ...
- For computing eigenvectors of A, let b be a random vector, project A onto $\mathcal{K}_k(A, b)$ and compute the eigen-decomposition of this reduced matrix. (Even better, pick b as a block of ℓ random vectors!)

Techniques covered — Krylov methods:

	A symmetric	A general
Linear system solve	Conjugate gradients (A is spd)	GMRES
Eigendecomposition	Lanczos	Arnoldi
	(three-term recurrence)	(k-term recurrence)

Techniques covered — Krylov methods:

We briefly discussed techniques for analyzing the convergence properties of Krylov methods. Most such methods can be interpreted as *polynomial approximation problems.* The idea is to minimize some polynomial $p(A)$ applied to b.
As a rule of thumb, Krylov method work well when:

- A is close to normal.
- A has a clustered spectrum / singular spectrum.

If A is "bad" (which is common) one looks for an accessible operator M such that $M^{-1}A$ is good. M is called a *pre-conditioner.*

Krylov methods are very common, and very powerful — standard tool of high-performance computing. Using them efficiently is part science, part "art" (=engineering, tricks, intuition, ...).

For practical use, note that while analyzing and constructing Krylov methods is quite complicated, most practitioners do not need to dig into the nitty gritty details: Simply do a brief literature survey to see which method might suit your problem and implement it, or, (strongly recommended) use a canned routine.

Techniques covered — (randomized) low-rank approximation:

Problem formulation:
Given an $n \times n$ matrix A, for a very large n (say $n \sim 10^6$),
we seek to compute a rank-k approximation, with $k \ll n$ (say $k \sim 10^2$ or 10^3),

$$\underset{n \times n}{\mathsf{A}} \approx \underset{n \times k}{\mathsf{E}} \underset{k \times n}{\mathsf{F}^*} = \sum_{j=1}^{k} \boldsymbol{e}_j \, \boldsymbol{f}_j^*.$$

Solving this problem leads to algorithms for computing:

- Singular Value Decomposition (SVD) / Principal Component Analysis (PCA).
  (Require $\{\boldsymbol{e}_j\}_{j=1}^k$ and $\{\boldsymbol{f}_j\}_{j=1}^k$ to be orthogonal sets.)

- Finding spanning columns or rows.
  (Require $\{\boldsymbol{e}_j\}_{j=1}^k$ to be columns of A, or require $\{\boldsymbol{f}_j^*\}_{j=1}^k$ to be rows of A.)

- Determine eigenvectors corresponding to leading eigenvalues.
  (Require $\boldsymbol{e}_j = \lambda_j \, \boldsymbol{f}_j$, and $\{\boldsymbol{f}_j\}_{j=1}^k$ to be orthonormal.)

- *etc*

Techniques covered — (randomized) low-rank approximation:

Idea of randomized sampling:

- Pick a random matrix Ω.
- Form a sample matrix $\mathsf{Y} = \mathsf{A}\,\Omega$.
- Orthonormalize the columns of Y to form $\mathsf{Q} = \texttt{orth}(\mathsf{Y})$.
- Project A onto the space spanned by the columns of Q to form $\mathsf{B} = \mathsf{Q}^*\mathsf{A}$.

End result: $\mathsf{A} \approx \mathsf{Q}\,\mathsf{B}$

Post processing yields SVD, eigenvalue decomposition, spanning columns, etc.

For improved accuracy, use variations of a power method: $\mathsf{Y} = (\mathsf{A}\,\mathsf{A}^*)^q\,\mathsf{A}\,\Omega$

Techniques covered — (randomized) low-rank approximation:

Benefits:

- Loosens communication constraints — parallelism, multicore, etc.
- Single pass methods possible.
- Can reduce complexity from $O(mnk)$ to $O(mn\log(k))$.

Primary application: Computing the SVD of a very large dense matrix (say of size $10^5 \times 10^5$) stored on disk.

Techniques covered — (randomized) low-rank approximation:

Note that this set of techniques is still slightly boutique.

They were included in class to bring attention to the power of *randomization* and as a pretext for discussing the importance of communication costs in an important standard environment.

(And, of course, to bring attention to research in our department — see June issue of *SIAM Review*!)

Brief review of different PDEs:

Concepts such as BIEs, FMM, FFT, etc, are useful for many of the standard equations:

- *Covered in class:*
  Laplace

- *Minor modifications required:*
  Elasticity, Stokes, Yukawa, low frequency Helmholtz

- *Significant modifications required:*
  Medium frequency Helmholtz and time-harmonic Maxwell.

- *Equations for which the techniques can form **a part** of a solver:*
  Heat equation. Convection equations. Wave equation in time domain. Navier-Stokes. (Note: much new physics that was not present for stationary elliptic problems.)