18.336/6.335 Fast Methods for Partial Differential and Integral Equations Spring 2013

February 7-14, 2013

Instructors: Laurent Demanet (laurent@math.mit.edu), Jacob White (white@mit.edu)
TA: Richard Zhang (ryz@mit.edu)
Office hours: Wednesday 2-4pm, room 2-392.

3 Fourier based fast methods for Ordinary Differential Equations

One convenient set of eigenfunctions for the Laplacian operator $\Delta = \nabla \cdot \nabla$ in $\mathbb{R}^m$ are the plane waves[1][2]:

$$\Delta \underbrace{e^{i2\pi\xi\cdot x}}_{\text{plane wave}} = -\left(2\pi|\xi|\right)^2 e^{i2\pi\xi\cdot x}$$

Consider the spectral decomposition of a matrix L, which introduces the transformation matrix P satisfying the following relation:

$$LP = P\Lambda.$$

The diagonal matrix Λ contains the eigenvalues of the matrix L. One important property of diagonalizable matrices is that any function g applied to the matrix L is equivalently applied to the eigenvalues λ under spectral decomposition:

$$g(L) = Pg(\lambda)P^{-1}.$$

In order to solve ODEs, we are often interested in the the inverse operation. The inverse of L can be computed under spectral decomposition according to the following:

$$L^{-1} = P\text{diag}(\lambda^{-1})P^{-1}.$$

Where the transformation matrix P is constructed using plane waves as the underlying eigenfunctions, P^{-1} is the Fourier transform:

$$\hat{f}(\xi) = \int e^{-i2\pi\xi\cdot x}\, f(x)\, \mathrm{d}x,$$

and P is the inverse Fourier transform[3]

$$f(x) = \int e^{i2\pi\xi\cdot x}\, \hat{f}(\xi)\, \mathrm{d}\xi.$$

Therefore, the inverse of the continuous Laplacian operator is the integral equation:

$$\Delta^{-1} f(x) = \int e^{i2\pi\xi\cdot x} \frac{-1}{4\pi^2|\xi|^2} \hat{f}(\xi)\, \mathrm{d}\xi.$$

As illustrated above, diagonalization via fast spectral decomposition is the core philosophy behind Fourier based fast methods. In practice, there are additional restrictions associated with the functions f and $\hat{f}$, associated with the solvability of the Poisson problem.

[1] Technically, plane waves are *generalized* eigenfunctions because they are of infinite energy

[2] Other eigenfunctions of Δ also exist, e.g. Bessel functions, spherical harmonics. Each set of eigenfunctions provides a spectral decomposition of Δ.

[3] Aside: if $L = L^*$ where $*$ denotes the Hermitian adjoint (transpose + complex conjugate), the matrix P is unitary: $P^{-1} = P^*$.

3.1 The Discretized Poisson Problem

Consider the simple ODE form (i.e. one-dimensional form) of Poisson's equation:

$$-\frac{d^2}{dx^2}u = f, \quad x \in [0,1].$$

The equation is usually complemented with one of three common boundary conditions:

1. Periodic boundary condition, where $u(0) = u(1)$.
2. Dirichlet boundary condition, where $u(0) = a$ and $u(1) = b$.
3. Neumann boundary condition, where $u'(0) = a$ and $u'(1) = b$.

In the following section we will discretize the interval of x into N points, thereby defining the discretized x_j:

$$x_j = jh, \quad h = \frac{1}{N}$$

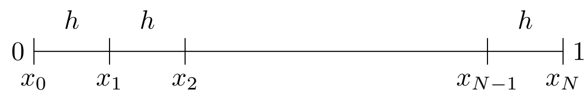

Let us define $u_j \approx u(x_j)$ in order to approximate the solution, and the discretized Laplacian matrix $\Delta_h \to -d^2/dx^2$, where:

$$(\Delta_h u)_j = \frac{2u_j - u_{j+1} - u_{j-1}}{h^2}.$$

Note that the *local truncation error* (LTE)[4] of this approximation is $\mathcal{O}(h^2)$. The properties of the Laplacian matrix are greatly affected by the associated boundary conditions!

1. Periodic boundary conditions. The periodic relation $u_0 = u_N$ is substituted into Δ_h, giving the following $N \times N$ discretized Laplacian matrix for the nodes $j = 0 \dots N-1$:

$$\Delta_h^{per} = \frac{1}{h^2}\begin{bmatrix} 2 & -1 & & & -1 \\ -1 & 2 & \ddots & & \\ & \ddots & \ddots & \ddots & \\ & & \ddots & 2 & -1 \\ -1 & & & -1 & 2 \end{bmatrix}.$$

The superscript "per" emphasizes the fact that this discretized Laplacian matrix is specific only to the periodic problem. Δ_h^{per} is symmetric, real, Toeplitz, positive semi-definite, but most importantly it is *Circulant.* A circulant matrix is where columns are cyclic shifts of one another. Circulant matrices are diagonalized by the DFT.

2. Dirichlet boundary conditions. Each boundary condition $u_0 = a, \quad u_N = b$ is substituted, giving the following $(N-1) \times (N-1)$ discretized Laplacian matrix for the nodes $j = 1 \dots N-1$:

$$\Delta_h^{D} = \frac{1}{h^2}\begin{bmatrix} 2 & -1 & & & \\ -1 & 2 & \ddots & & \\ & \ddots & \ddots & \ddots & \\ & & \ddots & 2 & -1 \\ & & & -1 & 2 \end{bmatrix}.$$

The matrix is no longer circulant, but it is now tridiagonal.

[4]For an explanation of LTE and its link to the global error, see Randall LeVeque's book on the finite differences method.

3. Neumann boundary conditions. The second-order accurate boundary conditions are $u_0 = u_{-1}, \quad u_N = u_{N+1}$, where u_{-1} and u_N are ficticious points introduced by extend the grid the distance of h in each direction. Substituting the second-order accurate central differences approximation:

$$\frac{d}{dx}u_0 \approx \frac{u_1 - u_{-1}}{2h}, \quad \frac{d}{dx}u_N \approx \frac{u_{N+1} - u_{N-1}}{2h},$$

into the discretized Laplacian eliminates the ghost points and yields the following $(N+1) \times (N+1)$ matrix for the nodes $j = 0 \dots N$:

$$\Delta_h^D = \frac{1}{h^2} \begin{bmatrix} 2 & -2 & & & \\ -1 & 2 & -1 & & \\ & \ddots & \ddots & \ddots & \\ & & -1 & 2 & -1 \\ & & & -2 & 2 \end{bmatrix}$$

The matrix is no longer Toeplitz, no longer symmetric, and it no longer makes sense to talk about positive-definiteness.

3.2 Diagonalization of the Periodic Problem

As described before, the one-dimensional plane wave $u_j : j = 0 \dots N-1$,

$$u_j = e^{i2\pi jk/N} = w^{jk}, \quad w = e^{i2\pi/N},$$

is an eigenvector of the discretized Laplacian Δ_h^{per}. It is easy to demonstrate this by applying the discretized Laplacian to u_j:

$$\begin{aligned} \left(h^2 \Delta_h^{per} u\right)_j &= 2w^{jk} - w^{(j-1)k} - w^{(j+1)k} \\ &= w^{jk}(2 - w^{-k} - w^k) \\ &= w^{jk}(2 - 2\cos\frac{2\pi k}{N}), \end{aligned}$$

since $w^{jk} = u_j$ re-emerges in the solution. The associated eigenvalue λ_k for u_j is:

$$\lambda_k = h^{-2}(2 - 2\cos\frac{2\pi k}{N}).$$

As long as k is discrete, $w^0 = w^{Nk} = 1$. The largest eigenvalue is:

$$\lambda_{N/2} = \frac{4}{h^2},$$

and further eigenvalues repeat according to $\lambda_{N-k} = \lambda_k$. This means the dimension of eigenspace is 2. Expanding out the eigenvector system, we obtain the IFT and DFT matrices:

$$F_{jk} = \frac{1}{\sqrt{N}} w^{jk}, \quad F_{jk}^{-1} = \frac{1}{\sqrt{N}} w^{jk}.$$

Because of this we can write:

$$\Delta_h^{per} = F \Lambda F^{-1}$$

where

$$\Lambda = \text{diag}(h^{-2}(2 - 2\cos\frac{2\pi k}{N})).$$

Substituting this relation into the original Poisson's equation yields:

$$
\begin{aligned}
F^*\Delta_{h,per}u &= F^*f \\
\Lambda F^*u &= F^*f \\
\lambda_k\hat{u}_k &= \hat{f}_k.
\end{aligned}
$$

Therefore, solving the periodic Poisson equation for u_j can be performed by obtaining $\hat{u}_k = \lambda_k^{-1}\hat{f}_k$ and inverting $\hat{u}_k$ with the inverse Fourier transform. However, consider the case when $k = 0$:

$$\lambda_0\hat{u}_0 = \hat{f}_0.$$

The eigenvalue λ_0 is zero, and the system has no solutions if $\hat{f}_0$ is anything other zero. This is the *compatibility condition* for the periodic Poisson problem. If $\hat{f}_0 = 0$, then $\hat{u}_0$ can be anything, but we set it to zero in order to obtain the least squares solution. This operation is then equivalent to the pseudo-inverse operation †, which is to invert everything in the range space of $\Delta_{h,per}$ and to set everything in the null space[5] to the least square solution of zero.

The fast Fourier algorithm for the periodic Poisson problem is the following:

1. Take the FFT,
2. Divide by λ_k for all non-zero eigenvalues, set to zero otherwise. Gives $\hat{u}_k$.
3. Take the IFFT gives back u_k.

3.3 Diagonalization of the Dirichlet Problem

Under the Dirichlet boundary conditions $u(0) = a, \quad u(1) = b$, the discretized Laplacian Δ_h^D is the following:

$$\Delta_h^D u = f, \quad \Delta_h^D = \frac{1}{h^2}\begin{bmatrix} 2 & -1 & & & \\ -1 & & & & \\ & & \ddots & & \\ & & & & -1 \\ & & & -1 & 2 \end{bmatrix}$$

The eigenvectors are sines of the plane waves (and not the cosines)[6]. The jth component of the k-th eigenvector is

$$v_k(j) = \sqrt{\frac{2}{N}}\sin(\frac{\pi jk}{N}),$$

and the associated eigenvalues are:

$$\lambda_k = \frac{1}{h^2}(2 - 2\cos\frac{\pi k}{N})$$

where $k = 1 \ldots N-1$. Because k goes from $1 \ldots k$, we have no singular eigenvalue. Instead of going through the cos twice, it only goes through it once. Eigenvalues are all positive, so Δ_h^D is positive definite.

Spectral decomposition

$$\Delta_h^D = S\Lambda S^*$$

where S is the inverse discrete sine transform and $S^*$is the discrete sine transform:

$$S_{jk} = \sqrt{\frac{2}{N}}\sin(\frac{\pi jk}{N})$$

S is its own transpose, its own inverse. This particular DST is known as DST-I. Fast unitary DST-I in $\mathcal{O}(N\log N)$ steps:

1. Do an odd extension, include the zeros on either side.
2. Apply length $2N$ FFT to the result.
3. Restrict $\hat{u_k}$ to $k = 1, \ldots, N-1$, and divide by $-i$.

[5] $[1, 1, \cdots, 1]$ spans the null space of $\Delta_{h,per}$

[6] The sines are a complete set, so they will represent all functions excluding boundary values, even the cosine function.

3.4 Diagonalization of the Neumann Problem

Here, the sum of columns is zero, and $[1, \dots, 1]^T$ spans the null space. The eigenvectors are all cosines:

$$v_k(j) = c_k \cos(\frac{\pi j k}{N}),$$

and the eigenvalues are:

$$\lambda_k = \frac{1}{h^2}(2 - 2\cos\frac{\pi k}{N}).$$

Therefore we can write:

$$\Delta_{h,N} = C\Lambda C^{-1}$$

where C is the inverse DCT (type I) and C^{-1} is the transpose. Note that $C^{-1} \neq C^*$.

Fast algorithm for type-I DCT:

- Do an even extension from $0 \dots N$ to $0 \dots 2N-1$.
- Perform length $2N$ FFT
- Restrict $\hat{u}_k$ to $k = 0 \dots N$ and divide by 2.

You can find a paper detailing everything on DST and DCT by Strang, SIAM reviews, 41-1, pp125-147 (1999)

3.5 Two-dimensional Poisson Equation on a Square

Now consider the two-dimensional, partial differential equations problem of the Poisson equation on a square:

$$\begin{aligned} -\Delta u &= f, \quad x \in \Omega = [0,1]^2 \\ u &= g, \quad x \in \partial\Omega \end{aligned}$$

Notation: N is the total number of grid points, n is the number of intervals along each direction. For the 2D Dirichlet problem, $N = (n-1)^2$.

The discretized Laplacian is the following:

$$(-\Delta_{h,D} u)_{j1,j2} = \frac{1}{h^2}[4u_{j1,j2} - u_{j1+1,j2} - u_{j1-1,j2} - u_{j1,j2+1} - u_{j1,j2-1}]$$

This is a second-order accurate approximation, meaning that the operator has a second order local trunction error. The variables must be placed in a vector with some rational order. If raster scan ordering is used, then the following structure will be seen in the matrix:

$$-\Delta_{h,D} = \frac{1}{h^2}\begin{bmatrix} T & -I & & & \\ -I & T & -I & & \\ & -I & T & \ddots & \\ & & \ddots & \ddots & -I \\ & & & -I & T \end{bmatrix} \in \mathbb{R}^{(n-1)^2 \times (n-1)^2}$$

where I is the identity matrix and T is the following matrix:

$$T = \begin{bmatrix} 4 & -1 & & & \\ -1 & 4 & -1 & & \\ & -1 & 4 & \ddots & \\ & & \ddots & \ddots & -1 \\ & & & -1 & 4 \end{bmatrix} \in \mathbb{R}^{(n-1)\times(n-1)}$$

Notice that if we define K to be the 1D discretized Laplacian:

$$K = \begin{bmatrix} 2 & -1 & & & \\ -1 & 2 & -1 & & \\ & -1 & 2 & \ddots & \\ & & \ddots & \ddots & -1 \\ & & & -1 & 2 \end{bmatrix} \in \mathbb{R}^{(n-1)\times(n-1)}$$

Then we can write the 2D discretized Laplacian is:

$$-\Delta_{h,D} = K \otimes I + I \otimes K$$

where $\otimes$ is the Kronecker product, defined as the following:

$$v \otimes w = \begin{bmatrix} v_1 w \\ v_2 w \\ v_3 w \\ \vdots \\ v_{n-1} w \end{bmatrix}$$

Lemma: $(A \otimes B)(v \otimes w) = (Av) \otimes (Bw)$. Remembering that the eigenvectors of the K are such that:

$$K v_k = \lambda_k v_k$$

we can write:

$$\begin{aligned} (K \otimes I + I \otimes K)(v_{k1} \otimes v_{k2}) &= (K \otimes I)(v_{k1} \otimes v_{k2}) + (I \otimes K)(v_{k1} \otimes v_{k2}) \\ &= \lambda_{k1} v_{k1} \otimes v_{k2} + \lambda_{k2} v_{k1} \otimes v_{k2} \\ &= (\lambda_{k1} + \lambda_{k2})(v_{k1} \otimes v_{k2}) \end{aligned}$$

Thus as shown, the eigenvalues of the 2D Laplacian is simply the sum of the eigenvalues of the 1D case. Since K is positive definite, the 2D $-\Delta_{h,D}$ must be stable. This also shows that the 2D Laplacian is diagonalized by the 2D DST S^{-1} and 2D IDST S :

$$\begin{aligned} -\Delta_{h,D} &= S \Lambda S^{-1} \\ -\Delta_{h,D}^{-1} &= S \Lambda^{-1} S^{-1} \end{aligned}$$

For a fast algorithm:

1. Expand f into $< f, v_{k1} \otimes v_{k2} >$. Do a 1D DST of each row, and then a 1D DST of each column. Each dimensional DST is $\mathcal{O}(n \times n \log n)$, meaning that the overall algorithm is $\mathcal{O}(n^2 \log n) = \mathcal{O}(N \log N)$.
2. Multiply by diag$(\frac{1}{\lambda_{k1} + \lambda_{k2}})$. One eigenvalue per point, this is also $\mathcal{O}(n^2) = \mathcal{O}(N)$
3. Inverse DST is just the DST.

4 Spectral methods

So far we have taken the derivatives by using the eigenvalues associated with the discretized Laplacian operator. However, we can do much better, since we can take derivatives of the eigenfunctions directly:

$$\begin{aligned} v_k(x) &= \sin(\pi k x) \\ -\frac{\partial^2}{\partial x^2} v_k(x) &= \pi^2 k^2 \sin(\pi k x) \end{aligned}$$

The continuous eigenvalues are therefore:

$$\lambda_{k,con} = \pi^2 k^2$$

Compare with the discretized operator:

$$\lambda_{k,disc} = \frac{1}{h^2}(2 - 2\cos(\frac{\pi k}{n}))$$

These functions do not look alike until they are plotted. Using Taylor expansions, we find that the agreement between the two functions is second order.

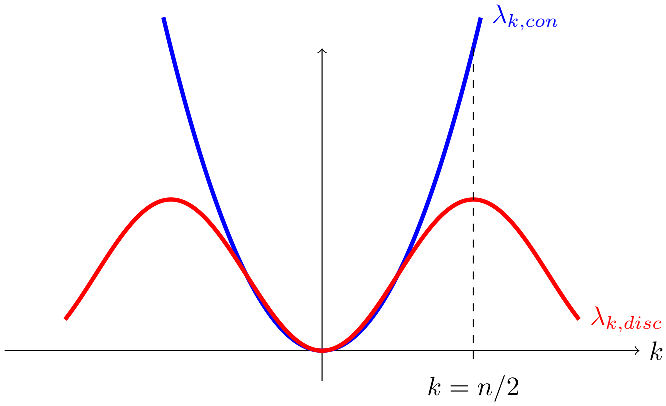

These results show that where a fast method expands a problem system to sines and cosines, then spectral accuracy can be obtained for free. We would get convergence at faster than polynomial rates, which is very fast indeed.

Aside: Use of the FFT for Convolutions

Given two discrete signals f_j and g_j where $j = 0 \ldots n-1$, the cyclic convolution is the following:

$$(f * g)_j = \sum_{j'=0}^{n-1} f_{j-j'} g_{j'}$$

when we write out the expansion:

$$f * g = \begin{bmatrix} f_0 & f_{n-1} & \cdots & f_1 \\ f_1 & f_0 & & \vdots \\ \vdots & & \ddots & \vdots \\ f_{n-1} & \cdots & \cdots & f_0 \end{bmatrix} g$$

The convolution matrix is circulant and can be diagonalized by the FFT.