

6.300: Signal Processing

Fast Fourier Transform

Agenda for Recitation

- History of the FFT
- Decimation-in-time FFT algorithm
- Finite-dimensional linear algebra and the FFT
- Polynomial multiplication with the FFT

History of the FFT

1805: Gauss develops a “method [that] greatly reduces the tediousness of mechanical calculations,” which will later come to be known as the **fast Fourier transform**.¹

1965: Cooley and **Tukey** publish “An Algorithm for the Machine Calculation of Complex Fourier Series.” This algorithm — the Cooley-Tukey FFT algorithm — made digital signal processing (DSP) practical.

Professor Alan V. Oppenheim: “You could think of it as the difference between BC and AD. The birth of the FFT was a very significant event.”^a

^a https://ethw.org/Oral-History:Alan_Oppenheim

¹ M.T. Heideman, D.H. Johnson, C.S. Burrus, and C. Truesdell. “Gauss and the History of the Fast Fourier Transform.” 1985.

History of the FFT

The FFT isn't really for pencil-and-paper calculations.

FFT Speedup

The small change in operation count for small N also explains why Gauss was not so excited about the method.


N	DFT	FFT	speed-up
2	4	2	2.0
4	16	8	2.0
8	64	24	2.7
16	256	64	4.0

So, we will try to **understand** conceptually why the FFT is so fast, rather than working through examples (with relatively small N) on the chalkboard.

Agenda for Recitation

- History of the FFT
- Decimation-in-time FFT algorithm
- Finite-dimensional linear algebra and the FFT
- Polynomial multiplication with the FFT

Decimation in Time



Graphic: Professor Denny Freeman (freeman@mit.edu)

Decimation in Time

The **signal flowgraph** on the previous slide is a graphical representation of the **radix-2 decimation-in-time FFT algorithm**. The big idea is to split an N -point DFT into a linear combination of $N/2$ -point DFTs.

$$X_N[k] = \frac{1}{2} \left(X_{N/2}^{\text{even}}[k] + r_N^k X_{N/2}^{\text{odd}}[k] \right)$$

Here, $r_N \triangleq e^{-j\frac{2\pi}{N}}$ denotes the **primitive N^{th} root of unity** — also called “**twiddle factor**,” which is more fun to say.

- Repeat until $N/2 = 1$, when we can't divide by 2 again.
- The DFT of a 1-point signal is the signal itself.
- “Glue” the 1-point DFTs together.

Directly computing the DFT requires $\mathcal{O}(N \log_2 N)$ operations. The FFT exploits the structure of the DFT to cut down the number of operations to $\mathcal{O}(N \log_2 N)$.

Agenda for Recitation

- History of the FFT
- Decimation-in-time FFT algorithm
- Finite-dimensional linear algebra and the FFT
- Polynomial multiplication with the FFT

Signal Processing and Linear Algebra

Signal processing is built on **linear algebra**.

For example, the analysis formulæ are inner products.

$$X(\omega) = \langle x(t) | e^{j\omega t} \rangle = \int_{-\infty}^{\infty} x(t) e^{-j\omega t} dt$$

So far, we have been doing linear algebra in an **infinite-dimensional space**, with $t \in \mathbb{R}$ or $n \in \mathbb{Z}$. If we restrict n to lie in $\{0, 1, 2, \dots, N-1\}$, then we'll be doing linear algebra in a **finite-dimensional space**.

Irving Kaplansky (1917–2006)

[Paul Halmos and I] share a philosophy about linear algebra: We think basis-free, we write basis-free, but when the chips are down we close the office door and compute with matrices like fury.

This is where **matrices** and **vectors** come into play.

Finite-Dimensional Linear Algebra

Letting $r = e^{-j\frac{2\pi}{N}}$ denote the **primitive N^{th} root of unity**, we can write the **DFT analysis equation** as follows.

$$X[k] = \frac{1}{N} \sum_{n=0}^{N-1} x[n] r^{kn}$$

Express the DFT analysis equation as a matrix-vector multiplication.

Fourier Matrix

The $N = 1$ Fourier matrix is hardly a matrix.

$$\mathbf{F}_1 = \begin{bmatrix} e^{-j\frac{2\pi}{1}0} \end{bmatrix} = [1]$$

The $N = 2$ Fourier matrix is simple.


$$\mathbf{F}_2 = \begin{bmatrix} e^{-j\frac{2\pi}{2}0} & e^{-j\frac{2\pi}{2}1} \\ e^{-j\frac{2\pi}{2}0} & e^{-j\frac{2\pi}{2}1} \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$$

The $N = 4$ Fourier matrix is relatively simple, too.

$$\mathbf{F}_4 = \begin{bmatrix} e^{-j\frac{2\pi}{4}0} & e^{-j\frac{2\pi}{4}1} & e^{-j\frac{2\pi}{4}2} & e^{-j\frac{2\pi}{4}3} \\ e^{-j\frac{2\pi}{4}0} & e^{-j\frac{2\pi}{4}1} & e^{-j\frac{2\pi}{4}2} & e^{-j\frac{2\pi}{4}3} \\ e^{-j\frac{2\pi}{4}0} & e^{-j\frac{2\pi}{4}2} & e^{-j\frac{2\pi}{4}4} & e^{-j\frac{2\pi}{4}6} \\ e^{-j\frac{2\pi}{4}0} & e^{-j\frac{2\pi}{4}3} & e^{-j\frac{2\pi}{4}6} & e^{-j\frac{2\pi}{4}9} \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -j & -1 & j \\ 1 & -1 & 1 & -1 \\ 1 & j & -1 & -j \end{bmatrix}$$


Fourier Matrix

Here's a visualization of the $N = 8$ Fourier matrix.




Fourier Matrix

Here's a visualization of the $N = 16$ Fourier matrix.




Fourier Matrix

Here's a visualization of the $N = 32$ Fourier matrix.



Fourier Matrix

Here's a visualization of the $N = 64$ Fourier matrix.



Counting Operations: Analysis

It takes $\mathcal{O}(N^2)$ operations to evaluate the sum

$$X[k] = \frac{1}{N} \sum_{n=0}^{N-1} x[n] e^{-jk \frac{2\pi}{N} n}$$

for $k \in \{0, 1, 2, \dots, N-1\}$.

- N different values of k
- $\mathcal{O}(N)$ operations per k

Likewise, the matrix-vector multiplication

$$\hat{\mathbf{x}} = \frac{1}{N} \mathbf{F} \mathbf{x}$$

takes $\mathcal{O}(N^2)$ operations.

Whichever way we look at it, direct computation of the DFT takes $\mathcal{O}(N^2)$ operations.

Counting Operations: Synthesis

Likewise, it takes $\mathcal{O}(N^2)$ operations to evaluate the sum

$$x[n] = \sum_{k=0}^{N-1} X[k] e^{jk \frac{2\pi}{N} n}$$

for $n \in \{0, 1, 2, \dots, N-1\}$.

However, in general, it takes $\mathcal{O}(N^3)$ operations to solve a matrix-vector equation like $\hat{\mathbf{x}} = \frac{1}{N} \mathbf{F} \mathbf{x}$ for $\mathbf{x}$. In general, inverting a matrix or solving a system of linear equations is $\mathcal{O}(N^3)$.

$$\mathbf{x} = \mathbf{F}^{-1} \hat{\mathbf{x}}$$

Something seems off: Shouldn't we be able to solve $\hat{\mathbf{x}} = \frac{1}{N} \mathbf{F} \mathbf{x}$ for $\mathbf{x}$ using no more than $\mathcal{O}(N^2)$ operations?

FFT as a Matrix Factorization

The big idea behind the **matrix factorization** approach is this: Multiplying a vector by a **sparse matrix** (i.e., mostly zeros) requires fewer than $\mathcal{O}(N^2)$ operations — perhaps on the order of **$\mathcal{O}(N)$ operations** instead.

FFT: Factor $\mathbf{F}_N$ into $\log_2 N$ sparse matrices, where multiplying by each sparse matrix requires $\mathcal{O}(N)$ operations. The total cost is **$\mathcal{O}(N \log_2 N)$** .

$$\mathbf{F}_N = \mathbf{F}_{\log_2 N} \cdots \mathbf{F}_2 \mathbf{F}_1 \mathbf{P}_N$$

- $\mathbf{F}_{\log_2 N}, \dots, \mathbf{F}_2, \mathbf{F}_1$ are sparse matrices.
- $\mathbf{P}_N$ is a permutation matrix.

FFT as a Matrix Factorization

Recursive block-matrix factorization:

$$\mathbf{F}_N = \begin{bmatrix} \mathbf{I}_{N/2} & \mathbf{D}_{N/2} \\ \mathbf{I}_{N/2} & -\mathbf{D}_{N/2} \end{bmatrix} \begin{bmatrix} \mathbf{F}_{N/2} & \\ & \mathbf{F}_{N/2} \end{bmatrix} \begin{bmatrix} \text{even-odd} \\ \text{permutation} \end{bmatrix}$$

Example with $N = 4$:

$$\mathbf{F}_4 = \begin{bmatrix} 1 & & 1 & \\ & 1 & & j \\ 1 & & -1 & \\ & 1 & & -j \end{bmatrix} \begin{bmatrix} 1 & 1 & & \\ 1 & -1 & & \\ & & 1 & 1 \\ & & 1 & -1 \end{bmatrix} \begin{bmatrix} 1 & & & \\ & & 1 & \\ & 1 & & \\ & & & 1 \end{bmatrix}$$

FFT algorithms let us compute (circular) convolutions in $\mathcal{O}(N \log_2 N)$ operations, rather than $\mathcal{O}(N^2)$ operations! Let's investigate how.

Fast Circular Convolution with the FFT

Let

$$\mathbf{x} = [x[0], x[1], x[2], x[3], \dots, x[N-1]]^\top$$

and

$$\mathbf{h} = [h[0], h[1], h[2], h[3], \dots, h[N-1]]^\top$$

hold the samples of $x[n]$ and $h[n]$, respectively. (Assume we zero-padded them to the same length if necessary.)

Circular convolution is a matrix-vector multiplication.

$$(\mathbf{x} \circledast \mathbf{h}) = \begin{bmatrix} h[0] & h[N-1] & h[N-2] & \cdots & h[1] \\ h[1] & h[0] & h[N-1] & \cdots & h[2] \\ h[2] & h[1] & h[0] & \cdots & h[3] \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ h[N-1] & h[N-2] & h[N-3] & \cdots & h[0] \end{bmatrix} \begin{bmatrix} x[0] \\ x[1] \\ x[2] \\ \vdots \\ x[N-1] \end{bmatrix}$$

Fast Circular Convolution with the FFT

Circular convolution is a matrix-vector multiplication.

$$(\mathbf{x} \circledast \mathbf{h}) = \underbrace{\begin{bmatrix} h[0] & h[N-1] & h[N-2] & \cdots & h[1] \\ h[1] & h[0] & h[N-1] & \cdots & h[2] \\ h[2] & h[1] & h[0] & \cdots & h[3] \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ h[N-1] & h[N-2] & h[N-3] & \cdots & h[0] \end{bmatrix}}_{\mathbf{H} = \text{circulant}(\mathbf{h})} \underbrace{\begin{bmatrix} x[0] \\ x[1] \\ x[2] \\ \vdots \\ x[N-1] \end{bmatrix}}_{\mathbf{x}}$$

The matrix $\mathbf{H} = \text{circulant}(\mathbf{h})$ has a very special structure: It is a **circulant matrix!** Each column (or row) is a circular shift of another column (or row) in the matrix.

Fast Circular Convolution with the FFT

LTI Systems

$$e^{jk\frac{2\pi}{N}n} \rightarrow \boxed{\text{LTI}} \rightarrow H[k]e^{jk\frac{2\pi}{N}n}$$

Recall: Exponentials are eigenfunctions of LTI systems.

- Eigenfunctions: $\phi_k[n] = e^{jk\frac{2\pi}{N}n}$
- Eigenvalues: $\lambda_k = H[k]$

Arrange samples of the eigenfunctions into vectors.

$$\phi_k = \begin{bmatrix} \phi_k[0] \\ \phi_k[1] \\ \phi_k[2] \\ \vdots \\ \phi_k[N-1] \end{bmatrix} = \begin{bmatrix} 1 \\ e^{jk\frac{2\pi}{N}} \\ e^{jk\frac{2\pi}{N}2} \\ \vdots \\ e^{jk\frac{2\pi}{N}(N-1)} \end{bmatrix} \quad \text{for } k \in \{0, 1, 2, \dots, N-1\}$$

Fast Circular Convolution with the FFT

Eigenpairs (ϕ_k, λ_k) satisfy the eigenequation.

$$\mathbf{H}\phi_k = \lambda_k\phi_k$$

Write out the eigenequation for $k \in \{0, \dots, N-1\}$.

$$\mathbf{H} \underbrace{\begin{bmatrix} | & | & & | \\ \phi_0 & \phi_1 & \cdots & \phi_{N-1} \\ | & | & & | \end{bmatrix}}_{\text{eigenvector matrix } \Phi} = \underbrace{\begin{bmatrix} | & | & & | \\ \phi_0 & \phi_1 & \cdots & \phi_{N-1} \\ | & | & & | \end{bmatrix}}_{\text{eigenvector matrix } \Phi} \underbrace{\begin{bmatrix} \lambda_0 & & & \\ & \ddots & & \\ & & \lambda_{N-1} \end{bmatrix}}_{\text{eigenvalue matrix } \Lambda}$$

This yields the **eigenvalue decomposition** of $\mathbf{H}$.

$$\mathbf{H}\Phi = \Phi\Lambda \iff \mathbf{H} = \Phi\Lambda\Phi^{-1} = \Phi\Lambda\Phi^*$$

$\Phi = \mathbf{F}^{-1}$ is the (inverse) Fourier matrix. $\Phi^{-1} = \Phi^* = \mathbf{F}$.

Fast Circular Convolution with the FFT

All circulant matrices have a full set of orthogonal eigenvectors — the Fourier basis vectors!

$$\phi_k = \begin{bmatrix} 1 \\ e^{jk\frac{2\pi}{N}} \\ e^{jk\frac{2\pi}{N}2} \\ \vdots \\ e^{jk\frac{2\pi}{N}(N-1)} \end{bmatrix} \quad \text{for } k \in \{0, 1, 2, \dots, N-1\}$$

The eigenvalues come from the analysis equation.

$$\lambda_k = H[k] = \frac{1}{N} \left(h[0] + h[1]e^{-jk\frac{2\pi}{N}} + \dots + h[N-1]e^{-jk\frac{2\pi}{N}(N-1)} \right)$$

Fast Circular Convolution with the FFT

Recall that circular convolution is a matrix-vector multiplication: $\mathbf{y} = (\mathbf{x} \circledast \mathbf{h}) = \mathbf{H}\mathbf{x}$, where $\mathbf{H} = \text{circulant}(\mathbf{h})$ is a circulant matrix.

Substitute the eigenvalue decomposition for $\mathbf{H}$.

$$\mathbf{y} = (\mathbf{x} \circledast \mathbf{h}) = \mathbf{H}\mathbf{x} = (\mathbf{\Phi}\mathbf{\Lambda}\mathbf{\Phi}^*)\mathbf{x}$$

The DFT $\hat{\mathbf{x}} = \mathbf{\Phi}^*\mathbf{x}$ takes $\mathcal{O}(N \log_2 N)$ operations.

Multiplying $\hat{\mathbf{x}}$ by $\mathbf{\Lambda}$ (diagonal matrix of eigenvalues, i.e., the $H[k]$ terms) takes only $\mathcal{O}(N)$ operations: $\hat{\mathbf{y}} = \mathbf{\Lambda}\hat{\mathbf{x}}$.

Finally, compute the IDFT: $\mathbf{y} = \mathbf{\Phi}\hat{\mathbf{y}}$. This is $\mathcal{O}(N \log_2 N)$.

Cost: $\mathcal{O}(N \log_2 N) + \mathcal{O}(N) + \mathcal{O}(N \log_2 N) \implies \mathcal{O}(N \log_2 N)$

Fast Circular Convolution with the FFT

Summary

Circulant matrices have a nice eigenvalue decomposition.

- Transform: $\hat{x} = \Phi^* x$. $\mathcal{O}(N \log_2 N)$
- Multiply by eigenvalues: $\hat{y} = \Lambda \hat{x}$. $\mathcal{O}(N)$
- Inverse transform: $y = (x \circledast h) = \Phi \hat{y}$. $\mathcal{O}(N \log_2 N)$

$$y = (x \circledast h) = \underbrace{(\Phi \Lambda \Phi^*)}_H x$$

FFT algorithms let us compute $y = (x \circledast h)$ in $\mathcal{O}(N \log_2 N)$ operations, rather than $\mathcal{O}(N^2)$ operations!

If $x[n]$ is length L and $h[n]$ is length M , we can **zero-pad** both to length $N = L + M - 1$ so that $(x \circledast h)[n] = (x * h)[n]$ for $n \in \{0, 1, 2, \dots, N - 1\}$.

Agenda for Recitation

- History of the FFT
- Decimation-in-time FFT algorithm
- Finite-dimensional linear algebra and the FFT
- Polynomial multiplication with the FFT

Polynomial Multiplication with the FFT

Let

$$p(x) = P[0] + P[1]x + P[2]x^2 + P[3]x^3 + P[4]x^4 + \dots$$

and

$$q(x) = Q[0] + Q[1]x + Q[2]x^2 + Q[3]x^3 + Q[4]x^4 + \dots$$

denote polynomials. The coefficients for $p(x)q(x)$ are given by the **convolution** of the coefficients $P[k]$ and $Q[k]$.

$$(P * Q)[k] = \sum_{m=-\infty}^{\infty} P[m]Q[k-m] = \sum_{m=-\infty}^{\infty} Q[m]P[k-m]$$

Notice that multiplication in one domain corresponds to convolution in the other domain: $p(x)q(x) \iff (P * Q)[k]$.

What if $x = e^{j\frac{2\pi}{N}n}$? Looks familiar ...

$$p'[n] \triangleq p(e^{j\frac{2\pi}{N}n}) = P[0] + P[1]e^{j\frac{2\pi}{N}n} + P[2]e^{j2\frac{2\pi}{N}n} + \dots$$

Lessons Learned

The **fast Fourier transform (FFT)** refers to a family of algorithms for efficiently computing the **DFT**.

FFT matrix factorization: Factor the $N \times N$ Fourier matrix $\mathbf{F}_N$ into $\log_2 N$ sparse matrices, where multiplying by each sparse matrix requires $\mathcal{O}(N)$ operations. The total cost is $\mathcal{O}(N \log_2 N)$.

$$\mathbf{F}_N = \mathbf{F}_{\log_2 N} \cdots \mathbf{F}_2 \mathbf{F}_1 \mathbf{P}_N$$

- $\mathbf{F}_{\log_2 N}, \dots, \mathbf{F}_2, \mathbf{F}_1$ are sparse matrices.
- $\mathbf{P}_N$ is a permutation matrix.

Circulant matrices have a nice eigenvalue decomposition. FFT algorithms let us compute $\mathbf{y} = (\mathbf{x} \circledast \mathbf{h})$ in $\mathcal{O}(N \log_2 N)$ operations, rather than $\mathcal{O}(N^2)$ operations!