

THE FAST FOURIER TRANSFORM ■ 10.3

Many applications of linear algebra take time to develop. It is not easy to explain them in an hour. The teacher and the author must choose between completing the theory and adding new applications. Generally the theory wins, because this course is the best chance to make it clear—and the importance of any one application seems limited. This section is almost an exception, because the importance of Fourier transforms is almost unlimited.

More than that, the algebra is basic. *We want to multiply quickly by F and F^{-1} , the Fourier matrix and its inverse.* This is achieved by the Fast Fourier Transform—the most valuable numerical algorithm in our lifetime.

The FFT has revolutionized signal processing. Whole industries are speeded up by this one idea. Electrical engineers are the first to know the difference—they take your Fourier transform as they meet you (if you are a function). Fourier's idea is to represent f as a sum of harmonics $c_k e^{ikx}$. The function is seen in *frequency space* through the coefficients c_k , instead of *physical space* through its values $f(x)$. The passage backward and forward between c 's and f 's is by the Fourier transform. Fast passage is by the FFT.

An ordinary product Fc uses n^2 multiplications (the matrix has n^2 nonzero entries). The Fast Fourier Transform needs only n times $\frac{1}{2} \log_2 n$. We will see how.

Roots of Unity and the Fourier Matrix

Quadratic equations have two roots (or one repeated root). Equations of degree n have n roots (counting repetitions). This is the Fundamental Theorem of Algebra, and to make it true we must allow complex roots. This section is about the very special equation $z^n = 1$. The solutions z are the “ n th roots of unity.” They are n evenly spaced points around the unit circle in the complex plane.

Figure 10.5 shows the eight solutions to $z^8 = 1$. Their spacing is $\frac{1}{8}(360^\circ) = 45^\circ$. The first root is at 45° or $\theta = 2\pi/8$ radians. *It is the complex number $w = e^{i\theta} = e^{i2\pi/8}$.* We call this number w_8 to emphasize that it is an 8th root. You could write it in terms of $\cos \frac{2\pi}{8}$ and $\sin \frac{2\pi}{8}$, but don't do it. The seven other 8th roots are w^2, w^3, \dots, w^8 , going around the circle. Powers of w are best in polar form, because we work only with the angle.

The fourth roots of 1 are also in the figure. They are $i, -1, -i, 1$. The angle is now $2\pi/4$ or 90° . The first root $w_4 = e^{2\pi i/4}$ is nothing but i . Even the square roots of 1 are seen, with $w_2 = e^{i2\pi/2} = -1$. Do not despise those square roots 1 and -1 . The idea behind the FFT is to go from an 8 by 8 Fourier matrix (containing powers of w_8) to the 4 by 4 matrix below (with powers of $w_4 = i$). The same idea goes from 4 to 2. By exploiting the connections of F_8 down to F_4 and up to F_{16} (and beyond), the FFT makes multiplication by F_{1024} very quick.

We describe the *Fourier matrix*, first for $n = 4$. Its rows contain powers of 1 and w and w^2 and w^3 . These are the fourth roots of 1, and their powers come in a




Figure 10.5 The eight solutions to $z^8 = 1$ are $1, w, w^2, \dots, w^7$ with $w = (1+i)/\sqrt{2}$.

special order:

$$F_4 = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & w & w^2 & w^3 \\ 1 & w^2 & w^4 & w^6 \\ 1 & w^3 & w^6 & w^9 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & i & i^2 & i^3 \\ 1 & i^2 & i^4 & i^6 \\ 1 & i^3 & i^6 & i^9 \end{bmatrix}.$$

The matrix is symmetric ($F = F^T$). It is *not* Hermitian. Its main diagonal is not real. But $\frac{1}{2}F$ is a **unitary matrix**, which means that $(\frac{1}{2}F^H)(\frac{1}{2}F) = I$:

The columns of F give $F^H F = 4I$. The inverse of F is $\frac{1}{4}F^H$ which is $\frac{1}{4}\overline{F}$.

The inverse changes from $w = i$ to $\bar{w} = -i$. That takes us from F to \overline{F} . When the Fast Fourier Transform gives a quick way to multiply by F_4 , it does the same for the inverse.

The unitary matrix is $U = F/\sqrt{n}$. We prefer to avoid that \sqrt{n} and just put $\frac{1}{n}$ outside F^{-1} . The main point is to multiply the matrix F times the coefficients in the Fourier series $c_0 + c_1 e^{ix} + c_2 e^{2ix} + c_3 e^{3ix}$:

$$Fc = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & w & w^2 & w^3 \\ 1 & w^2 & w^4 & w^6 \\ 1 & w^3 & w^6 & w^9 \end{bmatrix} \begin{bmatrix} c_0 \\ c_1 \\ c_2 \\ c_3 \end{bmatrix} = \begin{bmatrix} c_0 + c_1 + c_2 + c_3 \\ c_0 + c_1 w + c_2 w^2 + c_3 w^3 \\ c_0 + c_1 w^2 + c_2 w^4 + c_3 w^6 \\ c_0 + c_1 w^3 + c_2 w^6 + c_3 w^9 \end{bmatrix}. \quad (1)$$

The input is four complex coefficients c_0, c_1, c_2, c_3 . The output is four function values y_0, y_1, y_2, y_3 . The first output $y_0 = c_0 + c_1 + c_2 + c_3$ is the value of the Fourier series at $x = 0$. The second output is the value of that series $\sum c_k e^{ikx}$ at $x = 2\pi/4$:

$$y_1 = c_0 + c_1 e^{i2\pi/4} + c_2 e^{i4\pi/4} + c_3 e^{i6\pi/4} = c_0 + c_1 w + c_2 w^2 + c_3 w^3.$$

The third and fourth outputs y_2 and y_3 are the values of $\sum c_k e^{ikx}$ at $x = 4\pi/4$ and $x = 6\pi/4$. These are *finite* Fourier series! They contain $n = 4$ terms and they are evaluated at $n = 4$ points. Those points $x = 0, 2\pi/4, 4\pi/4, 6\pi/4$ are equally spaced.

The next point would be $x = 8\pi/4$ which is 2π . Then the series is back to y_0 , because $e^{2\pi i}$ is the same as $e^0 = 1$. Everything cycles around with period 4. In this world $2 + 2$ is 0 because $(w^2)(w^2) = w^0 = 1$. In matrix shorthand, F times \mathbf{c} gives a column vector \mathbf{y} . The four \mathbf{y} 's come from evaluating the series at the four x 's with spacing $2\pi/4$:

$$\mathbf{y} = F\mathbf{c} \text{ produces } y_j = \sum_{k=0}^3 c_k e^{ik(2\pi j/4)} = \text{the value of the series at } x = \frac{2\pi j}{4}.$$

We will follow the convention that j and k go from 0 to $n - 1$ (instead of 1 to n). The “zeroth row” and “zeroth column” of F contain all ones.

The n by n Fourier matrix contains powers of $w = e^{2\pi i/n}$:

$$F_n \mathbf{c} = \begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & w & w^2 & \cdots & w^{n-1} \\ 1 & w^2 & w^4 & \cdots & w^{2(n-1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & w^{n-1} & w^{2(n-1)} & \cdots & w^{(n-1)^2} \end{bmatrix} \begin{bmatrix} c_0 \\ c_1 \\ c_2 \\ \vdots \\ c_{n-1} \end{bmatrix} = \begin{bmatrix} y_0 \\ y_1 \\ y_2 \\ \vdots \\ y_{n-1} \end{bmatrix}. \quad (2)$$

F_n is symmetric but not Hermitian. *Its columns are orthogonal*, and $F_n \overline{F_n} = nI$. Then F_n^{-1} is $\overline{F_n}/n$. The inverse contains powers of $\overline{w}_n = e^{-2\pi i/n}$. Look at the pattern in F :

The entry in row j , column k is w^{jk} . Row zero and column zero contain $w^0 = 1$.

The zeroth output is $y_0 = c_0 + c_1 + \cdots + c_{n-1}$. This is the series $\sum c_k e^{ikx}$ at $x = 0$. When we multiply \mathbf{c} by F_n , we sum the series at n points. When we multiply \mathbf{y} by F_n^{-1} , we find the coefficients \mathbf{c} from the function values \mathbf{y} . The matrix F passes from “frequency space” to “physical space.” F^{-1} returns from the function values \mathbf{y} to the Fourier coefficients \mathbf{c} .

One Step of the Fast Fourier Transform

We want to multiply F times \mathbf{c} as quickly as possible. Normally a matrix times a vector takes n^2 separate multiplications—the matrix has n^2 entries. You might think it is impossible to do better. (If the matrix has zero entries then multiplications can be skipped. But the Fourier matrix has no zeros!) By using the special pattern w^{jk} for its entries, F can be factored in a way that produces many zeros. This is the **FFT**.

The key idea is to connect F_n with the half-size Fourier matrix $F_{n/2}$. Assume that n is a power of 2 (say $n = 2^{10} = 1024$). We will connect F_{1024} to F_{512} —or rather

to *two copies of* F_{512} . When $n = 4$, the key is in the relation between the matrices

$$F_4 = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & i & i^2 & i^3 \\ 1 & i^2 & i^4 & i^6 \\ 1 & i^3 & i^6 & i^9 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} F_2 & \\ & F_2 \end{bmatrix} = \begin{bmatrix} 1 & 1 & & \\ & i^2 & & \\ & & 1 & 1 \\ & & & i^2 \end{bmatrix}.$$

On the left is F_4 , with no zeros. On the right is a matrix that is half zero. The work is cut in half. But wait, those matrices are not the same. The block matrix with two copies of the half-size F is one piece of the picture but not the only piece. Here is the factorization of F_4 with many zeros:

$$F_4 = \begin{bmatrix} 1 & & 1 & \\ & 1 & & i \\ 1 & & -1 & \\ & 1 & & -i \end{bmatrix} \begin{bmatrix} 1 & 1 & & \\ 1 & i^2 & & \\ & & 1 & 1 \\ & & 1 & i^2 \end{bmatrix} \begin{bmatrix} 1 & & & \\ & 1 & & \\ & & 1 & \\ & & & 1 \end{bmatrix}. \quad (3)$$

The matrix on the right is a permutation. It puts the even c 's (c_0 and c_2) ahead of the odd c 's (c_1 and c_3). The middle matrix performs separate half-size transforms on the evens and odds. The matrix at the left combines the two half-size outputs—in a way that produces the correct full-size output $y = F_4 c$. You could multiply those three matrices to see that their product is F_4 .

The same idea applies when $n = 1024$ and $m = \frac{1}{2}n = 512$. The number w is $e^{2\pi i/1024}$. It is at the angle $\theta = 2\pi/1024$ on the unit circle. The Fourier matrix F_{1024} is full of powers of w . The first stage of the FFT is the great factorization discovered by Cooley and Tukey (and foreshadowed in 1805 by Gauss):

$$F_{1024} = \begin{bmatrix} I_{512} & D_{512} \\ I_{512} & -D_{512} \end{bmatrix} \begin{bmatrix} F_{512} & \\ & F_{512} \end{bmatrix} \begin{bmatrix} \text{even-odd} \\ \text{permutation} \end{bmatrix}. \quad (4)$$

I_{512} is the identity matrix. D_{512} is the diagonal matrix with entries $(1, w, \dots, w^{511})$. The two copies of F_{512} are what we expected. Don't forget that they use the 512th root of unity (which is nothing but w^2 !!) The permutation matrix separates the incoming vector c into its even and odd parts $c' = (c_0, c_2, \dots, c_{1022})$ and $c'' = (c_1, c_3, \dots, c_{1023})$.

Here are the algebra formulas which say the same thing as the factorization of F_{1024} :

101 (FFT) Set $m = \frac{1}{2}n$. The first m and last m components of $y = F_n c$ are combinations of the half-size transforms $y' = F_m c'$ and $y'' = F_m c''$. Equation (4) shows $Iy' + Dy''$ and $Iy' - Dy''$:

$$\begin{aligned} y_j &= y'_j + w_n^j y''_j, & j &= 0, \dots, m-1 \\ y_{j+m} &= y'_j - w_n^j y''_j, & j &= 0, \dots, m-1. \end{aligned} \quad (5)$$

Thus the three steps are: split c into c' and c'' , transform them by F_m into y' and y'' , and reconstruct y from equation (5).

You might like the flow graph in Figure 10.6 better than these formulas. The graph for $n = 4$ shows c' and c'' going through the half-size F_2 . Those steps are called “butterflies,” from their shape. Then the outputs from the F_2 's are combined using the I and D matrices to produce $y = F_4c$:




Figure 10.6 Flow graph for the Fast Fourier Transform with $n = 4$.

This reduction from F_n to two F_m 's almost cuts the work in half—you see the zeros in the matrix factorization. That reduction is good but not great. The full idea of the **FFT** is much more powerful. It saves much more than half the time.

The Full FFT by Recursion

If you have read this far, you have probably guessed what comes next. We reduced F_n to $F_{n/2}$. **Keep going to $F_{n/4}$.** The matrices F_{512} lead to F_{256} (in four copies). Then 256 leads to 128. *That is recursion.* It is a basic principle of many fast algorithms, and here is the second stage with four copies of $F = F_{256}$ and $D = D_{256}$:

$$\begin{bmatrix} F_{512} & \\ & F_{512} \end{bmatrix} = \begin{bmatrix} I & D \\ I & -D \\ & I & D \\ & I & -D \end{bmatrix} \begin{bmatrix} F & & \\ & F & \\ & & F & \\ & & & F \end{bmatrix} \begin{bmatrix} \text{pick } 0, 4, 8, \dots \\ \text{pick } 2, 6, 10, \dots \\ \text{pick } 1, 5, 9, \dots \\ \text{pick } 3, 7, 11, \dots \end{bmatrix}.$$

We will count the individual multiplications, to see how much is saved. Before the **FFT** was invented, the count was the usual $n^2 = (1024)^2$. This is about a million multiplications. I am not saying that they take a long time. The cost becomes large when we have many, many transforms to do—which is typical. Then the saving by the **FFT** is also large:

The final count for size $n = 2^l$ is reduced from n^2 to $\frac{1}{2}nl$.

The number 1024 is 2^{10} , so $l = 10$. The original count of $(1024)^2$ is reduced to $(5)(1024)$. The saving is a factor of 200. A million is reduced to five thousand. That is why the **FFT** has revolutionized signal processing.

Here is the reasoning behind $\frac{1}{2}nl$. There are l levels, going from $n = 2^l$ down to $n = 1$. Each level has $\frac{1}{2}n$ multiplications from the diagonal D 's, to reassemble the half-size outputs from the lower level. This yields the final count $\frac{1}{2}nl$, which is $\frac{1}{2}n \log_2 n$.

One last note about this remarkable algorithm. There is an amazing rule for the order that the c 's enter the **FFT**, after all the even-odd permutations. Write the numbers 0 to $n - 1$ in binary (base 2). *Reverse the order of their digits*. The complete picture shows the bit-reversed order at the start, the $l = \log_2 n$ steps of the recursion, and the final output y_0, \dots, y_{n-1} which is F_n times c . The book ends with that very fundamental idea, a matrix multiplying a vector.

Thank you for studying linear algebra. I hope you enjoyed it, and I very much hope you will use it. That is why the book was written. It was a pleasure.

Problem Set 10.3

- 1 Multiply the three matrices in equation (3) and compare with F . In which six entries do you need to know that $i^2 = -1$?
- 2 Invert the three factors in equation (3) to find a fast factorization of F^{-1} .
- 3 F is symmetric. So transpose equation (3) to find a new Fast Fourier Transform!
- 4 All entries in the factorization of F_6 involve powers of $w =$ sixth root of 1:

$$F_6 = \begin{bmatrix} I & D \\ I & -D \end{bmatrix} \begin{bmatrix} F_3 & \\ & F_3 \end{bmatrix} \begin{bmatrix} P \\ & P \end{bmatrix}.$$

Write down these three factors with 1, w , w^2 in D and powers of w^2 in F_3 . Multiply!

- 5 If $v = (1, 0, 0, 0)$ and $w = (1, 1, 1, 1)$, show that $Fv = w$ and $Fw = 4v$. Therefore $F^{-1}w = v$ and $F^{-1}v = \underline{\hspace{2cm}}$.
- 6 What is F^2 and what is F^4 for the 4 by 4 Fourier matrix?
- 7 Put the vector $c = (1, 0, 1, 0)$ through the three steps of the **FFT** to find $y = Fc$. Do the same for $c = (0, 1, 0, 1)$.
- 8 Compute $y = F_8c$ by the three **FFT** steps for $c = (1, 0, 1, 0, 1, 0, 1, 0)$. Repeat the computation for $c = (0, 1, 0, 1, 0, 1, 0, 1)$.
- 9 If $w = e^{2\pi i/64}$ then w^2 and \sqrt{w} are among the $\underline{\hspace{2cm}}$ and $\underline{\hspace{2cm}}$ roots of 1.
- 10 (a) Draw all the sixth roots of 1 on the unit circle. Prove they add to zero.
(b) What are the three cube roots of 1? Do they also add to zero?

- 11 The columns of the Fourier matrix F are the *eigenvectors* of the cyclic permutation P . Multiply PF to find the eigenvalues λ_1 to λ_4 :

$$\begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & i & i^2 & i^3 \\ 1 & i^2 & i^4 & i^6 \\ 1 & i^3 & i^6 & i^9 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & i & i^2 & i^3 \\ 1 & i^2 & i^4 & i^6 \\ 1 & i^3 & i^6 & i^9 \end{bmatrix} \begin{bmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \lambda_3 & \\ & & & \lambda_4 \end{bmatrix}.$$

This is $PF = F\Lambda$ or $P = F\Lambda F^{-1}$. The eigenvector matrix (usually S) is F .

- 12 The equation $\det(P - \lambda I) = 0$ is $\lambda^4 = 1$. This shows again that the eigenvalue matrix Λ is _____. Which permutation P has eigenvalues = cube roots of 1?
- 13 (a) Two eigenvectors of C are $(1, 1, 1, 1)$ and $(1, i, i^2, i^3)$. What are the eigenvalues?

$$\begin{bmatrix} c_0 & c_1 & c_2 & c_3 \\ c_3 & c_0 & c_1 & c_2 \\ c_2 & c_3 & c_0 & c_1 \\ c_1 & c_2 & c_3 & c_0 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} = e_1 \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} \quad \text{and} \quad C \begin{bmatrix} 1 \\ i \\ i^2 \\ i^3 \end{bmatrix} = e_2 \begin{bmatrix} 1 \\ i \\ i^2 \\ i^3 \end{bmatrix}.$$

- (b) $P = F\Lambda F^{-1}$ immediately gives $P^2 = F\Lambda^2 F^{-1}$ and $P^3 = F\Lambda^3 F^{-1}$. Then $C = c_0 I + c_1 P + c_2 P^2 + c_3 P^3 = F(c_0 I + c_1 \Lambda + c_2 \Lambda^2 + c_3 \Lambda^3) F^{-1} = FEF^{-1}$. That matrix E in parentheses is diagonal. It contains the _____ of C .
- 14 Find the eigenvalues of the “periodic” $-1, 2, -1$ matrix from $E = 2I - \Lambda - \Lambda^3$, with the eigenvalues of P in Λ . The -1 ’s in the corners make this matrix periodic:

$$C = \begin{bmatrix} 2 & -1 & 0 & -1 \\ -1 & 2 & -1 & 0 \\ 0 & -1 & 2 & -1 \\ -1 & 0 & -1 & 2 \end{bmatrix} \quad \text{has } c_0 = 2, c_1 = -1, c_2 = 0, c_3 = -1.$$

- 15 To multiply C times a vector x , we can multiply $F(E(F^{-1}x))$ instead. The direct way uses n^2 separate multiplications. Knowing E and F , the second way uses only $n \log_2 n + n$ multiplications. How many of those come from E , how many from F , and how many from F^{-1} ?
- 16 How could you quickly compute these four components of Fc starting from $c_0 + c_2, c_0 - c_2, c_1 + c_3, c_1 - c_3$? You are finding the Fast Fourier Transform!

$$Fc = \begin{bmatrix} c_0 + c_1 + c_2 + c_3 \\ c_0 + ic_1 + i^2 c_2 + i^3 c_3 \\ c_0 + i^2 c_1 + i^4 c_2 + i^6 c_3 \\ c_0 + i^3 c_1 + i^6 c_2 + i^9 c_3 \end{bmatrix}.$$