

Fast Fourier Transform Algorithms (MIT IAP 2008)

Prof. Steven G. Johnson, MIT Dept. of Mathematics

January 11, 2008

Fast Fourier transforms (FFTs), $O(N \log N)$ algorithms to compute a discrete Fourier transform (DFT) of size N , have been called one of the ten most important algorithms of the 20th century. They are what make Fourier transforms practical on a computer, and Fourier transforms (which express any function as a sum of pure sinusoids) are used in everything from solving partial differential equations to digital signal processing (e.g. MP3 compression) to multiplying large numbers (for computing π to 10^{12} decimal places). Although the applications are important and numerous, the FFT algorithms themselves reveal a surprisingly rich variety of mathematics that has been the subject of active research for 40+ years, and into which this lecture will attempt to dip your toes. The DFT and its inverse are defined by the following relation between N inputs x_n and N outputs y_k (all complex numbers):

$$\text{DFT}(x_n): y_k = \sum_{n=0}^{N-1} x_n e^{-\frac{2\pi i}{N}nk}, \quad (1)$$

$$\text{inverse DFT}(y_k): x_n = \frac{1}{N} \sum_{k=0}^{N-1} y_k e^{+\frac{2\pi i}{N}nk} \quad (2)$$

where $i = \sqrt{-1}$, recalling Euler's identity that $e^{i\phi} = \cos \phi + i \sin \phi$. Each of the N DFT outputs $k = 0, \dots, N-1$ is the sum of N terms, so evaluating this formula directly requires $O(N^2)$ operations.¹ The trick is to rearrange this computation to expose redundant calculations that we can factor out.

The most important FFT algorithm is called the Cooley-Tukey (C-T) algorithm, after the two authors who popularized it in 1965 (unknowingly re-inventing an algorithm known to Gauss in 1805). It works for any *composite* size $N = N_1 N_2$ by re-expressing the DFT of size N in terms of smaller DFTs of size N_1 and N_2 (which are themselves broken down, recursively, into smaller DFTs until the prime factors are reached). Effectively, C-T expresses the array x_n of length N as a "two-dimensional" array of size $N_1 \times N_2$ indexed by (n_1, n_2) , so that $n = N_1 n_2 + n_1$ (where $n_{1,2} = 0, \dots, N_{1,2} - 1$). Similarly, the output is expressed as a *transposed* 2d array, $N_2 \times N_1$, indexed by

¹Read " $O(N^2)$ " as "roughly proportional, for large N ." e.g. $15N^2 + 24N$ is $O(N^2)$. (Technically, I should really say $\Theta(N^2)$, but I'm not going to get that formal.)

(k_2, k_1) , so that $k = N_2 k_1 + k_2$. Substituted into the DFT above, this gives:

$$y_{N_2 k_1 + k_2} = \sum_{n_1=0}^{N_1-1} \left(\left\{ e^{-\frac{2\pi i}{N} n_1 k_2} \right\} \left[\sum_{n_2=0}^{N_2-1} e^{-\frac{2\pi i}{N_2} n_2 k_2} x_{N_1 n_2 + n_1} \right] \right) e^{-\frac{2\pi i}{N_1} n_1 k_1} \quad (3)$$

where we have used the fact that $e^{-2\pi i n_2 k_1} = 1$ (for any integers n_2 and k_1). Here, the outer sum is exactly a length- N_1 DFT of the (\dots) terms, one for each value of k_2 ; and the inner sum in $[\dots]$ is a length- N_2 DFT, one for each value of n_1 . The phase in the $\{\dots\}$ is called the "twiddle factor" (honest). Assuming that N has small (bounded) prime factors, this algorithm requires $O(N \log N)$ operations when carried out recursively — the key savings coming from the fact that we have exposed a repeated calculation: the $[\dots]$ DFTs need only be carried out *once* for *all* y_k outputs.

For a given N , there are many choices of factorizations (e.g. $12 = 3 \cdot 4$ and $4 \cdot 3$ give a different sequence of computations). Moreover, the transposition from input to output implies a data rearrangement process that can be accomplished in many ways. As a result, many different strategies for evaluating the C-T algorithm have been proposed (each with its own name), and the optimal approach is still a matter of active research. Commonly, either N_1 or N_2 is a small (bounded) constant factor, called the *radix*, and the approach is called decimation in time (DIT) for $N_1 = \text{radix}$ or frequency (DIF) for $N_2 = \text{radix}$. Textbook examples are typically radix-2 DIT (dividing x_n into two interleaved halves with each step), but serious implementations employ more sophisticated strategies.

There are many other FFT algorithms and there are also many different ways to view the *same* algorithms. One fruitful way is to view the DFT in terms of operations on *polynomials*. In particular, define a polynomial $x(z)$ by

$$x(z) = \sum_{n=0}^{N-1} x_n z^n. \quad (4)$$

Then

$$y_k = x(e^{-\frac{2\pi i}{N}k}) = x(z) \text{ mod } (z - e^{-\frac{2\pi i}{N}k}), \quad (5)$$

where $x(z) \bmod u(z)$ ($x(z)$ “modulo” $u(z)$) means the *remainder* of dividing $x(z)$ by $u(z)$. Since $u(z) \bmod u(z) = 0$, taking $x(z) \bmod u(z)$ is equivalent to setting $u(z) = 0$, which in this case means setting $z = e^{-\frac{2\pi i}{N}k}$.

The DFT corresponds to computing $x(z) \bmod (z - e^{-\frac{2\pi i}{N}k})$ for all $k = 0 \dots N - 1$, which would take $O(N^2)$ operations if done directly. The key observation, from a polynomial viewpoint, is that we can do this modulo operation *recursively* by *combining* the factors $(z - e^{-\frac{2\pi i}{N}k})$. In particular, it is easy to show that $x(z) \bmod u(z) = [x(z) \bmod u(z)v(z)] \bmod u(z)$ for any $u(z)$ and $v(z)$. This means that we can first compute $x(z)$ modulo the *product* of the factors, and then recursively evaluate the remainder by a *recursive factorization* of this product. But the product $\prod_k (z - e^{-\frac{2\pi i}{N}k}) = z^N - 1$, since the $e^{-\frac{2\pi i}{N}k}$ are just the N th roots of unity (solutions of $z^N - 1 = 0$). It follows that any recursive factorization of $z^N - 1$ into $N \log N$ bounded-degree factors gives us an $O(N \log N)$ FFT algorithm! In particular, the radix-2 Cooley-Tukey algorithm is *equivalent* to the recursive factorization (for N a power of 2): $z^N - a = (z^{N/2} - \sqrt{a})(z^{N/2} + \sqrt{a})$, where we start with $a = 1$ and end up with $a = e^{-i\frac{2\pi i}{N}k}$.

Different recursive factorizations of $z^N - 1$ lead to different FFT algorithms, one of which you will examine for homework. Many other FFT algorithms exist as well, from the “prime-factor algorithm” (1958) that exploits the Chinese remainder theorem for $\gcd(N_1, N_2) = 1$, to FFT algorithms that work for *prime* N , one of which we give below.

The core of the DFT is the constant $\omega_N = e^{-\frac{2\pi i}{N}}$; because this is a primitive root of unity ($\omega_N^N = 1$), any exponent of ω_N is evaluated *modulo* N . That is, $\omega_N^m = \omega_N^r$ where r is the remainder when we divide m by N ($r = m \bmod N$). A great body of number theory has been developed around such “modular arithmetic”, and we can exploit it to develop FFT algorithms different from C-T. For example, Rader’s algorithm (1968) allows us to compute $O(N \log N)$ FFTs of *prime* sizes N , by turning the DFT into a cyclic *convolution* of length $N - 1$, which in turn is evaluated by (non-prime) FFTs. Given a_n and b_n ($n = 0, \dots, N - 1$), their convolution c_n is defined by the sum

$$c_n = \sum_{m=0}^{N-1} a_m b_{n-m}, \quad (6)$$

where the convolution is *cyclic* if the $n - m$ subscript is “wrapped” periodically onto $0, \dots, N - 1$. This operation is central to digital filtering, differential equations, and other applications, and is evaluated in $O(N \log N)$ time by the *convolution theorem*: $c_n = \text{inverse FFT}(\text{FFT}(a_n) \cdot \text{FFT}(b_n))$. Now, back to the FFT...

For prime N , there exists a *generator* g of the multiplicative group modulo N : this means that $g^p \bmod N$ for $p = 0, \dots, N - 2$ produces all $n = 1, \dots, N - 1$ exactly once (but not in order!). Thus, we can write all non-zero n and k in the form $n = g^p$ and $k = g^{N-1-q}$ for some p and

q , and rewrite the DFT as

$$y_0 = \sum_{n=0}^{N-1} x_n, \quad (7)$$

$$y_{k \neq 0} = y_{g^{N-1-q}} = x_0 + \sum_{p=0}^{N-2} \omega_N^{g^{p+N-1-q}} x_{g^p}, \quad (8)$$

where (8) is exactly the cyclic convolution of $a_p = x_{g^p}$ with $b_p = \omega_N^{g^{N-1-p}}$. This convolution has non-prime length $N - 1$, and so we can evaluate it via the convolution theorem with FFTs in $O(N \log N)$ time (except for some unusual cases).

Further Reading

- D. N. Rockmore, “The FFT: An Algorithm the Whole Family Can Use,” *Comput. Sci. Eng.* **2** (1), 60 (2000). Special issue on “top ten” algorithms of century. See: <http://tinyurl.com/3wjkv> and <http://tinyurl.com/yvonp8>
- “Fast Fourier transform,” *Wikipedia: The Free Encyclopedia* (<http://tinyurl.com/5c6f3>). Edited by SGJ for correctness as of 10 Jan 2006 (along with subsidiary articles on C-T and other specific algorithms).
- “The Fastest Fourier Transform in the West,” a free FFT implementation obviously named by arrogant MIT graduate students. <http://www.fftw.org/>

Homework Problems

Problem 1: Prove that equation (2) really is the inverse of equation (1). Hint: substitute (1) into (2), interchange the order of the two sums, and sum the geometric series.

Problem 2: (a) Prove that for N a power of 2, we can recursively factorize $z^N - 1$ into polynomials of the form $z^M - 1$ and $z^{2M} + az^M + 1$ with a some real numbers and $|a| \leq 2$, for a decreasing sequence of M all the way down to $M = 1$. (The final quadratic factors for $M = 1$ can then be factored into conjugate pairs of roots of unity $e^{\frac{2\pi i}{N}k}$.) This gives an FFT algorithm due to Bruun (1978), distinct from Cooley-Tukey in that all of its multiplicative constants (a ’s) are *real* numbers until the very last step. **(b)** Apply this algorithm to write down the steps for a “Bruun” FFT of size $N = 8$, and count the number of required real additions and multiplications (not counting operations for x -independent constants like $2 \cdot \sqrt{2}$ that can be precomputed, and not counting trivial multiplications by ± 1 or $\pm i$). Compare this to the minimum known operation count of 56 total real additions and multiplications for $N = 8$ (achieved by the “split-radix” algorithm).