The Fourier transform is a conversion between two representations of a function: its list representation and its polynomial representation. Let's start with functions over the n-dimensional *Boolean cube* $\{+1,-1\}^n$. Its inputs are all 2^n bit strings of length n. We'll represent bits by the numbers $+1$ and -1, or $+,-$ in short. Here are three examples.

- The sum of two bits:

x_1x_2	$++$	$+-$	$-+$	$--$
$\mathrm{sum}_2(x_1,x_2)$	2	0	0	-2

- The maximum of two bits:

x_1x_2	$++$	$+-$	$-+$	$--$
$\mathrm{max}_2(x_1,x_2)$	$+1$	$+1$	$+1$	-1

- The majority of three bits:

$x_1x_2x_3$	$+++$	$++-$	$+-+$	$+--$	$-++$	$-+-$	$--+$	$---$
$\mathrm{maj}_3(x_1,x_2,x_3)$	$+1$	$+1$	$+1$	-1	$+1$	-1	-1	-1

These are the list representations of max_2 and maj_3. Any function on n-bit inputs can be defined by listing its 2^n evaluations in some predetermined order. It can also be specified as a *polynomial*:

$$\begin{aligned}\mathrm{sum}_2(x_1,x_2) &= x_1 + x_2\\ \mathrm{max}_2(x_1,x_2) &= \tfrac{1}{2} + \tfrac{1}{2}x_1 + \tfrac{1}{2}x_2 - \tfrac{1}{2}x_1x_2\\ \mathrm{maj}_3(x_1,x_2,x_3) &= \tfrac{1}{2}x_1 + \tfrac{1}{2}x_2 + \tfrac{1}{2}x_3 - \tfrac{1}{2}x_1x_2x_3.\end{aligned} \tag{1}$$

The monomials that are used to built up these polynomials are the 2^n *parity functions* $1, x_1, x_2, x_1x_2, x_1x_2x_3$, and so on. Every function $f(x_1,\ldots,x_n)$ can be uniquely represented as a linear combination of the 2^n possible parities in variables x_1 up to x_n.

The parities are indexed by the set S of variables that participate in it, e.g., x_1x_3 is indexed by the set $\{1,3\}$. We can also write down the polynomial representation by listing all 2^n coefficients $\hat{f}(S)$ of the parities, for example

S	$\varnothing$	$\{1\}$	$\{2\}$	$\{1,2\}$
$\widehat{\mathrm{sum}_2}(S)$	0	1	1	0
$\widehat{\mathrm{max}_2}(S)$	1/2	1/2	1/2	$-1/2$

and

S	$\varnothing$	$\{1\}$	$\{2\}$	$\{1,2\}$	$\{3\}$	$\{1,3\}$	$\{2,3\}$	$\{1,2,3\}$
$\widehat{\mathrm{maj}_3}(S)$	0	1/2	1/2	0	1/2	0	0	$-1/2$

These are the *Fourier representations* of sum_2, max_2, and maj_3, respectively. The Fourier transform is the conversion from the list representation of f to the Fourier representation $\hat{f}$.

This variant of the Fourier transform for functions over the Boolean cube is sometimes called the Fourier-Walsh transform. We'll talk shortly about functions over other domains.

1 The mathematics of the Fourier transform

Algebra

The polynomial representation of sum_2 was easy to figure out. After all sum_2 is defined to be the sum of its two inputs x_1 and x_2. But how did I calculate the Fourier transforms of max_2 and maj_3? The key fact

is that any function is a linear combination of *point functions* $\text{point}_{\mathbf{a}}$, namely functions that evaluate to one at a specific input $\mathbf{x} = \mathbf{a}$ and zero at all other inputs $\mathbf{x} \neq \mathbf{a}$.

For example, the $\max_2$ function is a linear combination of the four point functions point_{++}, point_{+-}, point_{-+} and point_{--} with coefficients

$$\max_2 = 1 \cdot \text{point}_{++} + 1 \cdot \text{point}_{+-} + 1 \cdot \text{point}_{-+} + (-1) \cdot \text{point}_{--}\,. \tag{2}$$

When we evaluate both sides at an input $\mathbf{x}$, say $\mathbf{x} = -+$, the corresponding point function $\text{point}_{-+}(\mathbf{x})$ is the only one that doesn't vanish. We can read off the value $\max_2(\mathbf{x})$ from its coefficient. This is precisely the list representation of $\max_2$, only written in different notation.

All we have to do now is figure out the polynomial/Fourier representation of the point functions. Once we have those we can add them up. The advantage of working with point functions is their multiplicativity. For example $\text{point}_{-+}(x_1, x_2) = \text{point}_{-}(x_1)\,\text{point}_{+}(x_2)$. We reduced the problem to finding a polynomial for the *univariate* point function $\text{point}_a(x)$ (when a and x are single bits). This is the linear function

$$\text{point}_a(x) = \frac{1 + ax}{2}.$$

When a and x are equal, both sides are $(1+1)/2 = 1$. When they are different, they are $(1-1)/2 = 0$. By multiplicativity,

$$\text{point}_{a_1,a_2}(x_1, x_2) = \frac{1 + a_1 x_1}{2} \cdot \frac{1 + a_2 x_2}{2}$$

and all we have to do is plug this into (2)

$$\max_2(x_1, x_2) = 1 \cdot \frac{1+x_1}{2} \cdot \frac{1+x_2}{2} + 1 \cdot \frac{1+x_1}{2} \cdot \frac{1-x_2}{2} + 1 \cdot \frac{1-x_1}{2} \cdot \frac{1+x_2}{2} - 1 \cdot \frac{1-x_1}{2} \cdot \frac{1-x_2}{2}.$$

and simplify to obtain formula (1).

This strategy works in general: The polynomial representation of a point function given by its 2^n evaluations $f(\mathbf{a})$ as $\mathbf{a}$ ranges over $\{-1,+1\}^n$ is obtained by simplifying the expression

$$f(\mathbf{x}) = \sum_{\mathbf{a} \in \{+1,-1\}^n} f(\mathbf{a})\,\text{point}_{\mathbf{a}}(\mathbf{x}) = \sum_{a_1,\ldots,a_n \in \{+1,-1\}} f(a_1,\ldots,a_n) \cdot \frac{1+a_1x_1}{2} \cdots \frac{1+a_nx_n}{2}. \tag{3}$$

Formula (3) gives one method of expanding f as a polynomial in its variables. Could other methods give different answers? Luckily, it turns out not. Representation (3) is unique.

Geometry

To understand why, we need to think geometrically. It helps to step back and look at a univariate f ($n = 1$). Such a function is specified by its two evaluations $(f(+), f(-))$. This is a two-dimensional vector. For example, $f(x) = 1 + \frac{1}{2}x$ corresponds to the vector $(3/2, 1/2)$. The point functions correspond to the standard basis $\text{point}_+ = (1,0)$ and $\text{point}_- = (0,1)$. In our example

$$f = \frac{3}{2} \cdot \text{point}_+ + \frac{1}{2} \cdot \text{point}_- = \frac{3}{2} \cdot \begin{bmatrix} 1 \\ 0 \end{bmatrix} + \frac{1}{2} \cdot \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

The parities 1 and x are themselves represented by the vectors $(+1,+1)$ and $(+1,-1)$. The polynomial representation of f is merely an expansion of the same vector f in this basis:

$$f = 1 \cdot 1 - \frac{1}{2} \cdot x = 1 \cdot \begin{bmatrix} +1 \\ +1 \end{bmatrix} + \frac{1}{2} \cdot \begin{bmatrix} +1 \\ -1 \end{bmatrix}.$$

The Fourier transform itself is nothing but a change of basis: It converts from a representation of f in the standard basis $\text{point}_{-1}, \text{point}_{+1}$ to a representation in the parity basis $1, x$ (see Figure 1).

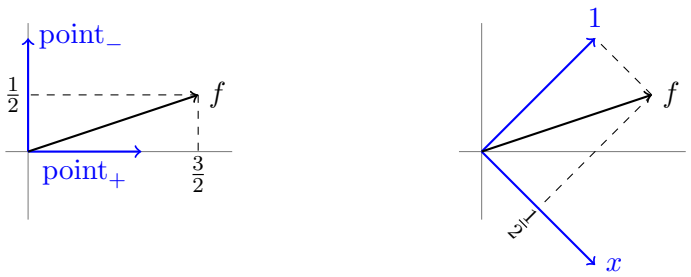


FIGURE 1: The Fourier representation as a change of basis.

The function $\max_2$ is specified by four values. Its list representation is the 4-dimensional vector $(+1,+1,+1,-1)$. This vector can no longer be visualized but the logic is the same. The standard basis vectors corresponding to the four coordinates are the point functions $\text{point}_{++} = (1,0,0,0)$, $\text{point}_{+-} = (0,1,0,0)$, $\text{point}_{-+} = (0,0,1,0)$, and $\text{point}_{--} = (0,0,0,1)$ and

$$\max_2 = \begin{bmatrix}+1\\+1\\+1\\-1\end{bmatrix} = +1\cdot\begin{bmatrix}1\\0\\0\\0\end{bmatrix} + 1\cdot\begin{bmatrix}0\\1\\0\\0\end{bmatrix} + 1\cdot\begin{bmatrix}0\\0\\1\\0\end{bmatrix} - 1\cdot\begin{bmatrix}0\\0\\0\\1\end{bmatrix}.$$

The polynomial expansion (1) represents the same vector in a different basis: the basis of parity functions $1 = (+1,+1,+1,+1), x_1 = (+1,+1,-1,-1), x_2 = (+1,-1,+1,-1), x_1x_2 = (+1,-1,-1,+1)$. Expression (1) is now the representation of $\max_2$ in this basis:

$$\begin{bmatrix}-1\\+1\\+1\\+1\end{bmatrix} = \frac{1}{2}\cdot\begin{bmatrix}+1\\+1\\+1\\+1\end{bmatrix} + \frac{1}{2}\cdot\begin{bmatrix}+1\\+1\\-1\\-1\end{bmatrix} + \frac{1}{2}\cdot\begin{bmatrix}+1\\-1\\+1\\-1\end{bmatrix} - \frac{1}{2}\cdot\begin{bmatrix}+1\\-1\\-1\\+1\end{bmatrix}.$$

From a geometric point of view, the Fourier transform is merely a change of basis formula. It converts from the the standard representation of f as a linear combination of the 2^n point functions $\text{point}_{++\cdots+}$, $\text{point}_{++\cdots-}$, …, $\text{point}_{--\cdots-}$ to the polynomial representation of f as a linear combination of the 2^n *parity functions*

$$\text{parity}_\varnothing = 1, \text{parity}_{\{1\}} = x_1, \text{parity}_{\{2\}} = x_2, \text{parity}_{\{1,2\}} = x_1x_2, \ldots.$$

As the *dimension* of both bases is the same and we have conversion formulas in both directions, linear algebra tells us that the parities must be linearly independent and therefore produce a unique Fourier representation.

If you did not follow this argument do not despair because geometry reveals a much more interesting reason behind the linear dependence of the parities: They are *mutually orthogonal*. When $n = 2$, for example, the dot product $x_1 \cdot (x_1x_2) = (+1,-1,+1,-1)\cdot(+1,-1,-1,+1)$ evaluates to zero. The same goes for the other five pairs of parities. There is an algebraic reason behind these cancellations: The dot product $x_1 \cdot (x_1x_2)$ is the sum of the products of the corresponding entries

$$x_1\cdot(x_1x_2) = \sum_{a_1,a_2} a_1(a_1a_2) = \sum_{a_1,a_2} a_1^2a_2 = \sum_{a_1,a_2\in\{-1,+1\}} a_2 = 0.$$

The x_1-part cancels out because $x_1^2 = 1$. This is where the cumbersome representation of bits by $+-$ pays off: When we evaluate the dot product of any two distinct parities, we sum up the evaluations of a third parity. This third parity keeps shifting between -1 and $+1$ and is therefore zero.

Lemma 1. *The 2^n parity functions are mutually orthogonal when viewed as 2^n-dimensional vectors.*

Orthogonality implies that parities are linearly independent as vectors. It tells us much more than that. The Fourier coefficients are the projections of f onto the basis of parities. For example

$$\widehat{\max_2}(\{2\}) = \frac{\max_2 \cdot x_2}{\|x_2\|^2} = \frac{(+1,+1,+1,-1)\cdot(+1,-1,+1,-1)}{\|(+1,-1,+1,-1)\|^2} = \frac{2}{4} = \frac{1}{2}.$$

All Fourier coefficients can be calculated in this way.

Probability

The last expression has a probabilistic interpretation: It is the average of the values obtained when f is shifted by the parity x_1x_3. In general,

$$\hat{f}(S) = \text{average value of } (f \cdot \text{parity}_S) = \frac{1}{2^n} \sum_{\mathbf{a}\in\{-1,+1\}^n} f(\mathbf{a})\,\text{parity}_S(\mathbf{a}). \tag{4}$$

When S is the empty set, $\text{parity}_\varnothing = 1$, so $\hat{f}(\varnothing)$ is nothing but the mean of f. For example, $\widehat{\max_2}(\varnothing) = (+1+1+1-1)/4 = 1/2$. If we had to summarize all of f by one number, this would be it. What do $\hat{f}(\{1\})$ and $\hat{f}(\{5,7\})$ tell us? We'll explain shortly.

Computation

Computing the Fourier Transform means taking the list of 2^n values $f(\mathbf{a})$ and producing the list of 2^n Fourier coefficients $f(S)$. Both the input and output have size $N = 2^n$ so it makes sense to express the complexity in terms of this N. Evaluating (4) would entail N^2 operations: It takes $N = 2^n$ additions to calculate each Fourier coefficient and there are $N = 2^n$ of them to calculate.

There is a faster recursive algorithm. It is easiest to describe it in the language of polynomials. Suppose that we have figured out the polynomial representations of the functions

$$f_+(x_1,\ldots,x_{n-1}) = f(x_1,\ldots,x_{n-1},+1) \qquad \text{and} \qquad f_-(x_1,\ldots,x_{n-1}) = f(x_1,\ldots,x_{n-1},-1).$$

How can we combine them into a single polynomial representation of f? What we want f to do is produce $f_+(x_1,\ldots,x_{n-1})$ when x_n is $+1$ and $f_-(x_1,\ldots,x_{n-1})$ when x_n is -1. We can resort to a tried and tested trick: Write f as a linear combination of two point functions.

$$\begin{aligned} f &= f_+ \cdot \text{point}_+(x_n) + f_- \cdot \text{point}_-(x_n) \\ &= f_+ \cdot \frac{1+x_n}{2} + f_- \cdot \frac{1-x_n}{2} \\ &= \frac{1}{2}(f_+ + f_-) + \frac{1}{2}(f_+ - f_-)x_n. \end{aligned}$$

The term $\frac{1}{2}(f_+ + f_-)$ contains those monomials that exclude x_n, while the term $\frac{1}{2}(f_+ - f_-)x_n$ contains those monomials that include it. The corresponding formula for the monomial/Fourier coefficients is

$$\hat{f}(S) = \frac{1}{2}(\hat{f}_+(S) + \hat{f}_-(S)) \quad \text{and} \quad \hat{f}(S\cup\{n\}) = \frac{1}{2}(\hat{f}_+(S) - \hat{f}_-(S)) \tag{5}$$

for every subset S of $\{1,\ldots,n-1\}$.

Algorithm *FFW* (Fast Fourier-Walsh Transform)
Input: A list representation of $f\colon \{-1,+1\}^n \to \mathbb{R}$.
1 If $n = 0$, output the value $f()$.
2 Calculate $\hat{f}_+ = FFW(f_+)$ and $\hat{f}_- = FFW(f_-)$.
3 Combine $\hat{f}_+$ and $\hat{f}_-$ into $\hat{f}$ using (5) and output it.

The correctness of the algorithm follows from induction by (5). The base case $n = 0$ holds because a function with no inputs is a constant, so its list and Fourier representations are identical.

Here is an example run of FFW when $f = \max_2$. The algorithm needs to calculate the Fourier-Walsh transforms of $f_+ = \max_2(x_1, +)$ and $f_- = \max_2(x_1, -)$ first. f_+ decomposes into $f_{++} = 1$ and $f_{+-} = 1$, so $f_+ = \frac{1}{2}(f_{++} + f_{+-}) + \frac{1}{2}(f_{++} - f_{+-})x_1 = 1$. Similarly, $f_- = x_1$. In step 3 they are combined into

$$f = \tfrac{1}{2}(f_+ + f_-) + \tfrac{1}{2}(f_+ - f_-)x_2 = \tfrac{1}{2}(1 + x_1) + \tfrac{1}{2}(1 - x_1)x_2 = \tfrac{1}{2} + \tfrac{1}{2}x_1 + \tfrac{1}{2}x_2 - \tfrac{1}{2}x_1x_2.$$

The complexity of FFW can be estimated by solving a recurrence. Let $T(n)$ be the number of operations (additions, subtractions, halvings) it takes to calculate an n-variate FFW. FFW solves an n-variate problem by splitting it into two $(n-1)$-variate ones (step 2) and then combining their answers. Combining takes two operations per subset S in (5) for a total of 2^n operations. Thus $T(n)$ equals $2T(n-1)$ plus 2^n, with base case $T(0) = 0$. This solves to $T(n) = n \cdot 2^n$, which is the same as $N \log N$. This is a huge improvement over N^2.

2 Polynomials, complexity, and approximation

A key goal of the theory of computation is to characterize the *computational* complexity of various functions that come up. The Fourier representation is sometimes a helpful indicator of complexity.

One natural complexity measure of a polynomial is its degree. Low-degree polynomials have relatively few "degrees of freedom." A linear function over the n-dimensional Boolean cube is completely specified by its mean $\hat{f}(\varnothing)$ and its n level-1 Fourier coefficients $\hat{f}(\{1\}), \ldots, \hat{f}(\{n\})$. More generally, a degree-d polynomial is determined by its Fourier coefficients of size d or less. When d is fixed and n grows, there are about $(n/d)^d$ such coefficients. They grow at a polynomial rate. This is much less than the 2^n values it takes to specify the function as a whole.

The compactness of polynomial representations can be exploited by algorithms. Suppose we are given input-output examples $(x, f(x))$ for some unknown function f. If f is a degree-d polynomial we can try to reconstruct the polynomial representing f by interpolating its Fourier coefficients from these examples. This amounts to solving a linear system, which can be accomplished in a number of steps that is polynomial in n^d, provided there are sufficiently many linearly independent equations in the system.

Some simple types of computations are naturally captured by low-degree polynomials. As an example, the function

$$f(x, y, z) = \text{“if } x = 1 \text{ then } y \text{ else } z\text{”}$$

has degree 2. It is unlike the 72% of the other Boolean functions on 3 inputs whose degree is 3. It has degree 2 because all point functions that show up in its list representation

$$f = y \cdot \text{point}_+(x) + z \cdot \text{point}_-(x) = y \cdot \frac{1+x}{2} + z \cdot \frac{1-x}{2}$$

have degree 2 or less.

A *decision tree* is an arbitrary nested if-then-else program of this form. The *depth* of the decision tree is the number of nested levels plus one. By the same reasoning any depth-d decision tree is a degree-d polynomial. Because it is a low-degree polynomial, some algorithmic tasks such as learning one from random examples are easier for decision trees than for more complex functions.

A more complex type of function is a formula in disjunctive normal form, namely an AND of terms, each of which is an OR of literals (variables or negated variables) such as

$$(x \text{ AND } (\text{NOT } y) \text{ AND } z) \text{ OR } (x \text{ AND } w) \text{ OR } (u \text{ AND } w \text{ AND } (\text{NOT } z)).$$

Terms could represent qualification requirements. For example a student is eligible for a scholarship if they are in an upper year and in CS and have a GPA of 3 or more, or if they are Francophone and do not live too close to home, or if they are not in an upper year but won a contest. Such formulas do not have

a low-degree representation. It turns out, however, that they can be "approximated" by polynomials of fairly low degree.

To illustrate how this works let's look at the formula

$$f = (x_1 \text{ AND } x_2) \text{ OR } (x_3 \text{ AND } x_4) \text{ OR } (x_5 \text{ AND } x_6).$$

Its complete polynomial representation is too long to write down. The lowest two levels are

$$f = -.15625 + .28125(x_1 + x_2 + x_3 + x_4 + x_5 + x_6) + \text{higher order terms.}$$

The consecutive terms in the polynomial expansion provided finer and finer approximations of f:

$$\begin{aligned} &\text{constant approximation } c = -.15625 \\ &\text{linear approximation } \ell(x) = -.15625 + .28125(x_1 + x_2 + x_3 + x_4 + x_5 + x_6) \end{aligned}$$

and so on. These approximations are best possible in the following sense. Among all constants $c = \hat{f}(0)$ is the one that minimizes the average $(f(\mathbf{x}) - c)^2$ of the square error. Among all linear functions, $\ell(x)$ is the one that minimizes the average $(f(\mathbf{x}) - \ell(\mathbf{x}))^2$ of the square error.

The reason is orthogonality. If we split f into a low-degree part plus a high-degree part, for instance

$$\text{maj}_3 = \underbrace{\tfrac{1}{2} + \tfrac{1}{2}x_1 + \tfrac{1}{2}x_2 + \tfrac{1}{2}x_3}_{\text{degree} \le 1} + \underbrace{\tfrac{1}{2}x_1x_2x_3}_{\text{degree} > 1}$$

then the two parts are orthogonal because all their constituent parities are. No matter how we approximate f by a low-degree polynomial, the high-degree part must be incurred as error.

There is another important consequence of the parities being orthogonal. By Pythagoras' theorem, the square length of f equals the sum of the square lengths of its projections along the parities:

$$\|f\|^2 = \text{sum of } \left(\frac{f \cdot \text{parity}_S}{\|\text{parity}_S\|} \right)^2$$

As all parities have length $\sqrt{2^n}$, this formula simplifies to

$$\frac{1}{2^n}\|f\|^2 = \text{sum of } \hat{f}(S)^2.$$

The left-hand side is the averaged sum of $f(\mathbf{x})^2$. This is *Parseval's identity*:

$$\text{average of } f(\mathbf{x})^2 = \text{sum of } \hat{f}(S)^2.$$

One weakness of polynomial degree as a measure of complexity is that dense parity functions, such as the parity of all the bits, do not qualify as simple. This is contrary to experience as parities are easy to compute in practice. Still, their high degree indicates a challenge for "analytic" algorithms like gradient descent.

3 The modular Fourier transform

The Fourier transform for functions on the Boolean cube can be useful for reasoning about computational and mathematical structures. There is another type of Fourier transform whose significance is indisputable in engineering.

The polynomial representation f with a one-bit input ($n = 1$, $N = 2$ in the notation of Section 1) is

$$f(x) = \hat{f}(0) + \hat{f}(1)x. \tag{6}$$

The coefficients $\hat{f}(0)$ and $\hat{f}(1)$ are determined by evaluating f at inputs -1 and $+1$. They can also be in principle extrapolated from any other two values of f, say $f(-3)$ and $f(2)$. A linear function is specified by any two of its values. They force its value at any other point. If we want to represent something with a third "degree of freedom", we can throw in a quadratic term:

$$f(x) = \hat{f}(0) + \hat{f}(1)x + \hat{f}(2)x^2. \tag{7}$$

The Boolean Fourier transform rested on some special features of the evaluation points $+1$ and -1 in (6). What made those points stand out over the rest is their *group structure*. Multiplication stays within the set.

$\cdot$	1	x
1	1	x
x	x	1

For quadratic polynomials, the same role is played by the *complex* cube roots of unity $1, \omega = e^{2\pi i/3}, \omega^2 = e^{4\pi i/3}$. They also form a group under complex multiplication:

$\cdot$	1	ω	ω^2
1	1	ω	ω^2
ω	ω	ω^2	1
ω^2	ω^2	1	ω

The arithmetic is the same as addition modulo 3, but in the exponent. The list representation for g is now specified by its values at the three roots of unity, for example:

x	1	ω	ω^2
$f(x)$	0	1	-2

Its corresponding vector representation is $(0, 2, -2)$. The basis functions $1, x, x^2$ are represented by the vectors

$$1 = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \qquad x = \begin{bmatrix} 1 \\ \omega \\ \omega^2 \end{bmatrix} = \begin{bmatrix} 1 \\ e^{2\pi i/3} \\ e^{4\pi i/3} \end{bmatrix} \qquad x^2 = \begin{bmatrix} 1 \\ \omega^2 \\ \omega \end{bmatrix} = \begin{bmatrix} 1 \\ e^{4\pi i/3} \\ e^{2\pi i/3} \end{bmatrix}.$$

They are orthogonal under the *complex* dot product

$$\mathbf{a} \cdot \overline{\mathbf{b}} = (a_1, a_2, a_3) \cdot (\overline{b_1}, \overline{b_2}, \overline{b_3}) = a_1\overline{b_1} + a_2\overline{b_2} + a_3\overline{b_3}.$$

The second component must be conjugated. (This is a common source of calculation errors.) Apart from this, the algebra, geometry, and probability carries over from the Boolean cube.

Let's calculate the Fourier transform of f to see this in action:

$$\begin{aligned}
\hat{f}(0) &= \text{average of } f = \frac{1}{3}\big(0 + 1 + (-2)\big) = -\frac{1}{3}, \\
\hat{f}(1) &= \text{average of } f \cdot \overline{x} = \frac{1}{3}\big(0 \cdot 1 + 1 \cdot \overline{\omega} + (-2) \cdot \overline{\omega^2}\big) = \frac{1}{3}(\omega^2 - 2\omega) = \frac{1}{6} + \frac{\sqrt{3}}{2} \cdot i, \\
\hat{f}(2) &= \text{average of } f \cdot \overline{x^2} = \frac{1}{3}\big(0 \cdot 1 + 1 \cdot \overline{\omega^2} + (-2) \cdot \overline{\omega}\big) = \frac{1}{3}(\omega - 2\omega^2) = \frac{1}{6} - \frac{\sqrt{3}}{2} \cdot i.
\end{aligned}$$

An annoyance of this Fourier transform are the complex numbers, even when the evaluations are real. One convenience is that roots of unity conjugate nicely. The conjugate of ω^j is ω^{-j}.

In general, the list representation of a degree-N polynomial f

$$f(x) = \hat{f}(0) + \hat{f}(1)x + \cdots + \hat{f}(N-1)x^{N-1} \tag{8}$$

is its evaluation at the N-th complex roots of unity, namely $(f(1), f(\omega), \ldots, f(\omega^{N-1}))$, where $\omega = e^{2\pi i/N}$. The Fourier coefficients are given by

$$\hat{f}(s) = \text{average of } (f \cdot \overline{x}^s) = \frac{1}{N}\big(f(1) + f(\omega)\cdot\omega^{-s} + f(\omega^2)\cdot\omega^{-2s} + \cdots + f(\omega^{N-1})\cdot\omega^{(N-1)s}\big). \tag{9}$$

The basis functions $x^s = (1, \omega^s, \omega^{2s}, \cdots, \omega^{(N-1)s})$ are orthogonal as vectors under the complex dot product. Their role is analogous to that of the parities parity_S on the Boolean cube. The size of set S corresponds to the magnitude of the number s. The zeroth coefficient $\hat{f}(0)$ is the average value of f, which is the best approximation among all constants c with respect to the average of $|f(x) - c|^2$. The first-level truncation $\hat{f}(0) + \hat{f}(1)x$ is the best approximation among all linear functions $\ell(x)$ with respect to the average of $|f(x) - \ell(x)|^2$, and so on.

4 The Fast Fourier Transform

Calculating the modular Fourier representation with formula (9) takes time quadratic in N, just like for the Fourier-Walsh transform over the Boolean cube. Again, there is a faster algorithm, provided N is a power of two. The strategy is to separately calculate the even and odd coefficients in (8). To be specific, let's take $N = 8$. Then f can be split into even and odd powers:

$$f(x) = \big(\hat{f}(0) + \hat{f}(2)x^2 + \hat{f}(4)x^4 + \hat{f}(6)x^6\big) + \big(\hat{f}(1)x + \hat{f}(3)x^3 + \hat{f}(5)x^5 + \hat{f}(7)x^7\big).$$

Those parts are the respective polynomial representations of $(f(x) + f(-x))/2$ and $(f(x) - f(-x))/2$:

$$\begin{aligned}\frac{f(x)+f(-x)}{2} &= \hat{f}(0) + \hat{f}(2)x^2 + \hat{f}(4)x^4 + \hat{f}(6)x^6\\ \frac{f(x)-f(-x)}{2} &= \hat{f}(1)x + \hat{f}(3)x^3 + \hat{f}(5)x^5 + \hat{f}(7)x^7\\ &= x(\hat{f}(1) + \hat{f}(3)x^2 + \hat{f}(5)x^4 + \hat{f}(7)x^6).\end{aligned}$$

This suggests setting up the functions

$$f_+(x^2) = \frac{f(x)+f(-x)}{2} \qquad \text{and} \qquad f_-(x^2) = \frac{f(x)-f(-x)}{2x}, \tag{10}$$

calculating their polynomial expansions recursively, and adding up the resulting polynomials.

As x ranges over the 8-th complex roots of unity $e^{2\pi ij/8}$, x^2 ranges over the 4-th roots $e^{2\pi ij/4}$. The functions f_+ and f_- must therefore have the intended Fourier expansions even modulo 4, namely

$$\begin{aligned}f_+(y) &= \hat{f}(0) + \hat{f}(2)y + \hat{f}(4)y^2 + \hat{f}(6)y^3, \qquad &\text{or } \hat{f}_+(j) = \hat{f}(2j)\\ f_-(y) &= \hat{f}(1) + \hat{f}(3)y + \hat{f}(5)y^2 + \hat{f}(7)y^3, \qquad &\text{or } \hat{f}_-(j) = \hat{f}(2j+1).\end{aligned}$$

Conversely,

$$\hat{f}(j') = \begin{cases}\hat{f}_+(j'/2), & \text{if } j' \text{ is even,}\\ \hat{f}_-((j'-1)/2), & \text{if } j' \text{ is odd.}\end{cases} \tag{11}$$

Algorithm *FFT* (The Fast Fourier Transform)
Input: N (a power of two) and the list $f(1), f(\omega), \ldots, f(\omega^{N-1})$.
1 If $N = 1$, output the value $f(1)$.
2 Set $\omega = e^{2\pi i/N}$.
3 Calculate the list representations of $f_+, f_-\colon \{1, \omega^2, \omega^4, \ldots, \omega^{2(N-1)}\} \to \mathbb{C}$ using (10).
4 Calculate $\hat{f}_+ = FFT(N/2, f_+)$ and $\hat{f}_- = FFT(N/2, f_-)$.
5 Construct $\hat{f}\colon \{0, \ldots, N-1\} \to \mathbb{C}$ using (11) and output it.

The FFT algorithm calculates the Fourier transform of f using $O(N \log N)$ operations with complex numbers (additions, subtractions, and scalings).

Example Let $N = 4$. The fourth roots of unity are 1, $\omega = i$, $\omega^2 = -1$ and $\omega^3 = -i$. Let's take

x	1	i	-1	$-i$
$g(x)$	1	0	-2	-1

We want to derive the polynomial representation

$$\begin{aligned} f(x) &= \hat{f}(0) + \hat{f}(1)x + \hat{f}(2)x^2 + \hat{f}(3)x^3 \\ &= (\hat{f}(0) + \hat{f}(2)x^2) + x(\hat{f}(1) + \hat{f}(3)x^2) \\ &= f_+(x^2) + xf_-(x^2), \end{aligned} \tag{12}$$

where $f_+(y)$ and $f_-(y)$ are functions of the second roots of unity 1 and -1. Their list representations are calculated in line 3 as

$$f_+(1) = \frac{f(1) + f(-1)}{2} = -.5 \qquad f_-(1) = \frac{f(1) - f(-1)}{2} = 1.5$$
$$f_+(-1) = \frac{f(i) + f(-i)}{2} = -.5 \qquad f_-(-1) = \frac{f(i) - f(-i)}{2i} = -.5i.$$

Even though f was real-valued, f_- is complex-valued.

In step 4 the FFT algorithm recursively calculates the polynomial representations of f_+ and f_-. As the domain of these functions is $-1, 1$ the outcome is the same as for the Fourier-Walsh transform. The representations are

$$f_+(y) = -.5, \qquad f_-(y) = (.75 - .25i) + (.75 + .25i) \cdot y.$$

Finally, in step 5 these representations are plugged into (12) to obtain the polynomial representation of f:

$$\begin{aligned} f(x) &= -.5 + x \cdot \big((.75 - .25i) + (.75 + .25i) \cdot x^2\big) \\ &= -.5 + (.75 - .25i) \cdot x + (.75 + .25i) \cdot x^3, \end{aligned}$$

or $\hat{f}(0) = -.5, \hat{f}(1) = .75 - .25i, \hat{f}(2) = 0, \hat{f}(3) = .75 + .25i$.

5 Variants

The Cosine Transform

In signal processing applications, the function $f(\omega^t)$ might represent a signal like the amplitude of a sound sampled at times $t = 0$, 1, up to $N - 1$:

$$\text{signal at time } t = f(\omega^t) = \hat{f}(0) + \hat{f}(1)\omega^t + \cdots + \hat{f}(N-1)\omega^{t(N-1)}, \quad \omega = e^{2\pi i/n}.$$

There are two annoyances with this representation. First, even though the signal is real-valued, the Fourier coefficients and the basis functions are complex. Second, the function $f(\omega^t)$ is periodic modulo n, but signals like sound waves are not naturally periodic. It is likely that there will be a discontinuity when the signal wraps around from $t = N - 1$ to $t = N \equiv 0$ (see Figure 2). Discontinuities create unnatural artifacts in the Fourier expansion.

There is a beautiful trick that eliminates both problems. Before taking its Fourier transform, concatenate the signal with its mirror image (see Figure 2). All the Fourier coefficients become real. The basis functions reduce to their real part, which is a vector of cosines:

$$g(\omega^t) = 2\hat{g}(0) + 4\hat{g}(1) \cdot \cos\Big(\frac{\pi t}{2N}\Big) + 4\hat{g}(1) \cdot \cos\Big(\frac{3\pi t}{2N}\Big) + \cdots + 4\hat{g}(N-1) \cdot \cos\Big(\frac{(2N-1)\pi t}{2N}\Big),$$

where $\omega = e^{\pi i t/2N}$ and t is odd modulo $4N$. The index s in $\hat{g}(s)$ represents the *frequency* of the corresponding cosine wave (see Figure 3). The low-frequency parts (small js) capture the stable part of the signal, while the high-frequency ones (large js) capture the oscillating part.

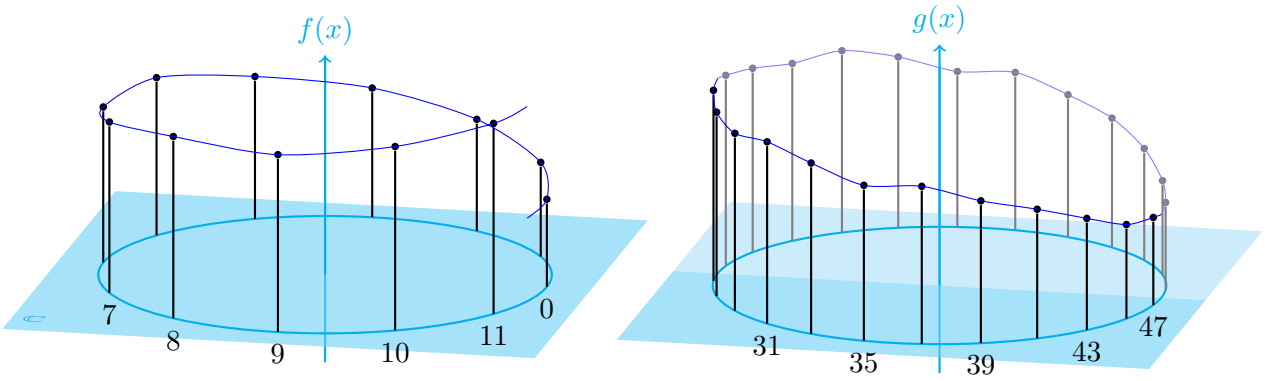


FIGURE 2: The cosine transform. (a) Evaluations of $f(e^{2\pi it/12})$, $t = 0, \ldots, 11$. The signal exhibits a discontinuity as it wraps from $t = 11$ to $t = 12 \equiv 0$. (b) g is obtained by concatenating f with its mirror image and shifting it so that it is symmetric about the real axis. The discontinuities are smoothed out and g is symmetric with respect to complex conjugation ($g(\overline{x}) = g(x)$).

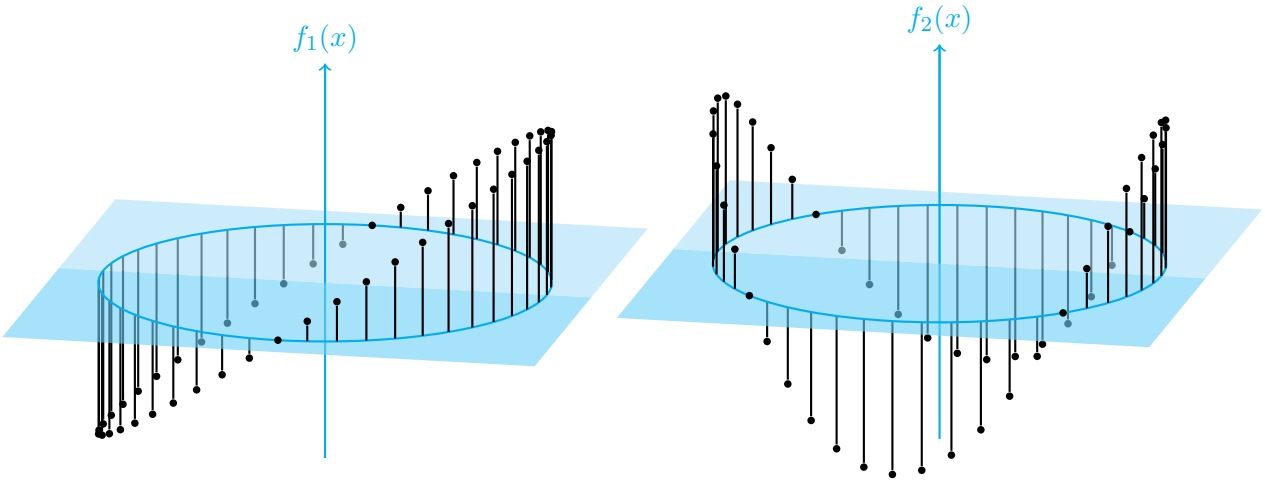


FIGURE 3: The basis functions (a) $g_1(\omega^t) = \cos(\pi t/2N)$ and (b) $g_2(\omega^t) = \cos(2\pi t/2N)$.

Two-dimensional transforms

There are Fourier transforms that apply to two and more-dimensional signals like images. They are constructed by taking the "product" of one-dimensional Fourier transfoms. Suppose we know a Fourier representation for functions in variable x and we also know a Fourier representation for functions in y. Then functions in x and y have a unique representation

$$f(x, y) = \sum_{s,t} \hat{f}(s, t) x^s y^t.$$

The vector representations of the monomials $x^s \cdot y^t$ are obtained by taking the outer product of the vectors representing x^i and y^j individually. The (i, j)-th entry of the vector representing $x^s y^t$ is the product of the i-th entry of x^s times the j-th entry of y^t. The Fourier coefficients are calculated by a formula analogous to (9): $\hat{f}(s, t)$ is the average of $f(x, y)x^{-s}y^{-t}$ for x, y ranging over their respective domains.

If x and y range over the N-th and M-th roots of unity, respectively, the Fourier representation of $f(x, y)$ has the form

$$\text{signal at position } (x, y) = f(e^{2\pi ix/N}, e^{2\pi iy/M}) = \sum_{j,k} \hat{f}(s, t) e^{2\pi i(xs/N + yt/M)}$$

with j ranging from 0 to $N-1$ and k ranging from 0 to $M-1$. The Fourier coefficient $\hat{f}(s,t)$ is obtained by averaging $f(e^{2\pi ix/N}, e^{2\pi iy/M})e^{-2\pi i(xs/N+yt/M)}$ over $x \in \{0,\ldots,N-1\}$ and $y \in \{0,\ldots,M-1\}$.

The Fourier-Walsh transform over the Boolean cube $\{-1,1\}^n$ is compatible with this construction: It is the n-th power of the Fourier transform for functions on a single bit $\{-1,1\}$.

Continuous transforms

It is sometimes useful to think of the signal as a continuous function. The amplitude of a sound wave can in principle be measured at any continuous instant t. Discrete-time approximations that are amenable to data processing can be obtained by sampling the signal at regularly spaced points. The more frequent the sampling, the more precise the approximation is. We would expect the same to be true for its Fourier transform.

Indeed, as N gets larger and larger, the Fourier transform modulo N approaches a continuous Fourier transform. Its inputs are functions f that take values on the continuous unit circle. Any "reasonable" f of this type can be expanded as a possibly infinite *Fourier series* $f(\omega) = \sum \hat{f}(j)\omega^j$. Here ω denotes any point on the complex unit circle and j ranges over all integers. The Fourier coefficient $\hat{f}(j)$ is the average of $f(\omega)\omega^{-j}$ with ω sampled uniformly from the unit circle.

Continuous Fourier representations are insightful for understanding properties like sensitivity to continuous noise: By how much does the value of f change typically if its input is perturbed by a small amount? Just like for functions over the Boolean cube, the higher frequency terms in the Fourier expansion yield more precise but also more complex approximations of f.

For "standard" mathematical functions f like the pulse and the Gaussian curve the Fourier coefficients can be calculated exactly via integration, often leading to insightful infinite series representations.

6 Quantum Fourier sampling

A quantum computer is an imaginary computational device. It is believed that if one is ever built, it will be dramatically more effective than any existing, classical computer for certain types of problems. Much of this excitement comes from the conjectured ability of quantum computers to "perform" Fourier transforms exponentially faster than classical computers can.

The Fast Fourier Transform is already very efficient. It runs in time $O(N \log N)$ on a N-input function. No algorithm can do better than time N because the least it has to do it inspect all values of f. Quantum algorithms, like classical ones, are subject to this law. There is little room for improvement in algorithms for *computing* the Fourier transform. (Having said that, the discovery of an $O(N)$-time classical or quantum algorithm for the Fourier Transform would be a tremendous breakthrough!)

There is a different problem related to Fourier transforms called Fourier Sampling. The starting point is Parseval's identity (**??**). Assume f is a function whose average square value is 1. Boolean ± 1-valued functions are like that. The squared Fourier coefficients specify a probability distribution over subsets S of $\{1,\ldots,n\}$: Set S is picked with probability $\hat{f}(S)^2$. For example, the distribution induced by maj_3 is uniform over the sets $\{1\}$, $\{2\}$, $\{3\}$, and $\{1,2,3\}$.

Fourier Sampling: Given the list representation of f, sample set S with probability $\hat{f}(S)^2$.

We would still expect that to solve Fourier Sampling, a computer would have to at the least inspect all N values of f. This is indeed the case for classical computers. It is no longer true for quantum computers. Quantum Fourier Sampling runs in $\log N$ elementary quantum steps!

It may be difficult to understand this exponential speedup without talking about quantum computers in detail but let me give a try. The state of a classical computer is determined by the contents of its memory. A computer with n bits of memory must be in one of the 2^n possible memory states. For example if $n = 2$ the possible states are 00, 01, 10, and 11. In contrast, a quantum computer can be in a *superposition* of these states. Mathematically, a superposition is any vector of unit norm in the 4-dimensional vector space

spanned by the standard basis vectors labeled by the four memory states. These four vectors are usually written as $|00\rangle, |01\rangle, |10\rangle, |11\rangle$ in ket notation. In general, a (pure) quantum state of an n-qubit quantum computer is a superposition of the 2^n standard basis states $|00\cdots0\rangle, |00\cdots1\rangle, |11\cdots1\rangle$.

Given a list representation of f, a quantum computer can create the quantum superposition state $|s\rangle = \frac{1}{2}(f(++)|00\rangle + f(-+)|10\rangle + f(+-)|01\rangle + f(--)|11\rangle)$ in one step. This is a valid quantum state, namely a unit vector, because its length is precisely the squared average of f. In general, the quantum computer should start in state

$$|s\rangle = \frac{1}{\sqrt{N}} \cdot \sum_{x \in \{\pm1\}^n} f(x)|x \text{ in 0/1 representation}\rangle. \tag{13}$$

Suppose it could process this state into the Fourier basis state $|t\rangle = \hat{f}(\varnothing)|00\rangle + \hat{f}(\{1\})|10\rangle + \hat{f}(\{2\})|01\rangle + \hat{f}(\{1,2\})|11\rangle$. Quantum physics tells us that the desired outcome—sample S with probability $\hat{f}(S)^2$—will be achieved merely by *observing* the quantum state: A measurement of a quantum state collapses it to a single (observable) outcome with probability proportional to the square of the coefficient in front of it (the amplitude).

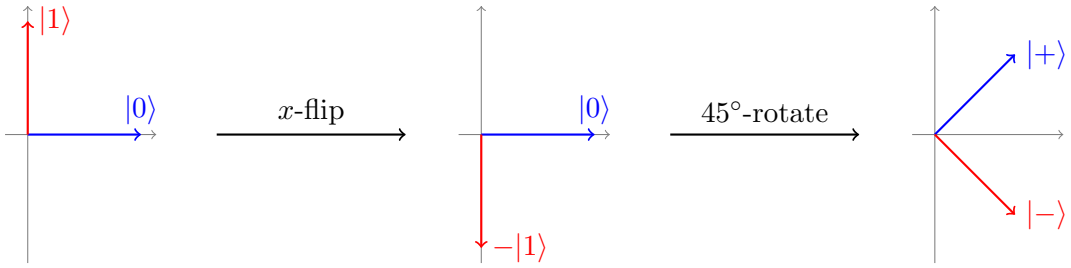
What remains to be done is to instruct the quantum computer to move from state $|s\rangle$ to state $|t\rangle$. By the product-ive nature of the Fourier transform, this can be accomplished by moving every qubit independently. (Algebraically, the n-dimensional Fourier transform is the n-th tensor power of the one-dimensional Fourier transform.) When $n = 1$, this means transitioning from the list-of-values basis

$$|0\rangle = \begin{bmatrix}1\\0\end{bmatrix} \qquad |1\rangle = \begin{bmatrix}0\\1\end{bmatrix}$$

to the Fourier basis

$$|+\rangle = \frac{1}{\sqrt{2}}\begin{bmatrix}+1\\+1\end{bmatrix} \qquad |-\rangle = \frac{1}{\sqrt{2}}\begin{bmatrix}+1\\-1\end{bmatrix}.$$

The "machine language" instructions of a quantum computer are *unitary transformations*. These are linear mappings that preserve unit length. Examples include rotations and sign flips. The desired change of basis can be realized by a x-sign flip followed by a 45-degree rotation:



The algebraic specification of the instruction is multiplication by the *Walsh-Hadamard matrix*

$$H = \frac{1}{\sqrt{2}}\begin{bmatrix}+1 & +1\\+1 & -1\end{bmatrix}.$$

Algorithm *QFS* (Quantum Fourier-Walsh Sampling)
Input: List representation of $f\colon \{\pm1\}^n \to \{\pm1\}$.
1 Prepare the quantum start state (13).
2 Apply H to each of the n qubits.
3 Measure the state and report the outcome.

A variant of Quantum Fourier Sampling works for the modular Fourier transform. One important application of it is period finding. Suppose the input function f is periodic modulo some unknown number q. The period shows in the Fourier transform of f: Most of its "weight" is concentrated around multiples of N/q, provided q is sufficiently smaller than N. The output of Quantum Fourier Sampling on such functions is likely to be very close to some multiple of N/q. It is possible to find a good approximation of q itself from a few such samples.

The most spectacular application of period finding is factoring. Given a product $N = pq$ of two primes it is possible to cook up a function that is almost periodic modulo one of the prime factors of N. Quantum Fourier Sampling then recovers one of the prime factors of N up to some small error. In 1994 Peter Shor showed how the error can be eliminated. His discovery dramatically raised the stakes in quantum computing. The ability to factor efficiently would break most existing encryption infrastructure.

Despite enormous well-funded efforts since then, a quantum computer is not yet in the cards. Many scientists believe it will one day be built. This quantum computer will be able to run the precise instructions required by algorithms like Quantum Fourier Sampling on inputs of arbitrarily large scale.

Some researchers believe that fundamental obstacles forbid the physical realization of quantum computers. Mathematician Gil Kalai is among the most outspoken ones. His main argument is that quantum systems are too susceptible to noise to do anything useful. How does he argue his case? You guessed it—the proof is in the Fourier Transform.