

Chapter 8

Fourier Analysis

We all use Fourier analysis every day without even knowing it. Cell phones, disc drives, DVDs, and JPEGs all involve fast finite Fourier transforms. This chapter discusses both the computation and the interpretation of FFTs.

The acronym FFT is ambiguous. The first F stands for both “fast” and “finite.” A more accurate abbreviation would be FFFT, but nobody wants to use that. In MATLAB the expression `fft(x)` computes the finite Fourier transform of any vector `x`. The computation is fast if the integer `n = length(x)` is the product of powers of small primes. We discuss this algorithm in section 8.6.

8.1 Touch-Tone Dialing

Touch-tone telephone dialing is an example of everyday use of Fourier analysis. The basis for touch-tone dialing is the Dual Tone Multi-Frequency (DTMF) system. The program `touchtone` demonstrates how DTMF tones are generated and decoded. The telephone dialing pad acts as a 4-by-3 matrix (Figure 8.1). Associated with each row and column is a frequency. These basic frequencies are

```
fr = [697 770 852 941];  
fc = [1209 1336 1477];
```

If `s` is a character that labels one of the buttons on the keypad, the corresponding row index `k` and column index `j` can be found with

```
switch s  
  case '*' , k = 4; j = 1;  
  case '0' , k = 4; j = 2;  
  case '#' , k = 4; j = 3;  
  otherwise,  
    d = s-'0'; j = mod(d-1,3)+1; k = (d-j)/3+1;  
end
```

September 21, 2013

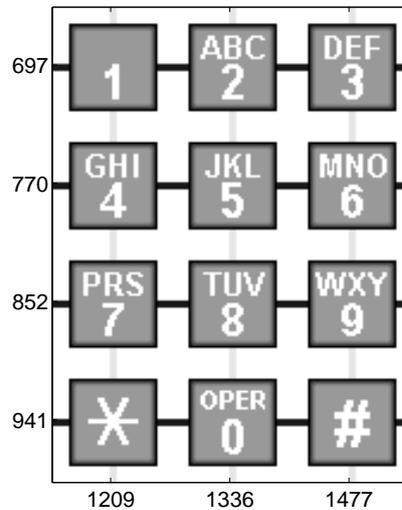


Figure 8.1. Telephone keypad.

A key parameter in digital sound is the sampling rate.

$$F_s = 32768$$

A vector of points in the time interval $0 \leq t \leq 0.25$ at this sampling rate is

$$t = 0:1/F_s:0.25$$

The tone generated by the button in position (k, j) is obtained by superimposing the two fundamental tones with frequencies $fr(k)$ and $fc(j)$.

$$\begin{aligned} y_1 &= \sin(2\pi \cdot fr(k) \cdot t); \\ y_2 &= \sin(2\pi \cdot fc(j) \cdot t); \\ y &= (y_1 + y_2)/2; \end{aligned}$$

If your computer is equipped with a sound card, the MATLAB statement

```
sound(y, Fs)
```

plays the tone.

Figure 8.2 is the display produced by `touchtone` for the '1' button. The top subplot depicts the two underlying frequencies and the bottom subplot shows a portion of the signal obtained by averaging the sine waves with those frequencies.

The data file `touchtone.mat` contains a recording of a telephone being dialed. Is it possible to determine the phone number by listening to the signal generated? The statement

```
load touchtone
```

loads a structure `y` into the workspace. The statement

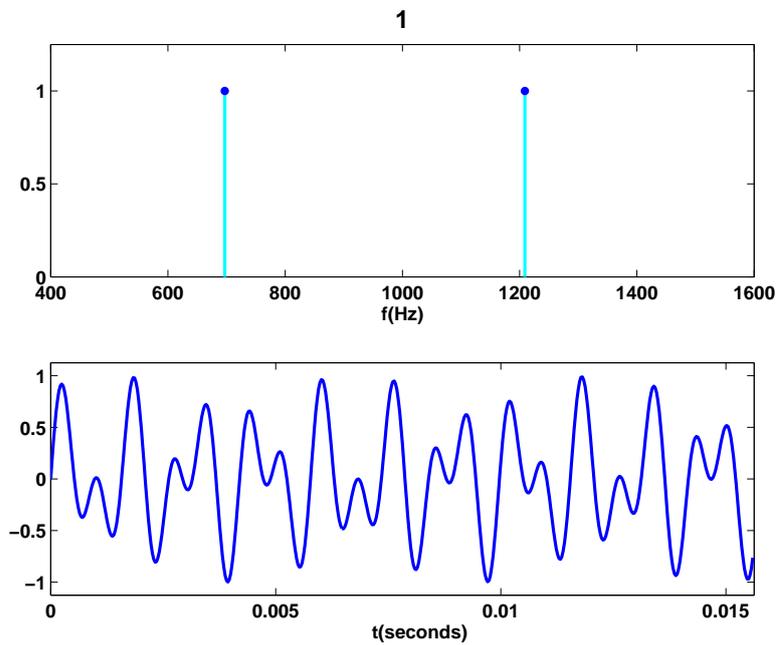


Figure 8.2. *The tone generated by the 1 button.*

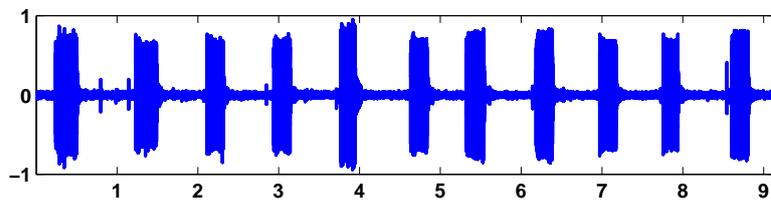


Figure 8.3. *Recording of an 11-digit telephone number.*

`y`

produces

```
y =
sig: [1x74800 int8]
fs: 8192
```

which shows that `y` has two fields, an integer vector `y.sig`, of length 74800, containing the signal, and a scalar `y.fs`, with the value 8192, which is the sample rate.

```
max(abs(y.sig))
```

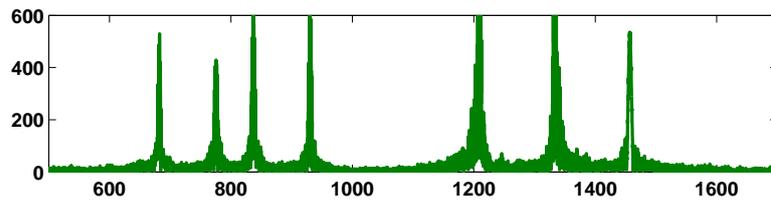


Figure 8.4. *FFT of the recorded signal.*

reveals that the elements of the signal are in the range $-127 \leq y_k \leq 127$. The statements

```
Fs = y.fs;
y = double(y.sig)/128;
```

save the sample rate and rescales the vector and converts it to double precision. The statements

```
n = length(y);
t = (0:n-1)/Fs
```

reproduce the sample times of the recording. The last component of `t` is 9.1307, indicating that the recording lasts a little over 9 s. Figure 8.3 is a plot of the entire signal.

This signal is noisy. You can even see small spikes on the graph at the times the buttons were clicked. It is easy to see that 11 digits were dialed, but, on this scale, it is impossible to determine the specific digits.

Figure 8.4 shows the magnitude of the FFT of the signal, which is the key to determining the individual digits.

The plot was produced with

```
p = abs(fft(y));
f = (0:n-1)*(Fs/n);
plot(f,p);
axis([500 1700 0 600])
```

The x -axis corresponds to frequency. The `axis` settings limit the display to the range of the DTMF frequencies. There are seven peaks, corresponding to the seven basic frequencies. This overall FFT shows that all seven frequencies are present someplace in the signal, but it does not help determine the individual digits.

The `touchtone` program also lets you break the signal into 11 equal segments and analyze each segment separately. Figure 8.5 is the display from the first segment.

For this segment, there are only two peaks, indicating that only two of the basic frequencies are present in this portion of the signal. These two frequencies come from the '1' button. You can also see that the waveform of a short portion of the first segment is similar to the waveform that our synthesizer produces for the

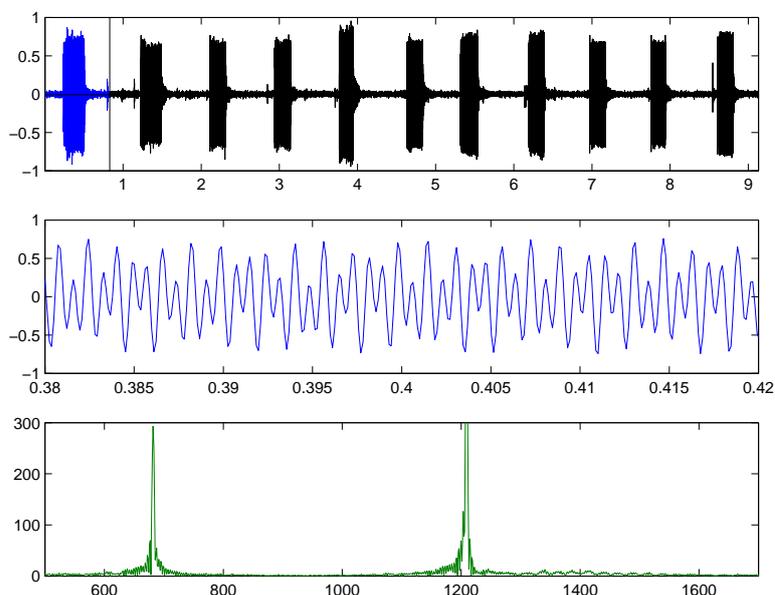


Figure 8.5. The first segment and its FFT.

'1' button. So we can conclude that the number being dialed in `touchtone` starts with a 1. Exercise 8.1 asks you to continue the analysis and identify the complete phone number.

8.2 Finite Fourier Transform

The finite, or discrete, Fourier transform of a complex vector y with n elements is another complex vector Y with n elements

$$Y_k = \sum_{j=0}^{n-1} \omega^{jk} y_j,$$

where ω is a complex n th root of unity:

$$\omega = e^{-2\pi i/n}.$$

In this chapter, the mathematical notation follows conventions common in signal processing literature where $i = \sqrt{-1}$ is the complex unit and j and k are indices that run from 0 to $n - 1$.

The Fourier transform can be expressed with matrix-vector notation:

$$Y = Fy,$$

where the Fourier matrix F has elements

$$f_{k,j} = \omega^{jk}.$$

It turns out that F is nearly its own inverse. More precisely, F^H , the complex conjugate transpose of F , satisfies

$$F^H F = nI,$$

so

$$F^{-1} = \frac{1}{n} F^H.$$

This allows us to invert the Fourier transform:

$$y = \frac{1}{n} F^H Y.$$

Hence

$$y_j = \frac{1}{n} \sum_{k=0}^{n-1} Y_k \bar{\omega}^{jk},$$

where $\bar{\omega}$ is the complex conjugate of ω :

$$\bar{\omega} = e^{2\pi i/n}.$$

We should point out that this is not the only notation for the finite Fourier transform in common use. The minus sign in the definition of ω after the first equation sometimes occurs instead in the definition of $\bar{\omega}$ used in the inverse transform. The $1/n$ scaling factor in the inverse transform is sometimes replaced with $1/\sqrt{n}$ scaling factors in both transforms.

In MATLAB, the Fourier matrix F can be generated for any given n by

```
omega = exp(-2*pi*i/n);
j = 0:n-1;
k = j';
F = omega.^(k*j)
```

The quantity $k*j$ is an *outer product*, an n -by- n matrix whose elements are the products of the elements of two vectors. However, the built-in function `fft` takes the finite Fourier transform of each column of a matrix argument, so an easier, and quicker, way to generate F is

```
F = fft(eye(n))
```

8.3 fftgui

The GUI `fftgui` allows you to investigate properties of the finite Fourier transform. If y is a vector containing a few dozen elements,

```
fftgui(y)
```

produces four plots.

```
real(y)      imag(y)
real(fft(y)) imag(fft(y))
```

You can use the mouse to move any of the points in any of the plots, and the points in the other plots respond.

Please run `fftgui` and try the following examples. Each illustrates some property of the Fourier transform. If you start with no arguments,

```
fftgui
```

all four plots are initialized to `zeros(1,32)`. Click your mouse in the upper left-hand corner of the upper left-hand plot. You are taking the `fft` of the first unit vector, with one in the first component and zeros elsewhere. This should produce Figure 8.6.

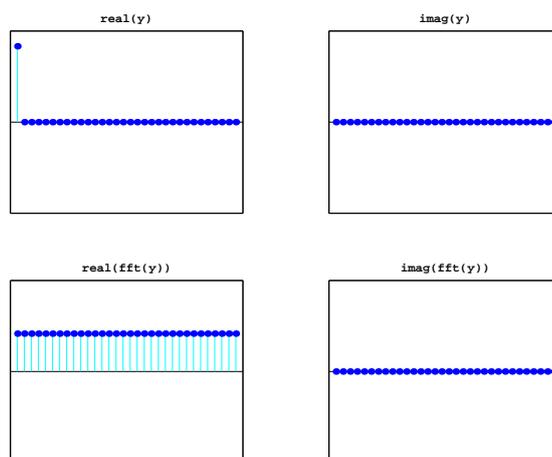


Figure 8.6. *FFT of the first unit vector is constant.*

The real part of the result is constant and the imaginary part is zero. You can also see this from the definition

$$Y_k = \sum_{j=0}^{n-1} y_j e^{-2ijk\pi/n}, \quad k = 0, \dots, n-1$$

if $y_0 = 1$ and $y_1 = \dots = y_{n-1} = 0$. The result is

$$Y_k = 1 \cdot e^0 + 0 + \dots + 0 = 1 \text{ for all } k.$$

Click y_0 again, hold the mouse down, and move the mouse vertically. The amplitude of the constant result varies accordingly.

Next try the second unit vector. Use the mouse to set $y_0 = 0$ and $y_1 = 1$. This should produce Figure 8.7. You are seeing the graph of

$$Y_k = 0 + 1 \cdot e^{-2ik\pi/n} + 0 + \dots + 0.$$

The n th root of unity can also be written

$$\omega = \cos \delta - i \sin \delta, \text{ where } \delta = 2\pi/n.$$

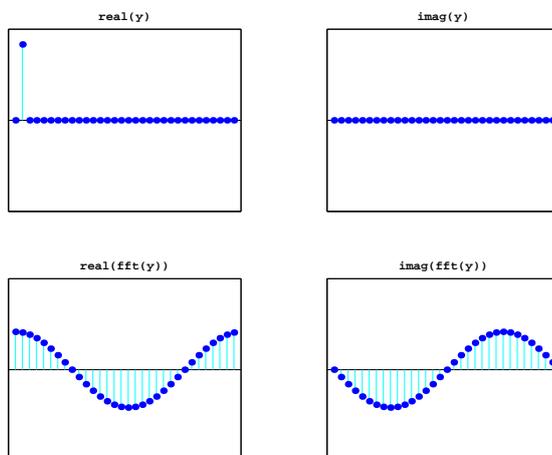


Figure 8.7. *FFT of the second unit vector is a pure sinusoid.*

Consequently, for $k = 0, \dots, n - 1$,

$$\text{real}(Y_k) = \cos k\delta, \quad \text{imag}(Y_k) = -\sin k\delta.$$

We have sampled two trig functions at n equally spaced points in the interval $0 \leq x < 2\pi$. The first sample point is $x = 0$ and the last sample point is $x = 2\pi - \delta$.

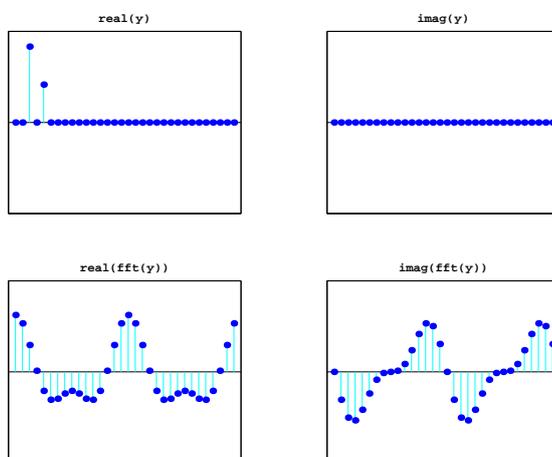


Figure 8.8. *FFT is the sum of two sinusoids.*

Now set $y_2 = 1$ and vary y_4 with the mouse. One snapshot is Figure 8.8. We have graphs of

$$\cos 2k\delta + \eta \cos 4k\delta \quad \text{and} \quad -\sin 2k\delta - \eta \sin 4k\delta$$

for various values of $\eta = y_4$.

The point just to the right of the midpoint of the x -axis is particularly important. It is known as the *Nyquist point*. With the points numbered from 0 to $n - 1$ for even n , it's the point with index $\frac{n}{2}$. If $n = 32$, it's point number 16. Figure 8.9 shows that the `fft` of a unit vector at the Nyquist point is a sequence of alternating $+1$'s and -1 's.

Now let's look at some symmetries in the FFT. Make several random clicks on the `real(y)` plot. Leave the `imag(y)` plot flat zero. Figure 8.10 shows an example. Look carefully at the two `fft` plots. Ignoring the first point in each plot, the real part is symmetric about the Nyquist point and the imaginary part is antisymmetric about the Nyquist point. More precisely, if y is any real vector of length n and $Y = \text{fft}(y)$, then

$$\begin{aligned} \text{real}(Y_0) &= \sum y_j, & \text{imag}(Y_0) &= 0, \\ \text{real}(Y_j) &= \text{real}(Y_{n-j}), & \text{imag}(Y_j) &= -\text{imag}(Y_{n-j}), \quad j = 1, \dots, n/2. \end{aligned}$$

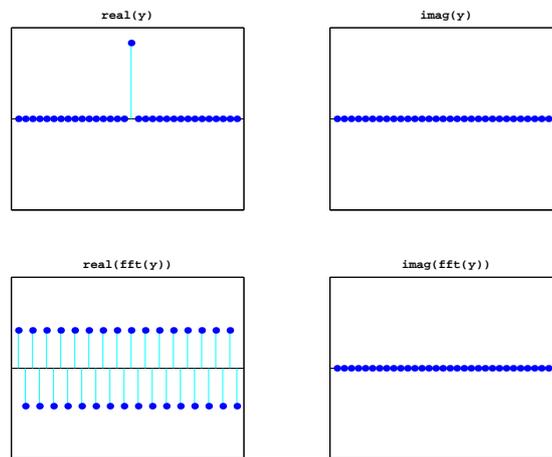


Figure 8.9. *FFT of the unit vector at the Nyquist point.*

8.4 Sunspots

For centuries, people have noted that the face of the sun is not constant or uniform in appearance, but that darker regions appear at random locations on a cyclical basis. This activity is correlated with weather and other economically significant terrestrial phenomena. In 1848, Rudolf Wolfer proposed a rule that combined the number and size of these sunspots into a single index. Using archival records, astronomers have applied Wolfer's rule to determine sunspot activity back to the year 1700. Today the sunspot index is measured by many astronomers, and the worldwide distribution of the data is coordinated by the Solar Influences Data Center at the Royal Observatory of Belgium [4].

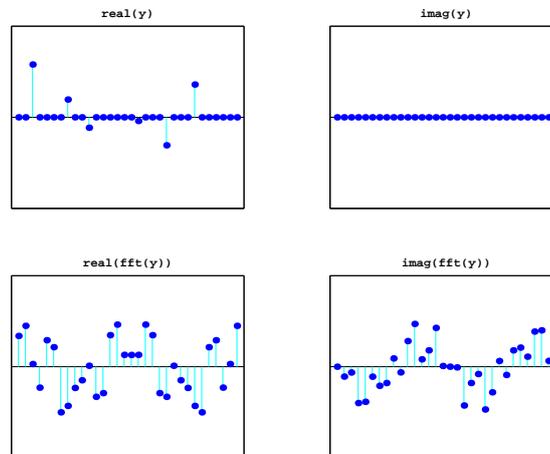


Figure 8.10. *Symmetry about the Nyquist point.*

The text file `sunspot.dat` in the MATLAB `demos` directory has two columns of numbers. The first column is the years from 1700 to 1987 and the second column is the average Wolfer sunspot number for each year.

```
load sunspot.dat
t = sunspot(:,1)';
wolfer = sunspot(:,2)';
n = length(wolfer);
```

There is a slight upward trend to the data. A least squares fit gives the trend line.

```
c = polyfit(t,wolfer,1);
trend = polyval(c,t);
plot(t,[wolfer; trend],'-',t,wolfer,'k.')
xlabel('year')
ylabel('Wolfer index')
title('Sunspot index with linear trend')
```

You can definitely see the cyclic nature of the phenomenon (Figure 8.11). The peaks and valleys are a little more than 10 years apart.

Now subtract off the linear trend and take the FFT.

```
y = wolfer - trend;
Y = fft(y);
```

The vector $|Y|^2$ is the *power* in the signal. A plot of power versus frequency is a *periodogram* (Figure 8.12). We prefer to plot $|Y|$, rather than $|Y|^2$, because the scaling is not so exaggerated. The sample rate for these data is one observation per year, so the frequency f has units of cycles per year.

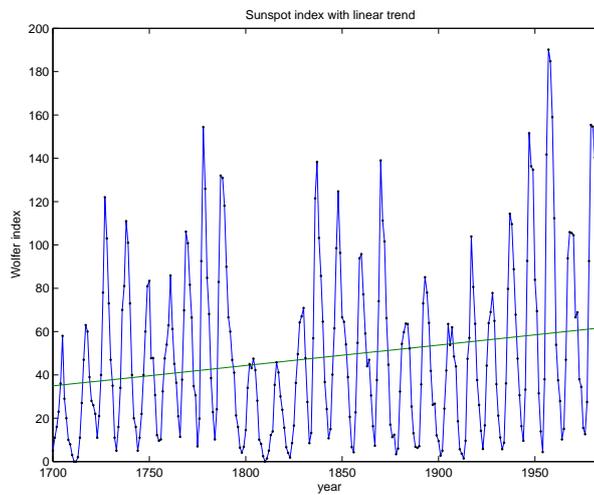


Figure 8.11. *Sunspot index.*

```

Fs = 1; % Sample rate
f = (0:n/2)*Fs/n;
pow = abs(Y(1:n/2+1));
pmax = 5000;
plot([f; f],[0*pow; pow],'c-', f,pow,'b.', ...
      'linewidth',2,'markersize',16)
axis([0 .5 0 pmax])
xlabel('cycles/year')
ylabel('power')
title('Periodogram')

```

The maximum power occurs near frequency = 0.09 cycles/year. We would like to know the corresponding period in years/cycle. Let's zoom in on the plot and use the reciprocal of frequency to label the x -axis.

```

k = 0:44;
f = k/n;
pow = pow(k+1);
plot([f; f],[0*pow; pow],'c-',f,pow,'b.', ...
      'linewidth',2,'markersize',16)
axis([0 max(f) 0 pmax])
k = 2:3:41;
f = k/n;
period = 1./f;
periods = sprintf('%5.1f|',period);
set(gca,'xtick',f)
set(gca,'xticklabel',periods)
xlabel('years/cycle')

```

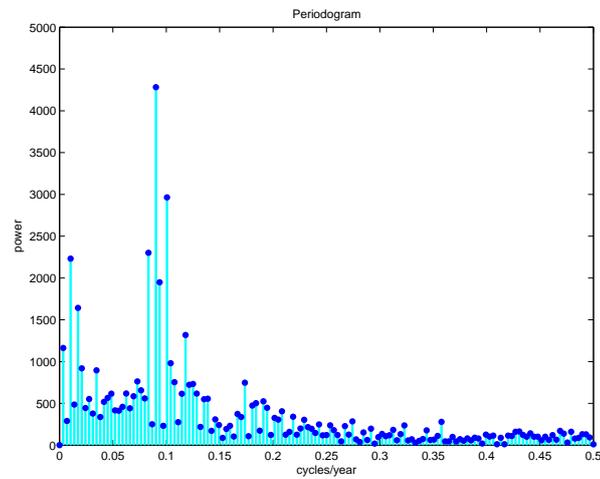


Figure 8.12. *Periodogram of the sunspot index.*

```
ylabel('power')
title('Periodogram detail')
```

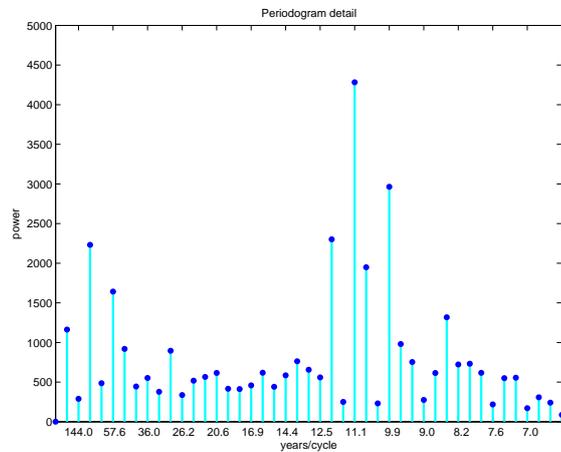


Figure 8.13. *Detail of periodogram shows 11-year cycle.*

As expected, there is a very prominent cycle with a length of about 11.1 years (Figure 8.13). This shows that, over the last 300 years, the period of the sunspot cycle has been slightly over 11 years.

8.5 Periodic Time Series

The tones generated by a touch-tone telephone and the Wolfer sunspot index are two examples of periodic time series, that is, functions of time that exhibit periodic behavior, at least approximately. Fourier analysis allows us to estimate the period from a discrete set of values sampled at a fixed rate. The following table shows the relationship between the various quantities involved in this analysis.

<code>y</code>	data
<code>F_s</code>	samples/unit-time
<code>n = length(y)</code>	number of samples
<code>t = (0:n-1)/F_s</code>	total time
<code>dt = 1/F_s</code>	time increment
<code>Y = fft(y)</code>	finite Fourier transform
<code>abs(Y)</code>	amplitude of FFT
<code>abs(Y).^2</code>	power
<code>f = (0:n-1)*(F_s/n)</code>	frequency, cycles/unit-time
<code>(n/2)*(F_s/n) = F_s/2</code>	Nyquist frequency
<code>p = 1./f</code>	period, unit-time/cycle

The periodogram is a plot of the FFT amplitude `abs(Y)`, or power `abs(Y).^2`, versus the frequency `f`. You only need to plot the first half because the second half is a reflection of the first half about the Nyquist frequency.

8.6 Fast Finite Fourier Transform

One-dimensional FFTs with a million points and two-dimensional 1000-by-1000 transforms are common. The key to modern signal and image processing is the ability to do these computations rapidly.

Direct application of the definition

$$Y_k = \sum_{j=0}^{n-1} \omega^{jk} y_j, \quad k = 0, \dots, n-1,$$

requires n multiplications and n additions for each of the n components of Y for a total of $2n^2$ floating-point operations. This does not include the generation of the powers of ω . A computer capable of doing one multiplication and addition every microsecond would require a million seconds, or about 11.5 days, to do a million-point FFT.

Several people discovered fast FFT algorithms independently and many people have since contributed to their development, but it was a 1965 paper by John Tukey of Princeton University and John Cooley of IBM Research that is generally credited as the starting point for the modern usage of the FFT.

Modern fast FFT algorithms have computational complexity $O(n \log_2 n)$ instead of $O(n^2)$. If n is a power of 2, a one-dimensional FFT of length n requires fewer than $3n \log_2 n$ floating-point operations. For $n = 2^{20}$, that's a factor of almost 35,000 faster than $2n^2$. Even if $n = 1024 = 2^{10}$, the factor is about 70.

With MATLAB 6.5 and a 700 MHz Pentium laptop, the time required for `fft(x)` if `length(x)` is $2^{20} = 1048576$ is about 1 s. The built-in `fft` function is based on FFTW, “The Fastest Fourier Transform in the West,” developed at M.I.T. by Matteo Frigo and Steven G. Johnson [2].

The key to the fast FFT algorithms is the fact that the square of the $2n$ th root of unity is the n th root of unity. Using complex notation,

$$\omega = \omega_n = e^{-2\pi i/n},$$

we have

$$\omega_{2n}^2 = \omega_n.$$

The derivation of the fast algorithm starts with the definition of the finite Fourier transform:

$$Y_k = \sum_{j=0}^{n-1} \omega^{jk} y_j, \quad k = 0, \dots, n-1.$$

Assume that n is even and that $k \leq n/2 - 1$. Divide the sum into terms with even subscripts and terms with odd subscripts.

$$\begin{aligned} Y_k &= \sum_{\text{even } j} \omega^{jk} y_j + \sum_{\text{odd } j} \omega^{jk} y_j \\ &= \sum_{j=0}^{n/2-1} \omega^{2jk} y_{2j} + \omega^k \sum_{j=0}^{n/2-1} \omega^{2jk} y_{2j+1}. \end{aligned}$$

The two sums on the right are components of the FFTs of length $n/2$ of the portions of y with even and odd subscripts. In order to get the entire FFT of length n , we have to do two FFTs of length $n/2$, multiply one of these by powers of ω , and concatenate the results.

The relationship between an FFT of length n and two FFTs of length $n/2$ can be expressed compactly in MATLAB. If `n = length(y)` is even,

```
omega = exp(-2*pi*i/n);
k = (0:n/2-1)';
w = omega .^ k;
u = fft(y(1:2:n-1));
v = w.*fft(y(2:2:n));
```

then

```
fft(y) = [u+v; u-v];
```

Now, if n is not only even but actually a power of 2, the process can be repeated. The FFT of length n is expressed in terms of two FFTs of length $n/2$, then four FFTs of length $n/4$, then eight FFTs of length $n/8$, and so on until we reach n FFTs of length one. An FFT of length one is just the number itself. If

$n = 2^p$, the number of steps in the recursion is p . There is $O(n)$ work at each step, so the total amount of work is

$$O(np) = O(n \log_2 n).$$

If n is not a power of two, it is still possible to express the FFT of length n in terms of several shorter FFTs. An FFT of length 100 is two FFTs of length 50, or four FFTs of length 25. An FFT of length 25 can be expressed in terms of five FFTs of length 5. If n is not a prime number, an FFT of length n can be expressed in terms of FFTs whose lengths divide n . Even if n is prime, it is possible to embed the FFT in another whose length can be factored. We do not go into the details of these algorithms here.

The `fft` function in older versions of MATLAB used fast algorithms if the length was a product of small primes. Beginning with MATLAB 6, the `fft` function uses fast algorithms even if the length is prime (see [2]).

8.7 ffttx

Our textbook function `fftttx` combines the two basic ideas of this chapter. If n is a power of 2, it uses the $O(n \log_2 n)$ fast algorithm. If n has an odd factor, it uses the fast recursion until it reaches an odd length, then sets up the discrete Fourier matrix and uses matrix-vector multiplication.

```
function y = ffttx(x)
%FFTTX Textbook Fast Finite Fourier Transform.
% FFTTX(X) computes the same finite Fourier transform
% as FFT(X). The code uses a recursive divide and
% conquer algorithm for even order and matrix-vector
% multiplication for odd order. If length(X) is m*p
% where m is odd and p is a power of 2, the computational
% complexity of this approach is O(m^2)*O(p*log2(p)).

x = x(:);
n = length(x);
omega = exp(-2*pi*i/n);

if rem(n,2) == 0
    % Recursive divide and conquer
    k = (0:n/2-1)';
    w = omega .^ k;
    u = ffttx(x(1:2:n-1));
    v = w.*fftttx(x(2:2:n));
    y = [u+v; u-v];
else
    % The Fourier matrix
    j = 0:n-1;
    k = j';
```

```
F = omega .^ (k*j);  
y = F*x;  
end
```

8.8 fftmatrix

The n -by- n matrix F generated by the MATLAB statement

```
F = fft(eye(n,n))
```

is a complex matrix whose elements are powers of the n th root of unity

$$\omega = e^{-2\pi i/n}$$

The statement

```
plot(fft(eye(n,n)))
```

connects the elements of each column of F and thereby generates a subgraph of the graph on n points. If n is prime, connecting the elements of all columns generates the complete graph on n points. If n is not prime, the sparsity of the graph of all columns is related to the speed of the FFT algorithm. The graphs for $n = 8, 9, 10,$ and 11 are shown in Figure 8.14.

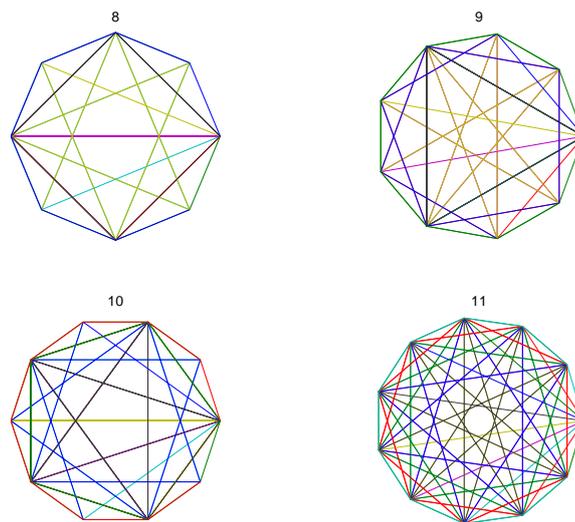


Figure 8.14. *Graphs of FFT matrix.*

Because $n = 11$ is prime, the corresponding graph shows all possible connections. But the other three values of n are not prime. Some of the links in their graphs are missing, indicating that the FFT of a vector with that many points can be computed more quickly.

The program `fftmatrix` allows you to investigate these graphs.

```
fftmatrix(n)
```

plots all the columns of the Fourier matrix of order n .

```
fftmatrix(n,j)
```

plots only the $j+1$ st column.

```
fftmatrix
```

defaults to `fftmatrix(10,4)`. In all cases, push buttons allow n , j and the choice between one or all columns be changed.

8.9 Other Fourier Transforms and Series

We have been studying the finite Fourier transform, which converts one finite sequence of coefficients into another sequence of the same length, n . The transform is

$$Y_k = \sum_{j=0}^{n-1} y_j e^{-2ij k \pi / n}, \quad k = 0, \dots, n-1.$$

The inverse transform is

$$y_j = \frac{1}{n} \sum_{k=0}^{n-1} Y_k e^{2ijk \pi / n}, \quad j = 0, \dots, n-1.$$

The Fourier integral transform converts one complex function into another. The transform is

$$F(\mu) = \int_{-\infty}^{\infty} f(t) e^{-2\pi i \mu t} dt.$$

The inverse transform is

$$f(t) = \int_{-\infty}^{\infty} F(\mu) e^{2\pi i \mu t} d\mu.$$

The variables t and μ run over the entire real line. If t has units of seconds, then μ has units of radians per second. Both of the functions $f(t)$ and $F(\mu)$ are complex valued, but in most applications the imaginary part of $f(t)$ is zero.

Alternative units use $\nu = 2\pi\mu$, which has units of cycles or revolutions per second. With this change of variable, there are no factors of 2π in the exponentials, but there are factors of $1/\sqrt{2\pi}$ in front of the integrals, or a single factor of $1/(2\pi)$ in the inverse transform. The MATLAB Symbolic Toolbox uses this alternative notation with the single factor in the inverse transform.

A Fourier series converts a periodic function into an infinite sequence of Fourier coefficients. Let $f(t)$ be the periodic function and let L be its period, so

$$f(t+L) = f(t) \text{ for all } t.$$

The Fourier coefficients are given by integrals over the period

$$c_j = \frac{1}{L} \int_{-L/2}^{L/2} f(t) e^{-2\pi i j t} dt, \quad j = \dots, -1, 0, 1, \dots$$

With these coefficients, the complex form of the Fourier series is

$$f(t) = \sum_{j=-\infty}^{\infty} c_j e^{2\pi i j t / L}.$$

A discrete time Fourier transform converts an infinite sequence of data values into a periodic function. Let x_k be the sequence, with the index k taking on all integer values, positive and negative.

The discrete time Fourier transform is the complex-valued periodic function

$$X(e^{i\omega}) = \sum_{k=-\infty}^{\infty} x_k e^{ik\omega}.$$

The sequence can then be represented as

$$x_k = \frac{1}{2\pi} \int_{-\pi}^{\pi} X(e^{i\omega}) e^{-ik\omega} d\omega, \quad k = \dots, -1, 0, 1, \dots$$

The Fourier integral transform involves only integrals. The finite Fourier transform involves only finite sums of coefficients. Fourier series and the discrete time Fourier transform involve both integrals and sequences. It is possible to “morph” any of the transforms into any of the others by taking limits or restricting domains.

Start with a Fourier series. Let L , the length of the period, become infinite and let j/L , the coefficient index scaled by the period length, become a continuous variable, μ . Then the Fourier coefficients c_j become the Fourier transform $F(\mu)$.

Again, start with a Fourier series. Interchanging the roles of the periodic function and the infinite sequence of coefficients leads to the discrete time Fourier transform.

Start with a Fourier series a third time. Now restrict t to a finite number of integral values, k , and restrict j to the same finite number of values. Then the Fourier coefficients become the finite Fourier transform.

In the Fourier integral transform context, Parseval’s theorem says

$$\int_{-\infty}^{+\infty} |f(t)|^2 dt = \int_{-\infty}^{+\infty} |F(\mu)|^2 d\mu.$$

This quantity is known as the *total power* in a signal.

8.10 Further Reading

Van Loan [5] describes the computational framework for the fast transforms. A page of links at the FFTW Web site [3] provides useful information.

Exercises

- 8.1. What is the telephone number recorded in `touchtone.mat` and analyzed by `touchtone.m`?

- 8.2. Modify `touchtone.m` so that it can dial a telephone number specified by an input argument, such as `touchtone('1-800-555-1212')`.
- 8.3. Our version of `touchtone.m` breaks the recording into a fixed number of equally spaced segments, each corresponding to a single digit. Modify `touchtone` so that it automatically determines the number and the possibly disparate lengths of the segments.
- 8.4. Investigate the use of the MATLAB functions `audiorecorder` and `audioplayer`, or some other system for making digital recordings. Make a recording of a phone number and analyze it with your modified version of `touchtone.m`.
- 8.5. Recall that the Fourier matrix F is the n -by- n complex matrix with elements

$$f_{k,j} = \omega^{jk},$$

where

$$\omega = e^{-2\pi i/n}.$$

Show that $\frac{1}{\sqrt{n}}F$ is unitary. In other words, show that F^H , the complex conjugate transpose of F , satisfies

$$F^H F = nI.$$

The notation here is a little unusual for matrices because the subscripts j and k run from 0 to $n - 1$, instead of from 1 to n .

- 8.6. What relationship between `n` and `j` causes `fftmatrix(n,j)` to produce a five-point star? What relationship produces a regular pentagon?
- 8.7. *El Niño*. The climatological phenomenon *el Niño* results from changes in atmospheric pressure in the southern Pacific ocean. The ‘Southern Oscillation Index’ is the difference in atmospheric pressure between Easter Island and Darwin, Australia, measured at sea level at the same moment. The text file `elnino.dat` contains values of this index measured on a monthly basis over the 14-year period 1962 through 1975. Your assignment is to carry out an analysis similar to the sunspot example on the *el Niño* data. The unit of time is one month instead of one year. You should find there is a prominent cycle with a period of 12 months and a second, less prominent, cycle with a longer period. This second cycle shows up in about three of the Fourier coefficients, so it is hard to measure its length, but see if you can make an estimate.
- 8.8. Train whistle. The MATLAB `demos` directory contains several sound samples. One of them is a train whistle. The statement

```
load train
```

gives you a long vector `y` and a scalar `Fs` whose value is the number of samples per second. The time increment is $1/Fs$ seconds.

If your computer has sound capabilities, the statement

```
sound(y,Fs)
```

plays the signal, but you don't need that for this problem.

The data do not have a significant linear trend. There are two pulses of the whistle, but the harmonic content of both pulses is the same.

(a) Plot the data with time in seconds as the independent variable.

(b) Produce a periodogram with frequency in cycles/second as the independent variable.

(c) Identify the frequencies of the six peaks in the periodogram. You should find that ratios between these six frequencies are close to ratios between small integers. For example, one of the frequencies is $5/3$ times another. The frequencies that are integer multiples of other frequencies are *overtones*. How many of the peaks are fundamental frequencies and how many are overtones?

- 8.9. Bird chirps. Analyze the `chirp` sound sample from the MATLAB `demod` directory. By ignoring a short portion at the end, it is possible to segment the signal into eight pieces of equal length, each containing one chirp. Plot the magnitude of the FFT of each segment. Use `subplot(4,2,k)` for $k = 1:8$ and the same axis scaling for all subplots. Frequencies in the range from roughly 400 Hz to 800 Hz are appropriate. You should notice that one or two of the chirps have distinctive plots. If you listen carefully, you should be able to hear the different sounds.

Bibliography

- [1] J. W. COOLEY AND J. W. TUKEY, *An algorithm for the machine calculation of complex Fourier series*, Mathematics of Computation, 19 (1965), pp. 297–301.
- [2] M. FRIGO AND S. G. JOHNSON, *FFTW: An adaptive software architecture for the FFT*, in Proceedings of the 1998 IEEE International Conference on Acoustics Speech and Signal Processing, 3 (1998), pp. 1381–1384.
<http://www.fftw.org>
- [3] M. FRIGO AND S. G. JOHNSON, *Links to FFT-related resources*.
<http://www.fftw.org/links.html>
- [4] SOLAR INFLUENCES DATA CENTER, *Sunspot archive and graphics*.
<http://sidc.oma.be>
- [5] C. VAN LOAN, *Computational Frameworks for the Fast Fourier Transform*, SIAM, Philadelphia, 1992.