

AMAT 415: Basic Signal Processing

Michael P. Lamoureux

Copyright 2008. Revision 17/04/08

Contents

1	Basic Complex Analysis	1
1.1	Complex numbers	1
1.1.1	Arithmetic	1
1.1.2	Geometry	2
1.1.3	Algebraic results	4
1.2	Elementary functions	5
1.2.1	The exponential function	7
1.2.2	The logarithm and complex powers	9
1.3	Analytic functions: doing calculus with complex numbers	10
1.4	The Cauchy-Riemann equations	12
1.5	Analytic functions	13
1.6	Power series	14
1.6.1	Standard series	15
1.7	Complex integrals	16
1.8	Some easy integrals	18
1.9	The Cauchy Theorems	19
1.10	Proofs of the Cauchy theorems	23
1.11	Applications of the Cauchy Integral Formula	23
2	Signals and systems	25
2.1	Sampling	25
2.2	Aliasing	26
2.3	Frequency aliasing	26
2.4	Sampled signals as a vector space	27
3	Convolving signals	29
3.1	Z transforms	29
3.2	Convolution	29
3.3	Convolution as matrix-vector multiplication	30
3.4	Convolution as matrix-matrix multiplication	30
3.5	Convolution by flipping and shifting	31
3.6	Convolution as a system	31
3.7	The shift operator	32
3.8	Signals as sums of δ^n 's.	33
4	Linear, shift-invariant systems	35
4.1	Impulse response of a LSI system, Z transform	36

5	Building practical LSI systems	37
5.1	FIR systems, minimum phase, maximum phase	37
5.2	Recursive (IIR) systems	38
5.3	Rational functions and IIR systems	38
5.4	Zeros and poles	40
5.5	Stability of recursive LSI systems	40
5.6	Example: a smoothing system. FIR and IIR	40
6	Frequency response of a LSI system	43
6.1	Frequency response of a simple smoother	44
6.2	A better three point smoother	45
6.3	Linear algebra notes: Eigenvalues and eigenvectors	46
6.4	A rant on radians	46
6.5	Symmetries of frequency response $H(\omega)$ in natural units.	47
6.6	Magnitude and phase response	50
7	Interlude: Some practical examples	53
7.1	MacRecorder	53
7.2	A 60Hz noise filter	54
8	Simple Filters	55
8.1	A single zero	55
8.2	A single pole	57
8.3	All-pass: matched zero/pole filters	58
8.4	Stability	58
9	Fourier transforms I	61
9.1	Defining the FT on sampled signals and on 1-periodic functions	61
9.2	Examples of FTs, on sequences	61
9.3	Examples of FTs, on 1-periodic functions	63
9.4	Properties of the Fourier transform	63
9.5	Some worked examples	64
9.6	Convolution, etc, on 1-periodic functions	67
9.7	Verifying the six properties	68
9.8	What does the FT mean?	69
9.9	How do we apply the FT?	70
10	Fourier transforms II	73
11	Fourier transforms III	75
11.1	The discrete FT as a matrix	75
11.2	What does the DFT mean?	76
11.3	The six FT properties, for the discrete FT	76
11.3.1	The normalized DFT	77
11.4	MATLAB quirks	78
11.5	Circular convolution	79
11.6	Zero padding	80
11.7	The Fast Fourier Transform	80

<i>CONTENTS</i>	iii
12 General Fourier transforms	83
13 Symmetries of the Fourier transform	85
14 The uncertainty principle	87
14.1 Boring math calculations: FT of Gaussians and boxcars	88
15 Calculus and the Fourier transform	91
15.1 More math: derivatives	92
16 Special Filters	93
17 FIR, IIR filters	97
18 Laplace transforms	99
19 Wavelet transforms	101
19.1 Dilations and translations of a wavelet	102
19.2 Analysis: wavelet coefficients	103
19.3 Synthesis: reconstructing the signal	103
19.4 Fast transforms	104
19.5 Multiresolution analysis	105
19.6 Application: data compression	105
APPENDIX 1: Collected wisdom	107
APPENDIX 2: Eigenvalues and eigenvectors	109
APPENDIX 3: Sums and integrals of exponentials	111
APPENDIX 4: Music and frequencies	113

List of Figures

1.1	Two complex numbers on the complex plane.	3
1.2	The complex number $z = 4 + 7i$ and its 4-th roots z_0, z_1, z_2, z_3 . Note the four roots are uniformly spread around a circle of radius $r = \sqrt[4]{ z }$. The polar angle of z_0 is one-quarter the polar angle of z	5
1.3	Two strips in complex plane, mapped under the exponential, to curving strips. . . .	8
1.4	A curve in the complex plane, connecting A to B . Integrating by breaking the curve up into short segments, and summing the segments.	17
1.5	A circle of radius R , which we will integrate around. The branch cut indicates where the log function fails to be analytic.	18
1.6	Some curves. a) is a simple oriented curve, which is not closed. b) is a simple closed curve, oriented counterclockwise. c) is a closed curve, but not simple as it intersects itself. d) is a closed curve which wraps around the point x twice. e) is also a simple closed curve, although it is pretty wiggly.	20
1.7	Curves A and B can be continuously deformed one into the other, without leaving the shaded region. They are homotopic. However, curve C cannot be deformed into A or B , because it doesn't go around the hole.	21
5.1	FIR smoother: A boxcar input \mathbf{x} and its smoothed output \mathbf{y}	41
5.2	IIR smoother: A boxcar input \mathbf{x} and its smoothed output \mathbf{y}	42
6.1	Plot of the frequency response $H(\omega)$ for a three point smoother.	45
6.2	Plot of the frequency response $H(\omega)$ for a better three point smoother.	46
6.3	Sine waves at 1 rad per second (left) and 10 rad per second (right). These are not natural units!	47
6.4	Sine waves at 1 cycle per second (left) and 10 cycles per second (right). This is natural.	47
6.5	1-periodicity of the frequency response curve $H(\omega)$ for a 3 point smoother.	48
6.6	The frequency response curve of a 3 point smoother. Three standard plots.	49
6.7	Input/output response with the amplitude and phase effect.	51
8.1	Amplitude response for a filter with zero at $Z = 1$ ($\omega = 0$) and at $Z = i$ ($\omega = -.25$)	56
8.2	Amplitude response for a filter with zero at $Z = 1$ and at $Z = 2$	56
8.3	Phase response for a filter with zero at $Z = 2$ (min phase), and at $Z = 1/2$ (max phase).	56
8.4	Amplitude response for a filter with pole at $Z = 1$ ($\omega = 0$) and at $Z = -i$ ($\omega = .25$)	57
8.5	Amplitude response for a filter with pole at $Z = 1$ and at $Z = 2$	58
8.6	Phase response for a filter with pole at $Z = 2$ (min phase), and at $Z = 1/2$ (max phase).	58
8.7	Amplitude and phase response for a filter with matched zero and pole pair.	59

9.1	The Fourier transform of a long stretch of ones. ($N = 10$)	62
9.2	The Fourier transform of a boxcar, supported on interval $[-1/3, 1/3]$. The FT is a time series sequence, that is a set of uniform samples of the sinc function.	64
9.3	Plot of signal made up of sine waves, and the FT showing the frequency content.	70
9.4	Plot of spiky signal, and the FT showing the frequency content.	71
9.5	Smoothing out a noisy signal by zeroing out the high frequencies.	71
13.1	A real signal made up of sine waves, and the FT showing symmetry in frequency content.	85
14.1	A wide Gaussian signal, and its Fourier transform, which is narrow.	87
14.2	A narrow Gaussian signal, and its Fourier transform, which is wide.	88
14.3	A wide boxcar signal, and its Fourier transform, which is narrow.	89
14.4	A narrow boxcar signal, and its Fourier transform, which is wide.	90
16.1	A signal, filtered with a low-pass brick wall, to get the smoothed output.	93
16.2	A wave being phase rotated through one complete cycle.	95
19.1	Four different examples of wavelets, or small waves.	101
19.2	A wavelet $\phi(t)$ centered at zero, its translate $\phi(t - 3)$, and the translate/dilate $\frac{1}{\sqrt{2}}\phi(\frac{t-7}{2})$. Note the dilated one is twice as wide, and 70% of the height of the others.	102
19.3	A plot of the wavelet coefficients for a chirp signal. See how it clearly indicates a signal which starts on the left at a low frequency (high scale a) and ends up at the right with a high frequency (low scale a).	104

Chapter 1

Basic Complex Analysis

1.1 Complex numbers

Complex numbers are numbers of the form $x + yi$, where x and y are real numbers, and i is the square root of minus one. For instance, $2 + 3i$ is a complex number, as is $\sqrt{2} + \pi i$. We can add, subtract, multiply, and divide complex numbers just like regular real numbers, with the usual laws of algebra holding (commutative, associative, distributive laws, etc). We usually write $z = x + yi$ for a generic complex number, where x is the real part of the number, and y is the imaginary part of the number. It is convenient to define the basic functions

$$\Re(z) = x \quad \text{the real part of } z \quad (1.1)$$

$$\Im(z) = y \quad \text{the imaginary part of } z \quad (1.2)$$

$$|z| = \sqrt{x^2 + y^2} \quad \text{the absolute value (or modulus) of } z \quad (1.3)$$

$$\bar{z} = x - yi \quad \text{the conjugate of } z \quad (1.4)$$

and observe some simple relations like $z\bar{z} = |z|^2$, $|\bar{z}| = |z|$, and $-|z| \leq \Re(z) \leq |z|$.

1.1.1 Arithmetic

Addition and subtraction are done by adding the real and imaginary parts, while multiplication is done by usual distribution and using the fact that $i^2 = -1$. Thus we can compute

$$(2 + 3i) + (4 + 5i) = 5 + 8i \quad (1.5)$$

$$(2 + 3i) - (4 + 5i) = (2 - 4) + (3 - 5)i = -2 - 2i \quad (1.6)$$

$$(2 + 3i) * (4 + 5i) = 2 * 4 + 2 * 5i + 3i * 4 + 3i * 5i \quad (1.7)$$

$$= 8 + 10i + 12i + 15i^2 = (8 - 15) + (10 + 12)i = -7 + 22i \quad (1.8)$$

Division of a complex number by a real number is easy, just do it by components,

$$(4 + 5i)/2 = 2 + 2.5i. \quad (1.9)$$

Division by a complex number is only slightly more complicated, just use the formula

$$\frac{z_1}{z_2} = \frac{z_1 \bar{z}_2}{|z_2|^2} \quad (1.10)$$

so, for instance

$$\frac{2+3i}{4+5i} = \frac{(2+3i)(4-5i)}{|4+5i|^2} = \frac{23+2i}{16+25} = \frac{23}{41} + \frac{2}{41}i. \quad (1.11)$$

If this seems weird to you, remember that the point of division is to be the reverse of multiplication. So we can check the answer, by multiplying

$$\left(\frac{23}{41} + \frac{2}{41}i\right)(4+5i) = (4*23-5*2)/41 + ((23*5+2*4)/41)i = (82/41) + (123/41)i = 2+3i \quad (1.12)$$

which is the original numerator.

The point of all this is that with these definitions, the set of all complex numbers forms a field, which means we can do all the usual arithmetic tricks with this set of numbers. The only thing we don't have is an order, so it doesn't make sense to ask whether $2+3i < 4+5i$.

You should also note how the arithmetic operations combine with conjugation. For instance, it is easy to verify that

$$\overline{z_1 + z_2} = \overline{z_1} + \overline{z_2} \quad (1.13)$$

$$\overline{z_1 - z_2} = \overline{z_1} - \overline{z_2} \quad (1.14)$$

$$\overline{(z_1 z_2)} = \overline{z_1} \overline{z_2} \quad (1.15)$$

$$\overline{(z_1/z_2)} = \overline{z_1}/\overline{z_2} \quad (1.16)$$

$$\overline{\overline{z}} = z \quad (1.17)$$

$$\overline{z} = z \quad \text{if and only if } z \text{ is real} \quad (1.18)$$

$$\Re(z) = \frac{z + \overline{z}}{2} \quad (1.19)$$

$$\Im(z) = \frac{z - \overline{z}}{2i} \quad (\text{note the } i \text{ on the bottom}) \quad (1.20)$$

1.1.2 Geometry

It is useful to draw complex numbers as points on a plane, with the x and y parts the usual horizontal and vertical components. See Figure 1. Thus complex numbers “look” just like vectors in the plane. Addition and subtraction behave just like vector addition and subtraction. The absolute value of a complex number is just its usual length as a vector. We also see immediately that the triangle rule for addition holds,

$$|z_1 + z_2| \leq |z_1| + |z_2|, \quad (1.21)$$

since it is true for vectors.

From this graphical representation, it is clear that we write any complex number in polar form, with

$$z = r(\cos(\theta) + i \sin(\theta)) \quad (1.22)$$

where the radial distance $r = |z|$ is just the absolute value of z , and $\phi = \arg(z)$ is the angle measured from the x axis to the radial vector pointing to z . The sum $\cos(\theta) + i \sin(\theta)$ comes up so often that we give it a special notation, using the exponential form

$$e^{i\theta} = \cos(\theta) + i \sin(\theta). \quad (1.23)$$




Figure 1.1: Two complex numbers on the complex plane.

Note that the usual property of exponentials (exp of a sum is the product of the exp's) summarizes the familiar trig formulas for the sum of angles. That is, we have three equivalent formulas:

$$e^{i(\theta_1 + \theta_2)} = e^{i\theta_1} e^{i\theta_2} \quad (1.24)$$

$$\cos(\theta_1 + \theta_2) + i \sin(\theta_1 + \theta_2) = [\cos(\theta_1) + i \sin(\theta_1)][\cos(\theta_2) + i \sin(\theta_2)] \quad (1.25)$$

$$\cos(\theta_1 + \theta_2) + i \sin(\theta_1 + \theta_2) = [\cos(\theta_1) \cos(\theta_2) - \sin(\theta_1) \sin(\theta_2)] + \quad (1.26)$$

$$i[\sin(\theta_1) \cos(\theta_2) + \cos(\theta_1) \sin(\theta_2)] \quad (1.27)$$

What is amazing is now you don't have to remember those complicated sine, cosine laws that you learned in trig: they can all be derived from these simple exponential rules.

The form $z = re^{i\theta}$ is called the polar representation of the complex number. Notice we can do multiplication in this form, so

$$z_1 z_2 = (r_1 e^{i\theta_1})(r_2 e^{i\theta_2}) = (r_1 r_2)(e^{i\theta_1} e^{i\theta_2}) = (r_1 r_2) e^{i(\theta_1 + \theta_2)}. \quad (1.28)$$

Thus we see that when two complex numbers are multiplied, their lengths (abs. value) multiply, and their angles (arg) just add. We can sum up several related properties as follows:

$$|z_1 * z_2| = |z_1| * |z_2| \quad (1.29)$$

$$|z_1 / z_2| = |z_1| / |z_2| \quad (1.30)$$

$$\arg(z_1 * z_2) = \arg(z_1) + \arg(z_2) \quad \text{mod } 2\pi \quad (1.31)$$

$$\arg(z_1 / z_2) = \arg(z_1) - \arg(z_2) \quad \text{mod } 2\pi \quad (1.32)$$

$$\arg(\bar{z}) = -\arg(z) \quad \text{mod } 2\pi \quad (1.33)$$

These comments “mod 2π ” point out one difficulty that the angle from the x axis is not uniquely defined, since we can wrap around the origin many times as we work out the angle. There are several conventions used when defining the argument function. One standard convention is to choose the angle measurement so that we always have

$$0 \leq \arg(z) < 2\pi; \quad (1.34)$$

another convention is to choose

$$-\pi < \arg(z) \leq \pi. \quad (1.35)$$

Mathematicians usually use the first convention. Computer scientists, physical scientists, and others, often use the second convention. MATLAB uses the second convention. The problems with calculating modulo 2π remains in any convention you choose.

From the geometric picture, it is tempting to define the arg function as

$$\arg(z) = \arctan(y/x), \quad (1.36)$$

however, this is only valid on the right half of the complex plane, where $x > 0$. So be careful.

To summarize, the graphical representation of complex numbers on the two dimensional plane tells us a lot about how complex numbers behave under calculations. Addition and subtraction is just like with vectors. Multiplication of two complex numbers just multiplies their vector lengths, and adds the polar angles. You can think about arithmetic operations on complex numbers as simple geometric movements.

1.1.3 Algebraic results

Since we know how multiplication works, we have a simple formula for computing powers of a complex number. This is called **de Moivre's formula**. With $z = re^{i\theta}$, then

$$z^n = r^n e^{in\theta} \quad \text{for any integer } n. \quad (1.37)$$

From this, we can also compute the n -th roots of any complex number. Again, with $z = re^{i\theta}$, the n -th roots of z are given as

$$z_k = \sqrt[n]{r} e^{i(\theta/n + 2\pi k/n)}, \quad k = 0, 1, 2, \dots, n-1. \quad (1.38)$$

The reason this works of course follows from de Moivre's formula, since we see that

$$z_k^n = (\sqrt[n]{r})^n e^{in(\theta/n + 2\pi k/n)} = r e^{i\theta + 2\pi k} = r e^{i\theta} = z. \quad (1.39)$$

Figure 2 gives a nice geometric interpretation of this root formula, for the case of finding fourth roots. Notice the four roots are uniformly spread around a circle of fixed radius. This happens in general case of finding n -th roots of a complex number.

The numbers $z_k = e^{i2\pi k/n}$, $k = 0, 1, 2, \dots, n-1$ are called the n -th roots of unity, and are the n distinct solutions to the equation

$$z^n = 1. \quad (1.40)$$

As in the previous example, these n complex roots of unity are uniformly spread around the circle of radius one in the complex plane, and include the trivial real root $z_0 = 1$. We will see these roots of unity many times in the course, in particular with the discrete Fourier transform.

A deep result in algebra, which we can prove using complex analysis, is that EVERY polynomial equation of the form

$$z^n + a_1 z^{n-1} + a_2 z^{n-2} + \dots + a_{n-1} z + a_n = 0 \quad (1.41)$$

can be solved, finding roots z_1, z_2, \dots, z_n to the equation. Thus the polynomial can always be factored in the form

$$z^n + a_1 z^{n-1} + a_2 z^{n-2} + \dots + a_{n-1} z + a_n = (z - z_1)(z - z_2)(z - z_2) \dots (z - z_n). \quad (1.42)$$

This result is called the **Fundamental Theorem of Algebra** and holds whether the coefficients $a_1 \dots a_n$ are real or complex. The quadratic formula we learned in high school shows how to solve this in the simple case of $n = 2$. There are also formulas for $n = 3, 4$, but for higher degree polynomial, the roots may be found using numerical methods.




Figure 1.2: The complex number $z = 4 + 7i$ and its 4-th roots z_0, z_1, z_2, z_3 . Note the four roots are uniformly spread around a circle of radius $r = \sqrt[4]{|z|}$. The polar angle of z_0 is one-quarter the polar angle of z .

1.2 Elementary functions

We are interested in defining functions that take any complex number z and compute a new complex function. For instance,

$$f(z) = z^3 + 23z^2 + (2 + 3i)z + (4 + 5i) \quad (1.43)$$

is a simple example of a *polynomial function* that maps complex numbers to complex numbers. A *rational function* is the quotient of two polynomials, such as the function

$$f(z) = \frac{z^3 + 23z^2 + (2 + 3i)z + (4 + 5i)}{z^2 + 1}. \quad (1.44)$$

This last function is undefined at the points $z = \pm i$, since we can't divide by zero. But we still consider it a perfectly useful complex valued function, with a domain that include all but two complex numbers.

The complex exponential function is defined on complex number $z = x + iy$ as

$$e^z = e^x e^{iy} = e^x (\cos(y) + i \sin(y)). \quad (1.45)$$

From this, we may define the usual trig functions as

$$\cos(z) = \frac{e^{iz} + e^{-iz}}{2} \quad (1.46)$$

$$\sin(z) = \frac{e^{iz} - e^{-iz}}{2i} \quad (1.47)$$

$$\tan(z) = \frac{\sin(z)}{\cos(z)} = \frac{1}{i} \frac{e^{iz} - e^{-iz}}{e^{iz} + e^{-iz}} \quad (1.48)$$

$$\sec(z) = \frac{1}{\cos(z)} = \frac{2}{e^{iz} + e^{-iz}} \quad (1.49)$$

$$\csc(z) = \frac{1}{\sin(z)} = \frac{2i}{e^{iz} - e^{-iz}} \quad (1.50)$$

$$\cot(z) = \frac{\cos(z)}{\sin(z)} = \frac{i}{1} \frac{e^{iz} + e^{-iz}}{e^{iz} - e^{-iz}}. \quad (1.51)$$

The hyperbolic trig function are defined similarly, with

$$\cosh(z) = \frac{e^z + e^{-z}}{2} \quad (1.52)$$

$$\sinh(z) = \frac{e^z - e^{-z}}{2} \quad (1.53)$$

$$\tanh(z) = \frac{\sinh(z)}{\cosh(z)} = \frac{e^z - e^{-z}}{e^z + e^{-z}} \quad (1.54)$$

$$\operatorname{sech}(z) = \frac{1}{\cosh(z)} = \frac{2}{e^z + e^{-z}} \quad (1.55)$$

$$\operatorname{csch}(z) = \frac{1}{\sinh(z)} = \frac{2}{e^z - e^{-z}} \quad (1.56)$$

$$\operatorname{coth}(z) = \frac{\cosh(z)}{\sinh(z)} = \frac{e^z + e^{-z}}{e^z - e^{-z}}. \quad (1.57)$$

What we see is that all the trig and hyperbolic function are defined in terms of the exponential function. So a careful examination of this one function is in order. We will do that in the next section.

It is worth checking that the trig identities still hold. For instance, we can check that

$$\sin(z_1 + z_2) = \sin(z_1) \cos(z_2) + \cos(z_1) \sin(z_2). \quad (1.58)$$

Do this as an exercise, using the definition of sine and cosine above, and the exponential laws.

It is also interesting to note that the trig functions are defined almost exactly the same as the hyperbolic functions, except for the careful placement of a factor of i . Thus, we see that on pure imaginary numbers of the form iy , we have the interesting identities

$$\sin(iy) = i * \sinh(y) \quad (1.59)$$

$$\cos(iy) = \cosh(y) \quad (1.60)$$

$$\tan(iy) = i * \tanh(y) \quad (1.61)$$

$$\sec(iy) = \operatorname{sech}(y) \quad (1.62)$$

$$\csc(iy) = -i * \operatorname{csch}(y) \quad (1.63)$$

$$\cot(iy) = -i * \operatorname{coth}(y) \quad (1.64)$$

$$(1.65)$$

Thus, the trig functions, applied to pure imaginary numbers, evaluate to the usual (real) hyperbolic functions, with the possible inclusion of an imaginary factor. (You may have noticed the even functions like \cos, \sec don't pick up any imaginary factor, while the odd functions do.)

It turns out we can use trig identities to compute our complex functions in terms of the more familiar real functions. For instance, to compute $\sin(z)$ for complex $z = x + iy$, we write

$$\sin(z) = \sin(x + iy) = \sin(x) \cos(iy) + \cos(x) \sin(iy) = \sin(x) \cosh(y) + i \cos(x) \sinh(y). \quad (1.66)$$

That is, $\sin(z)$ is simply defined in terms of products of real sine and cosine functions with real hyperbolic sine and cosine functions.

This sounds confusing. It's not. The two key facts are 1) everything is defined by the exponential, and 2) all the trig identities you used to know still apply to the complex valued functions. So for instance, it is true that

$$\tan^2(z) + 1 = \sec^2(z) \quad (1.67)$$

holds for all complex numbers z .

1.2.1 The exponential function

The exponential of a complex number $z = x + iy$ is simply defined as

$$e^z = e^x(\cos(y) + i\sin(y)). \quad (1.68)$$

The motivation for this comes from the Taylor series expansion for e^x , which we remember from first year calculus to be

$$e^x = 1 + \frac{x}{1!} + \frac{x^2}{2!} + \frac{x^3}{3!} + \frac{x^4}{4!} + \frac{x^5}{5!} + \cdots \quad (1.69)$$

Replacing x with the pure imaginary number iy , we see that by multiplying out the terms in the power series, and collecting the real and imaginary parts, we have

$$e^{iy} = 1 + \frac{iy}{1!} + \frac{(iy)^2}{2!} + \frac{(iy)^3}{3!} + \frac{(iy)^4}{4!} + \frac{(iy)^5}{5!} + \cdots \quad (1.70)$$

$$= 1 + i\frac{y}{1!} - \frac{y^2}{2!} - i\frac{y^3}{3!} + \frac{y^4}{4!} + \cdots \quad (1.71)$$

$$= \left(1 - \frac{y^2}{2!} + \frac{y^4}{4!} - \cdots\right) + i\left(\frac{y}{1!} - \frac{y^3}{3!} + \frac{y^5}{5!} + \cdots\right) \quad (1.72)$$

$$= \cos(y) + i\sin(y) \quad (1.73)$$

where the last equality comes from recognizing the power series expansion for sine and cosine. The definition for general complex numbers $z = x + iy$ follows by forcing the exponential of a sum to be the product of two exponentials.

Some basic properties of the exponential function include:

- e^z is never equal to zero; however, it can reach every other complex number;
- $e^{\pi i/2} = i, e^{\pi i} = -1, e^{3\pi i/2} = -i, e^{2\pi i} = 1$;
- $e^{z_1+z_2} = e^{z_1}e^{z_2}$;
- $e^{z+2\pi i} = e^z$. That is, the function is periodic with period $2\pi i$;
- $|e^{x+iy}| = e^x$;
- e^{x+iy} is inside the unit circle if $x < 0$, on the unit circle if $x = 0$, outside the unit circle if $x > 0$;
- $e^z = 1$ if and only if $z = 2n\pi i$ for some integer n .

Each of these properties is easily checked, using a simple real-valued trig calculation. For instance, in the second item, we have $e^{\pi i/2} = \cos(\pi/2) + i \sin(\pi/2) = 0 + i * 1 = i$, as required.

It is useful to see how the exponential function maps regions the complex plane; the real part x determines a stretching factor e^x while the imaginary part iy determines a rotation by angle y around the origin. Figure 3 shows some simple regions in \mathbb{C} and the images under the complex plane. The basic idea is that a vertical box shape gets stretched out in the horizontal direction, while the vertical part gets wrapped around the origin.




Figure 1.3: Two strips in complex plane, mapped under the exponential, to curving strips.

The trig and hyperbolic functions are defined in terms of this complex exponential function, as described in the last section. It is easy to verify that the usual trig identities hold for all complex numbers, such as

- $\sin^2(z) + \cos^2(z) = 1$;
- $\sin(z_1 + z_2) = \sin(z_1) \cos(z_2) + \cos(z_1) \sin(z_2)$;
- $\cos(z_1 + z_2) = \cos(z_1) \cos(z_2) - \sin(z_1) \sin(z_2)$;
- $\tan^2(z) + 1 = \sec^2(z)$;

for all complex numbers z . For instance, to verify the first identity, we note:

$$\sin^2(z) + \cos^2(z) = \left(\frac{e^{iz} - e^{-iz}}{2i} \right)^2 + \left(\frac{e^{iz} + e^{-iz}}{2} \right)^2 \quad (1.74)$$

$$= \left(\frac{e^{2iz} - 2 + e^{-2iz}}{-4} \right) + \left(\frac{e^{2iz} + 2 + e^{-2iz}}{4} \right) \quad (1.75)$$

$$= \frac{2 + 2}{4} = 1, \quad (1.76)$$

as desired.

There are analogous identities for the hyperbolic functions. Rather than memorizing the identities, it is easiest just to notice the close connection between the hyperbolic and trig functions, namely that you go from trig to hyperbolic by introducing a factor of i . Thus,

- $\cosh(z) = \cos(iz)$;
- $\sinh(z) = \frac{1}{i} \sin(iz)$;
- $\tanh(z) = \frac{1}{i} \tan(iz)$;
- $\operatorname{sech}(z) = \cos(iz)$;
- $\operatorname{csch}(z) = i \csc(iz)$;
- $\coth(z) = i \cot(iz)$.

Again, these are easily verified by following the definitions given in terms of the exponential function. Thus, it is useful to remember that the exponential function is fundamental, and all the other elementary functions are defined in terms of it, and their properties follow from exponential properties.

1.2.2 The logarithm and complex powers

The logarithm function is defined as the inverse of the exponential function. More precisely, we write

$$\log(z) \equiv \ln |z| + i \operatorname{arg}(z), \quad (1.77)$$

where $\ln |z|$ is the natural logarithm of the absolute value of z , and $\operatorname{arg}(z)$ is the argument of z , also known as the phase angle. Remember that the argument is not uniquely defined, so we have to pick a branch cut for arg , and hence a choice of branch cut for \log . Usually, we assume a convention like

$$0 \leq \operatorname{arg}(z) < 2\pi \quad (1.78)$$

or

$$-\pi < \operatorname{arg}(z) \leq \pi. \quad (1.79)$$

To see that this is an inverse, note that if we write z in polar form, $z = |z|e^{i\theta}$, then we see θ is the arg of z and the real valued $\ln |z|$ will exponentiate to $|z|$, so

$$e^{\log z} = e^{\ln |z| + i \operatorname{arg}(z)} = e^{\ln |z|} e^{i \operatorname{arg}(z)} = |z| e^{i\theta} = z, \quad (1.80)$$

as desired. Note of course that the \log of zero is undefined.

The usual \log identities follow, except of course we have the problem of the 2π ambiguity in the argument. Thus we have

- $\log(1) = 0$;
- $\log(z_1 z_2) = \log(z_1) + \log(z_2) \bmod 2\pi i$;
- $\log(z^{-1}) = -\log(z) \bmod 2\pi i$;
- $\log(z^n) = n \log(z) \bmod 2\pi i$, for any integer n .

This last identity $\log(z^n) = n \log(z)$ we would like to hold with n replaced by any complex number.

To do so, we have to first define the complex power of a complex number, which we do as

$$z^a \equiv e^{a \log(z)}. \quad (1.81)$$

Note that implicit in this definition is the choice of branch cut for the log function. Be very careful about this, it means something apparently simple like $(-i)^i$ is not uniquely defined. For instance, in our first convention for angles, we have $\arg(-i) = 3\pi/2$ and so

$$(-i)^i = e^{i \log(-i)} = e^{i(\ln 1 + 3\pi i/2)} = e^{-3\pi/2} = 0.0090..., \quad (1.82)$$

while in our second convention, $\arg(-i) = -\pi/2$ and so

$$(-i)^i = e^{i \log(-i)} = e^{i(\ln 1 - \pi i/2)} = e^{\pi/2} = 4.8105.... \quad (1.83)$$

Freaky.

1.3 Analytic functions: doing calculus with complex numbers

Definition 1 A function $f(z)$ is differentiable at z_0 if the limit

$$\lim_{z \rightarrow z_0} \frac{f(z) - f(z_0)}{z - z_0} \text{ exists.} \quad (1.84)$$

When this limit exists, we denote it by $f'(z_0)$, which is called derivative of f at z_0 .

For this definition to work precisely, we require that the function $f(z)$ be defined at least in some small disk around the point z_0 . This way, the limit allows z to approach z_0 from any direction.¹

As an example, consider the simple function $f(z) = z^2$. Then we see

$$\lim_{z \rightarrow z_0} \frac{z^2 - z_0^2}{z - z_0} = \lim_{z \rightarrow z_0} \frac{(z - z_0)(z + z_0)}{z - z_0} = \lim_{z \rightarrow z_0} (z + z_0) = 2z_0. \quad (1.85)$$

Thus we have that the derivative of z^2 at z_0 is just $2z_0$. Since this works at any point z_0 in the complex plane, we see that $f(z) = z^2$ is differentiable everywhere, with derivative $f'(z) = 2z$. In other words,

$$\frac{d}{dz} z^2 = 2z. \quad (1.86)$$

Similarly, we can find the derivative of $f(z) = z^n$ by computing

$$\lim_{z \rightarrow z_0} \frac{z^n - z_0^n}{z - z_0} = \lim_{z \rightarrow z_0} \frac{(z - z_0)(z^{n-1} + z^{n-2}z_0 + \cdots + z z_0^{n-2} + z_0^{n-1})}{z - z_0} \quad (1.87)$$

$$= \lim_{z \rightarrow z_0} (z^{n-1} + z^{n-2}z_0 + \cdots + z z_0^{n-2} + z_0^{n-1}) \quad (1.88)$$

$$= n z_0^{n-1}, \quad (1.89)$$

which is what we expected. That is, the derivative of z^n is the function $n z^{n-1}$.

From this definition of derivative, it is easy to check that the usual sum, difference, product, and quotient rules for derivatives hold. That is

¹The limit is defined the same way as with real numbers: as z gets close to z_0 , the ratio inside the limit gets close to some fixed number L , which will be called $f'(z_0)$.

Theorem 1 Suppose f, g are complex, differentiable functions at point z_0 . Then so are the sum, difference, product, and quotient of f and g , with

$$(f + g)'(z_0) = f'(z_0) + g'(z_0) \quad (1.90)$$

$$(f - g)'(z_0) = f'(z_0) - g'(z_0) \quad (1.91)$$

$$(f * g)'(z_0) = f'(z_0) * g(z_0) + f(z_0) * g'(z_0) \quad (1.92)$$

$$\left(\frac{f}{g}\right)'(z_0) = \frac{f'(z_0)g(z_0) - f(z_0)g'(z_0)}{g(z_0)^2} \quad (1.93)$$

We also have the chain rule,

$$f(g(z))'|_{z=z_0} = f'(g(z_0))g'(z_0). \quad (1.94)$$

The proofs for this are exactly as in real calculus. Thus, we can now differentiate simple rational functions, so for example:

$$\frac{d}{dz} \left(\frac{z}{1+z^2} \right) = \frac{1(1+z^2) - z(2z)}{(1+z^2)^2} = \frac{1-z^2}{(1+z^2)^2}. \quad (1.95)$$

To compute the derivative of the exponential is a bit tricky. We start by computing the derivative at $z_0 = 0$, which involves computing the limit

$$\lim_{z \rightarrow 0} \frac{e^z - e^0}{z - 0}. \quad (1.96)$$

This is tricky. We try to simplify by taking the limit over restricted set of directions, say $z = x$ with x real. Then

$$\lim_{z \rightarrow 0} \frac{e^z - e^0}{z - 0} = \lim_{x \rightarrow 0} \frac{e^x - e^0}{x - 0} = \frac{d}{dx} e^x|_{x=0} = 1, \quad (1.97)$$

which we know is 1 from our usual real calculus derivatives. Similarly, if we restrict to the imaginary axis, and take the limit with $z = iy$ a pure imaginary, then

$$\lim_{z \rightarrow 0} \frac{e^z - e^0}{z - 0} = \lim_{y \rightarrow 0} \frac{e^{iy} - e^0}{iy} = \lim_{y \rightarrow 0} \frac{\cos(y) + i \sin(y) - 1}{iy} = \lim_{y \rightarrow 0} \frac{\sin(y)}{y} - i \frac{\cos(y) - 1}{y} = 1 + i0, \quad (1.98)$$

where again the last two limits come from real calculus.

Since we checked the limits in two directions, we can hope this limit holds:

$$\lim_{z \rightarrow 0} \frac{e^z - e^0}{z - 0} = 1. \quad (1.99)$$

Then, we can compute the derivative at any point z_0 as

$$\lim_{z \rightarrow z_0} \frac{e^z - e^{z_0}}{z - z_0} = \lim_{z \rightarrow z_0} \frac{e^{z_0}(e^{z-z_0} - e^0)}{z - z_0} = e^{z_0} \lim_{z \rightarrow z_0} \frac{e^{z-z_0} - e^0}{z - z_0} = e^{z_0} \cdot 1 = e^{z_0}. \quad (1.100)$$

That is, the derivative of the exponential is again the exponential.

Since the trig, hyperbolic, and log functions are all defined in terms of the exponential, we can find the derivatives knowing what we have already. Thus, for instance, we get derivatives just as with real functions, so

- $\frac{d}{dz} e^z = e^z$;

- $\frac{d}{dz} \sin(z) = \cos(z)$, since $\frac{d}{dz} \frac{e^{iz} - e^{-iz}}{2i} = \frac{ie^{iz} + ie^{-iz}}{2i} = \frac{e^{iz} + e^{-iz}}{2}$;
- $\frac{d}{dz} \cos(z) = -\sin(z)$, since $\frac{d}{dz} \frac{e^{iz} + e^{-iz}}{2} = \frac{ie^{iz} - ie^{-iz}}{2} = -\frac{e^{iz} - e^{-iz}}{2i}$;
- $\frac{d}{dz} \tan(z) = \sec^2(z)$;
- $\frac{d}{dz} \log(z) = \frac{1}{z}$ but look out for the branch cut;
- $\frac{d}{dz} z^n = nz^{n-1}$ for any integer n ;
- $\frac{d}{dz} z^a = az^{a-1}$ for any complex number a , but watch the branch cut;

etc.

Similarly, we can define antiderivatives in the obvious manner, so $-\cos(z) + C$ is the general antiderivative of the sine function, since the derivative of $-\cos(z) + C$ is $\sin(z)$.

A function is said to be *analytic* on an open set² in the complex plane if it is differentiable at every point in that set. All the examples above are analytic on the whole set where they are defined. Except for the case of log and complex powers, where they are analytic everywhere except the branch cut. A function is *entire* if it is analytic on the whole complex plane.

However, we are cheating. We are looking at a lot of nice functions that are nicely differentiable. But you should be warned: there are lots of smooth, complex valued functions on the complex plane that are not analytic. We look at this in the next section.

1.4 The Cauchy-Riemann equations

Any complex valued function can be split into its real and imaginary parts, with

$$f(z) = u(x, y) + iv(x, y), \text{ where } x + iy = z. \quad (1.101)$$

There is a close connection between the partial derivatives of u and v and the complex derivative of f . In fact, we have the following:

Theorem 2 *A function $f(z) = u(x, y) + iv(x, y)$ is (complex) differentiable at $z_0 = x_0 + iy_0$ if and only if u and v satisfy the Cauchy-Riemann equations at z_0 ; that is,*

$$\frac{\partial u}{\partial x}(x_0, y_0) = \frac{\partial v}{\partial y}(x_0, y_0), \quad \frac{\partial u}{\partial y}(x_0, y_0) = -\frac{\partial v}{\partial x}(x_0, y_0). \quad (1.102)$$

In this case, the derivative of f at z_0 is

$$f'(z_0) = \frac{\partial u}{\partial x}(x_0, y_0) + i \frac{\partial v}{\partial x}(x_0, y_0). \quad (1.103)$$

Let's see an example first. The function $f(z) = z^2$ is differentiable everywhere. Writing it as real plus imaginary, we see

$$f(z) = z^2 = (x^2 - y^2) + i(2xy) \quad (1.104)$$

so we have $u(x, y) = x^2 - y^2$ and $v(x, y) = 2xy$. Taking partial derivatives, we see

$$\frac{\partial u}{\partial x} = 2x = \frac{\partial v}{\partial y} \text{ and } \frac{\partial u}{\partial y} = -2y = -\frac{\partial v}{\partial x}. \quad (1.105)$$

²An open set is one that does not contain any boundary points. Each point in the set is surrounded by a small disk inside the same set.

Thus u, v satisfy the Cauchy-Riemann equations, and we can verify that

$$f'(z) = 2z = 2x + i(2y) = \frac{\partial u}{\partial x} + i \frac{\partial v}{\partial x}, \quad (1.106)$$

as required.

So, how do we prove this theorem. Well, frankly I'm not too concerned with proofs here. However, one direction of the "if and only if" is easy. If f is differentiable, you can look at the definition of derivative as a limit, and restrict to the direction $z = z_0 + x$, where x is a small real parameter. Taking the limit as $x \rightarrow 0$ will give the equation

$$f'(z_0) = \frac{\partial u}{\partial x}(x_0, y_0) + i \frac{\partial v}{\partial x}(x_0, y_0). \quad (1.107)$$

On the other hand, restricting to the direction $z = z_0 + iy$, where y is a small real parameter, and taking the limit as $y \rightarrow 0$ will give

$$f'(z_0) = \frac{1}{i} \left(\frac{\partial u}{\partial y}(x_0, y_0) + i \frac{\partial v}{\partial y}(x_0, y_0) \right). \quad (1.108)$$

Equating these last two equations yields the Cauchy-Riemann equations.

Proving the result in the reverse direction is quite a bit harder.

The C-R equations puts a very strong restriction on our choice of functions u and v . So for instance, the functions $u(x, y) = x^2 + y^2, v(x, y) = 2xy$ does not give a complex differentiable function, since

$$\frac{\partial u}{\partial y} = 2y \neq -\frac{\partial v}{\partial x}. \quad (1.109)$$

That is, the pair fails the second C-R equation, and thus the complex function

$$f(z) = (x^2 + y^2) + i(2xy) \quad (1.110)$$

is not differentiable in the complex sense. (Even though u and v are very nice, differentiable real functions on the plane.)

This points out that complex analytic functions are very special. In fact, they have some really remarkable properties, which we list in the next section.

1.5 Analytic functions

Suppose $f : D \rightarrow \mathbb{C}$ is an analytic function on an open set D inside the complex plane. That is, it is a complex valued function whose derivative is defined (by the limit) everywhere inside of D . Here are some of the properties of f .

- f is infinitely differentiable. That is, once the first derivative exists, $f'(z)$, then so does the second derivative $f''(z)$, the third derivative $f'''(z)$, and so on, to any order of derivative.
- The real and imaginary parts of $f(z) = u(x, y) + iv(x, y)$ are also infinitely differentiable. We usually say they are smooth function.
- The real and imaginary parts u, v are harmonic functions, that is

$$\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} = 0 \quad \frac{\partial^2 v}{\partial x^2} + \frac{\partial^2 v}{\partial y^2} = 0 \quad (1.111)$$

everywhere on D .

- f has a uniquely defined antiderivative on any simply connected³ subset of D , up to an

³A region is simply connected if it is in one piece, but has no holes in it.

additive constant.

- At any point z_0 , the function f can be expanded in a power series of the form

$$f(z) = \sum_{n=0}^{\infty} a_n (z - z_0)^n, \quad (1.112)$$

which converges in a disk of some radius R centered at z_0 . The radius R is at least as big as the radius of the biggest disk that fits in the region D , with center at z_0 .

- The coefficients in the power series expansion are given by the derivative of f , with

$$a_n = \frac{f^{(n)}(z_0)}{n!}. \quad (1.113)$$

We will come back to these properties later. I want to point out how special complex analytic functions are. Contrast them to real-valued differentiable functions, which may be once differentiable but not twice, and may not have convergent power series anywhere.

1.6 Power series

In high school we learn the geometric series:

$$\frac{1}{1-x} = 1 + x + x^2 + x^3 + \cdots. \quad (1.114)$$

What this means is we can compute the fraction on the left by adding up the numbers in the infinite series on the right. So, for instance, when $x = 1/2$, we get

$$\frac{1}{1-1/2} = \frac{1}{1/2} = 2 = 1 + \frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \frac{1}{16} + \cdots. \quad (1.115)$$

If that doesn't make sense to you, it is because you are not used to thinking in binary. A similar result happens in decimal, by taking $x = .1$, so

$$\frac{1}{1-0.1} = 1 + 0.1 + (0.1)^2 + (0.1)^3 + (0.1)^4 + \cdots. \quad (1.116)$$

That is,

$$1/0.9 = 1.11111\dots = 1 + .1 + .01 + .001 + .0001 + \cdots. \quad (1.117)$$

To be precise, we really should write the geometric series as a limit of partial sums, so we write

$$\frac{1}{1-x} = \lim_{N \rightarrow \infty} (1 + x + x^2 + x^3 + \cdots + x^N). \quad (1.118)$$

There is nothing special about x being real; we get a similar series with complex number z , so

$$\frac{1}{1-z} = \lim_{N \rightarrow \infty} (1 + z + z^2 + z^3 + \cdots + z^N). \quad (1.119)$$

However, this limit only converges for complex numbers z with $|z| < 1$.

Why? Well, using a little algebra, we can find

$$S_N = 1 + z + z^2 + z^3 + \cdots + z^N \quad (1.120)$$

$$zS_N = z + z^2 + z^3 + \cdots + z^{N+1} \quad (1.121)$$

$$S_N - zS_N = 1 - z^{N+1} \quad (1.122)$$

$$\text{so } (1 - z)S_N = 1 - z^{N+1}. \quad (1.123)$$

Thus we get

$$S_N = \frac{1 - z^{N+1}}{1 - z}, \quad (1.124)$$

and if we try to take the limit as $N \rightarrow \infty$, the fraction on the right will converge only if $|z| < 1$, in which case $z^{N+1} \rightarrow 0$, so

$$\lim S_N = \frac{1 - 0}{1 - z}. \quad (1.125)$$

This is typical behaviour for a power series. Given a sequence of complex numbers a_0, a_1, a_2, \dots , we can try to define an analytic function in the form

$$f(z) = a_0 + a_1z + a_2z^2 + a_3z^3 + a_4z^4 + \cdots. \quad (1.126)$$

There will be a disk in the complex plane, of the form $\{z \in \mathbb{C} : |z| < R\}$ where the series converges for all z inside the disk, and diverge for all z outside the disk. It is possible that $R = 0$, or $R = \infty$, but you get the idea. How you compute R is a bit involved, so we skip that.

Inside this disk, the function is analytic, so we can its derivative by differentiating term by term,

$$f'(z) = a_1 + 2a_2z + 3a_3z^2 + 4a_4z^3 + \cdots, \quad (1.127)$$

and also compute an antiderivative,

$$F(z) = \frac{a_0}{1}z + \frac{a_1}{2}z^2 + \frac{a_2}{3}z^3 + \frac{a_3}{4}z^4 + \frac{a_4}{5}z^5 + \cdots. \quad (1.128)$$

The derivative series, and antiderivative series, will have the same radius of convergence as the original series.

1.6.1 Standard series

Geometric series:

$$\frac{1}{1 - z} = 1 + z + z^2 + z^3 + z^4 + \cdots, \text{ for } |z| < 1. \quad (1.129)$$

Integrate the above, to get the log series

$$-\log(1 - z) = \frac{z}{1} + \frac{z^2}{2} + \frac{z^3}{3} + \frac{z^4}{4} + \frac{z^5}{5} + \cdots, \text{ for } |z| < 1. \quad (1.130)$$

Replace the z with $-z^2$ in the geometric series to get

$$\frac{1}{1 + z^2} = 1 - z^2 + z^4 - z^6 + z^8 + \cdots, \text{ for } |z| < 1. \quad (1.131)$$

Integrate this series in z^2 to obtain the series for arctan,

$$\arctan(z) = \frac{z}{1} - \frac{z^3}{3} + \frac{z^5}{5} - \frac{z^7}{7} + \frac{z^9}{9} + \cdots, \text{ for } |z| < 1. \quad (1.132)$$

The exponential series is given by

$$\exp(z) = 1 + \frac{z}{1!} + \frac{z^2}{2!} + \frac{z^3}{3!} + \frac{z^4}{4!} + \cdots, \text{ for all } z. \quad (1.133)$$

As an exercise, check that the derivative of this series gives back the same series. (HINT: just differentiate the series term-by-term.)

The cosine series is found by expressing the function as a sum of two exponentials:

$$\begin{aligned} \cos(z) &= \frac{1}{2}(e^{iz} + e^{-iz}) \\ &= \frac{1}{2}\left(1 + \frac{iz}{1!} + \frac{(iz)^2}{2!} + \frac{(iz)^3}{3!} + \frac{(iz)^4}{4!} + \cdots \right. \\ &\quad \left. + 1 + \frac{-iz}{1!} + \frac{(-iz)^2}{2!} + \frac{(-iz)^3}{3!} + \frac{(-iz)^4}{4!} + \cdots\right) \\ &= 1 - \frac{z^2}{2!} + \frac{z^4}{4!} - \frac{z^6}{6!} + \cdots \text{ for all } z. \end{aligned}$$

Similarly, the sine series can be found using exponentials, or by integrating the cosine series, so

$$\sin(z) = \frac{z}{1!} - \frac{z^3}{3!} + \frac{z^5}{5!} - \frac{z^7}{7!} + \cdots \text{ for all } z. \quad (1.134)$$

As an exercise, find the series expansion for hyperbolic cosine, and hyperbolic sine. (ANSWER: you get the same as the sine and cosine expansions, except all the negative terms become positive.)

1.7 Complex integrals

Just as we learned to do derivatives for complex functions, we can also do integrals. However, because we have a whole plane of complex numbers to work in, the integral of a function $f(z)$ from endpoints A to B must also specify a specific path γ connecting A and B . Thus the integral

$$\int_{\gamma} f(z) dz \quad (1.135)$$

include three pieces of information:

- the integrand $f(z)$, which is a function of a complex variable;
- the path γ , which is a continuous, piecewise differentiable curve in the plane; and
- the endpoint A, B , which are taken as specified by the path γ .

To compute the integral, we chop up the curve γ into N pieces, each piece with endpoints z_{k-1}, z_k ($k = 1, 2, \dots, N$), and do a Riemann sum, and finally taking the limit as N goes to infinity. So we define the integral as a limit

$$\int_{\gamma} f(z) dz = \lim_{N \rightarrow \infty} \sum_{k=1}^N f(z_k)(z_k - z_{k-1}). \quad (1.136)$$

Figure 4 shows an example of this, with the curve γ a semicircular arc connecting the point $A = 1$ to the point $B = i$. If we integrate the function $f(z) = 1$ along this curve, we see the Riemann sum just adds a bunch of short vectors spanning the curve, giving as the sum the vector that goes from 1 to i . Thus the integral will be

$$\int_{\gamma} 1 dz = i - 1. \quad (1.137)$$




Figure 1.4: A curve in the complex plane, connecting A to B . Integrating by breaking the curve up into short segments, and summing the segments.

In practice, we evaluate an integral by parametrizing the curve as a function $\gamma(t)$ from the real line to the complex plane, and using the formula

$$\int_{\gamma} f(z) dz = \int_a^b f(\gamma(t)) \gamma'(t) dt. \quad (1.138)$$

So, for instance, with the previous example, we can parametrize the semicircle as

$$\gamma(t) = \cos(t) + i \sin(t), \quad 0 \leq t \leq \pi/2, \quad (1.139)$$

$$\gamma'(t) = -\sin(t) + i \cos(t), \quad (1.140)$$

and get

$$\int_{\gamma} 1 dz = \int_0^{\pi/2} 1(-\sin(t) + i \cos(t)) dt = \int_0^{\pi/2} -\sin(t) dt + i \int_0^{\pi/2} \cos(t) dt = -1 + i, \quad (1.141)$$

which is the same result we saw as before using Riemann sums.

The third way to compute an integral is to use the Fundamental Theorem of Calculus. If we have an analytic function $F(z)$ with derivative $F'(z) = f(z)$, then

$$\int_{\gamma} f(z) dz = F(B) - F(A), \quad (1.142)$$

PROVIDED that the curve γ lies in a region on which $F(z)$ is analytic.⁴ In the previous example, the integrand $f(z) \equiv 1$ has antiderivative $F(z) \equiv z$, which is analytic everywhere, so we can compute

$$\int_{\gamma} 1 dz = F(i) - F(1) = i - 1, \quad (1.143)$$

⁴This analytic condition is very important, which we will see later.

which again agrees with our previous calculation.

Let's do a slightly harder integral, just to see again how the parameterization and antiderivative ideas work. We let γ be the same semicircle connecting 1 to i , take $f(z) = z$, and note its antiderivative is $F(z) = z^2/2$. Thus we have

$$\int_{\gamma} z dz = F(i) - F(1) = \frac{i^2}{2} - \frac{1^2}{2} = -1. \quad (1.144)$$

We should get the same answer by parameterizing the curve and solving two real integrals. So let's use the same sin,cos parameterization for γ , and compute

$$\int_{\gamma} z dz = \int_0^{\pi/2} \gamma(t) \gamma'(t) dt \quad (1.145)$$

$$= \int_0^{\pi/2} (\cos(t) + i \sin(t)) * (-\sin(t) + i \cos(t)) dt \quad (1.146)$$

$$= \int_0^{\pi/2} -2 \cos(t) \sin(t) dt + i \int_0^{\pi/2} \cos^2(t) - \sin^2(t) dt = -1, \quad (1.147)$$

as expected.

1.8 Some easy integrals

We take the function $f(z) = z^n$ and integrate it around the circle of radius R centered at the origin. That is, the curve γ starts at the point $z = R$, loops around the circle to travel a distance of $2\pi R$, and returns to $z = R$. Figure 5 shows such a circle (ignore the branch cut for now.)




Figure 1.5: A circle of radius R , which we will integrate around. The branch cut indicates where the log function fails to be analytic.

The parameterization for γ is given by

$$\gamma(t) = R \cos(t) + iR \sin(t) = Re^{it} \quad 0 \leq t \leq 2\pi \quad (1.148)$$

$$\gamma'(t) = -R \sin(t) + iR \cos(t) = iRe^{it} \quad (1.149)$$

and so the integral is

$$\int_{\gamma} z^n dz = \int_0^{2\pi} (Re^{it})^n iRe^{it} dt = i \int_0^{2\pi} (Re^{it})^{n+1} dt \quad (1.150)$$

$$= iR^{n+1} \int_0^{2\pi} e^{(n+1)it} dt = 0 \quad (1.151)$$

$$= iR^{n+1} \int_0^{2\pi} [\cos((n+1)t) + i \sin((n+1)t)] dt = 0, \text{ when } n \neq -1 \quad (1.152)$$

which is zero since the cos, sin functions are periodic, and their negative and positive parts cancel out in the integral. Now this is not surprising from the Fundamental Theorem of Calculus, since the function $f(z) = z^n$ has antiderivative $F(z) = \frac{1}{n+1} z^{n+1}$ and so we expect that by evaluating at the endpoints, we have

$$\int_{\gamma} z^n dz = F(R) - F(R) = 0. \quad (1.153)$$

So why is the case $n = -1$ any different? We note first that using the parameterization as above, we obtain

$$\int_{\gamma} z^{-1} dz = iR^0 \int_0^{2\pi} [\cos(0t) + i \sin(0t)] dt = i \int_0^{2\pi} 1 dt = 2\pi i, \quad (1.154)$$

which is certainly a non-zero answer. If we try to use the fundamental theorem of calculus, we note that the antiderivative of $f(z) = z^{-1}$ is $\log(z)$, so why don't we get

$$\int_{\gamma} z^{-1} dz = \log(r) - \log(R) = 0? \quad (1.155)$$

Well, the problem is that the function $\log(z)$ has to have a branch cut somewhere, and this branch cut slices through the curve γ . So the antiderivative is not analytic everywhere on γ , and so the fundamental theorem does not apply. Figure 5 shows the circle of radius R , and the branch cut, and you see how they slice through.

Now, frankly, this is bizarre. The function $f(z) = z^{-1}$ is perfectly well-defined and analytic on the curve γ . It turns out, though, that somehow the integral “sees” inside the curve and notices there is a place where z^{-1} is not analytic. Namely, it is undefined at the point $z = 0$. It is also interesting to notice that the integral always gives $2\pi i$, no matter what size the radius of the circle is.

The next section discusses the general results related to this phenomena.

1.9 The Cauchy Theorems

The Cauchy theorems tell us what happens when we integrate an analytic function around a simple closed curve. A lot of mathematics is done in describing precisely what a “simple closed curve” is, what the inside of a closed curve is, and what it means to deform one curve into another. We will keep things easy here and just take some intuitive ideas as to what these are.

A *simple closed curve* is a curve in the complex plane that is continuous, piecewise differentiable, does not intersect itself, and forms a loop (the end equals the start.) The inside of a closed curve




Figure 1.6: Some curves. a) is a simple oriented curve, which is not closed. b) is a simple closed curve, oriented counterclockwise. c) is a closed curve, but not simple as it intersects itself. d) is a closed curve which wraps around the point x twice. e) is also a simple closed curve, although it is pretty wiggly.

are all those points on the plane that are separated from infinity by the curve. Figure 6 gives a few examples of curves, some simple, some closed.

The first of the Cauchy theorems tells us that integrating an analytic function around a closed curve gives zero. Its precise statement is this:

Theorem 3 (Cauchy's theorem) *Suppose γ is a simple closed curve, and $f(z)$ is analytic both on γ and at every point inside γ . Then*

$$\int_{\gamma} f(z) dz = 0. \quad (1.156)$$

The next Cauchy theorem is the deformation theorem, and it concerns what happens when you have two curves that can be “deformed” on in to the other. The definition of what it means for two curves to be similar in this precise sense is as follows:

Two closed curves are *homotopic in a region D* if one can be continuously deformed into the other without ever leaving the region D . The idea is to stretch out one curve into the other, without breaking it, and without passing through any holes in the region. Figure 7 shows some curves which are homotopic, and not homotopic, in a region with a hole in it.

The deformation theorem simply says that integrating the same analytic function around two different, but homotopic curves, will give you the same value. The precise statement is as follows:

Theorem 4 (Deformation theorem) *Suppose γ_1 is a simple closed curve, that can be continuously deformed into another curve γ_2 inside a region D , and $f(z)$ is analytic on all of this region. Then*

$$\int_{\gamma_1} f(z) dz = \int_{\gamma_2} f(z) dz. \quad (1.157)$$

So, for instance, in Figure 7, integrating a function like $1/z$ around curve A will give you the same answer as integrating around curve B, provided the point 0 is inside the hole of the shaded region. In fact, from the last section, we know the value of the integral is $2\pi i$ for either curve. For curve C, though, the integral will be zero, by the first Cauchy theorem.

The third Cauchy theorem tells us that by introducing a simple singularity in the integrand, we can find the value of an analytic function at a point *inside* a curve γ by integrating *on* the curve.




Figure 1.7: Curves A and B can be continuously deformed one into the other, without leaving the shaded region. They are homotopic. However, curve C cannot be deformed into A or B, because it doesn't go around the hole.

The singularity we introduce is by replacing $f(z)$ with $f(z)/(z - z_0)$. The statement of the theorem is as follows:

Theorem 5 (Cauchy Integral Formula) *Suppose γ is a simple closed curve, oriented counterclockwise, and $f(z)$ is analytic both on γ and at every point inside γ . If z_0 is any point inside the curve, then*

$$f(z_0) = \frac{1}{2\pi i} \int_{\gamma} \frac{f(z)}{z - z_0} dz. \quad (1.158)$$

In particular, the value of f at any point z_0 inside the curve is completely determined by the values of f on the curve.

By differentiating this formula with respect to z_0 , we get immediately the Cauchy Integral Formula for derivatives:

Theorem 6 (CIF for derivatives) *Suppose γ is a simple closed curve, oriented counterclockwise, and $f(z)$ is analytic both on γ and at every point inside γ . If z_0 is any point inside the curve, then the k -th derivative is given as*

$$f^{(k)}(z_0) = \frac{k!}{2\pi i} \int_{\gamma} \frac{f(z)}{(z - z_0)^{k+1}} dz. \quad (1.159)$$

In particular, the value of f and all its derivatives at any point z_0 inside the curve is completely determined by the values of f on the curve.

As a simple application of the Cauchy integral theorem, we get Liouville's theorem which tells us that an analytic function which is defined everywhere on the complex plane either blows up at infinity (i.e. is unbounded), or it is constant:

Theorem 7 (Liouville's Theorem) *Suppose $f(z)$ is entire, and bounded on \mathbb{C} . Then $f(z)$ must be constant.*

The proof of this is so fast, we can do it here. By the CIF for derivatives, we compute the first derivative of f at an arbitrary point z_0 , as the integral

$$f'(z_0) = \frac{1}{2\pi i} \int_{\gamma} \frac{f(z)}{(z - z_0)^2} dz, \quad (1.160)$$

where we take γ to be a circle of radius R centered at z_0 . Since f is bounded, the integrand $f(z)/(z - z_0)^2$ is no bigger (in absolute value) than M/R^2 , where M is the bound on $|f(z)|$ and R^2 is the length of $(z - z_0)^2$. Thus we have

$$|f'(z_0)| \leq \frac{1}{2\pi} \int_{\gamma} \frac{M}{R^2} |dz| \leq \frac{1}{2\pi} \frac{M}{R^2} \text{length}(\gamma) = \frac{M}{R}, \quad (1.161)$$

since the length of γ is just $2\pi R$. We notice that the inequality $|f'(z_0)| \leq M/R$ holds for all values of radius R , so we must have $f'(z_0) = 0$. But now, since z_0 was arbitrary, we have $f'(z) = 0$ for all values of z . Thus, the derivative of f is zero everywhere, and so f must be a constant. QED.

Notice that this is much different than the case of functions on the real line. For instance, the functions $\cos(x)$ and $\frac{1}{1+x^2}$ are bounded, smooth functions on the whole real line. But they are not constant.

Corollary 8 (The Fundamental Theorem of Algebra) *Any polynomial in z must have a root in the complex plane. In particular, any polynomial of degree n can be factored as the product of n linear terms, with*

$$z^n + a_{n-1}z^{n-1} + a_{n-2}z^{n-2} + \cdots + a_1z + a_0 = (z - z_1)(z - z_2) \cdots (z - z_n), \quad (1.162)$$

where z_1, z_2, \dots, z_n are the n roots of the polynomial.

The proof is easy. We write

$$f(z) = \frac{1}{z^n + a_{n-1}z^{n-1} + a_{n-2}z^{n-2} + \cdots + a_1z + a_0}, \quad (1.163)$$

which is analytic everywhere that the polynomial is non-zero. If the polynomial is never zero, then this is an entire function. As $z \rightarrow \infty$, the function behaves like $1/z^n$, and so the limit at infinity is zero. Thus the entire function $f(z)$ is bounded on the complex plane. And thus by Liouville's theorem, it is constant. Which is obviously impossible (for $n \geq 1$.)

Thus it must be the case that the polynomial is zero at some point z_1 . Thus we have shown the existence of at least one root of the polynomial.

The rest is algebra. We can factor out the root, so

$$z^n + a_{n-1}z^{n-1} + a_{n-2}z^{n-2} + \cdots + a_1z + a_0 = (z - z_1)(z^{n-1} + b_{n-2}z^{n-2} + b_{n-3}z^{n-3} + \cdots + b_1z + b_0), \quad (1.164)$$

which is a linear term times a degree $n - 1$ polynomial. By the previous argument, this next polynomial has some root z_2 . Thus we can factor out another linear term, $(z - z_2)$. Repeat, until all n roots have been extracted. QED.

You might ask yourself: how can you prove the fundamental theorem of algebra using only algebra, and not analysis. Off the top of my head, I don't know!

1.10 Proofs of the Cauchy theorems

Cauchy theorem - use Green's theorem, and the Cauchy-Riemann equations.

Deformation theorem - connect the curves, and make one big curve where the function is analytic inside, and use the 1st Cauchy theorem.

CIF - reduce to computing $\int \frac{f(z)-f(z_0)}{z-z_0}$.

CIF for derivatives: just differentiate under the integral sign.

1.11 Applications of the Cauchy Integral Formula

Compute $\int_{-\infty}^{\infty} \frac{1}{1+x^2} dx = \pi$ using a Cauchy type integral. Also try $\int_{-\infty}^{\infty} \frac{\sin(x)}{x} dx$.

Chapter 2

Signals and systems

In signal processing, we are mainly interested in *signals* and *systems*. A typical signal might be a sound, a bit of music, a photographic image, a seismic vibration, or just about any physical phenomena that can be measured over time and/or space. Usually, we represent a signal by a function $f(t)$, where t is time, and $f(t)$ is the value measured at time t . For an image, we would have a function of two variables $f(x, y)$, where f represents the intensity of the image at position x, y in the plane, say.

A *system* takes one signal in, and outputs another signal. It could be a physical system: an earthquake in India starts a signal (vibrations) on one side of the earth, the earth transmits the vibrations to the other side (the system), and a new signal is felt in Calgary (the received vibrations). It could be an electrical system: a sound is picked up by a microphone (the input signal), the signal is passed to a stereo amplifier (the system), and the resulting amplified signal is output to the speakers (the output signal). It could be a computational system: a string of numbers is input to a computer, the computer churns away on the numbers (adding, subtracting, multiplying, etc – the system), and a string of numbers is output by the computer. It could be a combination of such systems: a digital camera captures a real image through its lens, the intensities are converted to a function $f(x, y)$, and then the computer mucks around with the values of the function to compute a sharper image, represented by a new function $g(x, y)$. Here, the input is the real image, the system is the camera/computer, the output is the function $g(x, y)$.

The point is: signals are function, and systems operate on signals.

2.1 Sampling

The signal $f(t)$ is a function on the real line \mathbb{R} . When we compute with a computer, usually we can't know everything about the function, or store it all on the computer. So we just evaluate the signal at a sequence of times t_n , and define a vector $\mathbf{x} = (\dots, x_{-2}, x_{-1}, x_0, x_1, x_2, \dots)$ with components

$$x_n = f(t_n), \quad \text{for all } n \in \mathbb{Z}. \quad (2.1)$$

For reasons that have mostly to do with engineering technology, we usually take the time samples t_n to be uniformly spaced. That is, we have a sequence of numbers separated by a uniform step Δt , and write $t_0 = 0, t_1 = \Delta t, t_2 = 2\Delta t, \dots$, and so the vector \mathbf{x} has components

$$x_n = f(n\Delta t), \quad \text{for all } n \in \mathbb{Z}. \quad (2.2)$$

Although the vector \mathbf{x} is infinitely long, that is it has infinitely many components, it is important to realize that most of the component x_n are just zero. Why? Because we can't measure back to

time minus infinity, or forward to plus infinity. So at some point we stop measuring, and can just assume everything else is zero.

The vector \mathbf{x} is called a sampled signal. Δt is called the sampling interval. $1/\Delta t$ is called the sampling rate.

2.2 Aliasing

The problem with sampling is that you lose information in the process. Two different functions $f(t)$ and $g(t)$ might get sampled and produce the same vector \mathbf{x} . For instance, suppose f is a sine wave, $f(t) = \sin(t)$ and g is the zero function, $g(t) = 0$. These are two very different signals. But, with the sampling interval $\Delta t = \pi$, we see that

$$\begin{aligned} x_n &= f(n\pi) = \sin(n\pi) \\ &= 0 \\ &= g(n\pi), \end{aligned}$$

so f and g get sampled to appear as the same vector \mathbf{x} , which happens to be the zero vector. This is called *aliasing*: one signal appears the same as another.

2.3 Frequency aliasing

There is a special kind of aliasing, where two sinusoidal signals appear the same under sampling. A complex sinusoid is a function of the form

$$f(t) = e^{2\pi i \omega t} = \cos(2\pi \omega t) + i \sin(2\pi \omega t). \quad (2.3)$$

This signal is periodic, which repeats itself at a rate of ω cycles per unit time. Eg. $f(t) = e^{2\pi i 60t}$ represents a 60 Hertz signal, where t is measured in seconds. Two signals

$$f(t) = e^{2\pi i \omega_1 t} \quad g(t) = e^{2\pi i \omega_2 t}, \quad (2.4)$$

will get aliased at a sample interval Δt if

$$f(n\Delta t) = g(n\Delta t) \text{ for all integers } n. \quad (2.5)$$

Equivalently,

$$(e^{2\pi i \omega_1 \Delta t})^n = (e^{2\pi i \omega_2 \Delta t})^n \text{ for all integers } n, \quad (2.6)$$

or more simply, if $e^{2\pi i \omega_1 \Delta t} = e^{2\pi i \omega_2 \Delta t}$. A bit of algebra shows this happens if $\omega_1 - \omega_2 = N/\Delta t$, for some integer N .

Thus, two sinusoids get aliased if the difference of their frequencies $\omega_1 - \omega_2$ is a multiple of the sampling rate $1/\Delta t$.

Usually, we are interested in measuring signals with frequencies in some interval $[-F, F]$. So, for instance, you might like to measure frequencies ω_1, ω_2 in $[-400\text{Hz}, 400\text{Hz}]$. The difference $\omega_1 - \omega_2$ could be as big as 800Hz. In order for this to not be a multiple of the sample rate, we have to choose a sample rate bigger than 800Hz. That is,

$$800\text{Hz} \leq \frac{1}{\Delta t}. \quad (2.7)$$

Being a little big lazy, we might choose the sample rate to be 1000Hz, and so the sampling interval is $\Delta t = .001$ second (a millisecond).

2.4 Sampled signals as a vector space

The sampled signal

$$\mathbf{x} = (\dots, x_{-2}, x_{-1}, x_0, x_1, x_2, \dots) \quad (2.8)$$

is supposed to look like a vector, just as you learned in linear algebra. It just happens to be infinitely long. You can still treat as you would regular vectors: add, subtract, pointwise multiply, scalar multiply, and take inner products. For instance, write

$$\begin{aligned} \mathbf{x} &= (\dots, 0, 0, 1, 2, 3, 0, 0, \dots) \\ \mathbf{y} &= (\dots, 0, 0, 4, 5, 6, 0, 0, \dots) \\ \mathbf{x} + \mathbf{y} &= (\dots, 0, 0, 5, 7, 9, 0, 0, \dots) \\ \mathbf{y} - \mathbf{x} &= (\dots, 0, 0, 3, 3, 3, 0, 0, \dots) \\ \mathbf{x} \cdot \mathbf{y} &= (\dots, 0, 0, 4, 10, 18, 0, 0, \dots) \\ 10\mathbf{x} &= (\dots, 0, 0, 10, 20, 30, 0, 0, \dots) \\ 10\mathbf{y} &= (\dots, 0, 0, 40, 50, 60, 0, 0, \dots) \\ \langle \mathbf{x}, \mathbf{y} \rangle &= 1 \cdot 4 + 2 \cdot 5 + 3 \cdot 6 = 32. \end{aligned}$$

It is a nuisance to keep writing all these zeros in the infinite vectors, so sometimes we shorten things to condensed vectors, like $\mathbf{x} = (1, 2, 3)$ and $\mathbf{y} = (4, 5, 6)$. In this form, we have the 0-th entry in \mathbf{x} as the first number that appears in the short vector. So $x_0 = 1, x_1 = 2, x_2 = 3$, and the rest are zero.

Chapter 3

Convolving signals

3.1 Z transforms

The Z transform of the signal

$$\mathbf{x} = (\dots, 0, 0, 1, 2, 3, 0, 0, \dots) \quad (3.1)$$

is the polynomial

$$X(Z) = 1 + 2Z + 3Z^2. \quad (3.2)$$

The Z transform of the signal

$$\mathbf{y} = (\dots, 0, 0, 4, 5, 6, 0, 0, \dots) \quad (3.3)$$

is the polynomial

$$Y(Z) = 4 + 5Z + 6Z^2. \quad (3.4)$$

In general, for a signal $\mathbf{x} = (\dots, x_{-2}, x_{-1}, x_0, x_1, x_2, \dots)$, the Z transform is the polynomial

$$X(Z) = \sum_n x_n Z^n. \quad (3.5)$$

Notice that if there are non-zero x_n with $n < 0$, the polynomial could include negative powers of Z. So, for instance, if

$$x_{-1} = 3, x_0 = 5, x_1 = 7, \quad (3.6)$$

then

$$X(Z) = 3Z^{-1} + 5 + 7Z. \quad (3.7)$$

3.2 Convolution

Any two signals \mathbf{x}, \mathbf{y} can be combined in a special operation called *convolution*. An easy way to define the convolution is using the Z transform. So, for example, if

$$\begin{aligned} \mathbf{x} &= (\dots, 0, 0, 1, 2, 3, 0, 0, \dots) \\ \mathbf{y} &= (\dots, 0, 0, 4, 5, 6, 0, 0, \dots), \end{aligned}$$

we have Z transforms $X(Z) = 1 + 2Z + 3Z^2$, $Y(Z) = 4 + 5Z + 6Z^2$. Take the product of the two polynomials,

$$(1 + 2Z + 3Z^2)(4 + 5Z + 6Z^2) = 4 + 13Z + 28Z^2 + 27Z^3 + 18Z^4, \quad (3.8)$$

which we can recognize as the Z transform of the vector

$$(\dots, 0, 0, 4, 13, 28, 27, 18, 0, 0, \dots) \quad (3.9)$$

which we define as the convolution of \mathbf{x} with \mathbf{y} . That is,

$$\mathbf{x} * \mathbf{y} = (\dots, 0, 0, 4, 13, 28, 27, 18, 0, 0, \dots). \quad (3.10)$$

There is a general formula for convolution: if you think about how polynomial multiplication works, it is pretty easy to see that the n -th entry in the vector $\mathbf{x} * \mathbf{y}$ will be a sum of terms like $x_j y_k$, where $j + k = n$. In other words, we can write

$$(\mathbf{x} * \mathbf{y})_n = \sum_k x_{n-k} y_k. \quad (3.11)$$

Because we know how polynomial multiplication works, we can observe that convolution works in either order and gives the same answer: $\mathbf{x} * \mathbf{y} = \mathbf{y} * \mathbf{x}$. Also, the convolution operation distributes over addition: $\mathbf{x} * (\mathbf{y} + \mathbf{z}) = (\mathbf{x} * \mathbf{y}) + (\mathbf{x} * \mathbf{z})$: because multiplication of polynomials distributes over addition. You can also move in scalar constants quite freely, so for instance, $3(\mathbf{x} * \mathbf{y}) = (3\mathbf{x}) * \mathbf{y} = \mathbf{x} * (3\mathbf{y})$, which is clear from the summation formula defining convolution.

3.3 Convolution as matrix-vector multiplication

Notice we can organize a matrix and a vector to get the same result as the convolution in the last section. We just do the simple \mathbf{x}, \mathbf{y} example:

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 2 & 1 & 0 & 0 & 0 \\ 3 & 2 & 1 & 0 & 0 \\ 0 & 3 & 2 & 1 & 0 \\ 0 & 0 & 3 & 2 & 1 \end{bmatrix} \begin{bmatrix} 4 \\ 5 \\ 3 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 4 \\ 13 \\ 28 \\ 27 \\ 18 \end{bmatrix}. \quad (3.12)$$

The matrix representing \mathbf{x} is constant along diagonals. That is called a Toeplitz matrix. \mathbf{y} is organized as a column vector. The usual matrix-vector product that we learned in linear algebra gives the resulting column vector representing $\mathbf{x} * \mathbf{y}$.

Notice that although \mathbf{x}, \mathbf{y} had only three non-zero entries, the matrix had to be 5 by 5, and the vector 5 by 1, in order for the arithmetic to work out. Similarly, for vectors with N non-zero entries (all in a row), we need matrices of size $2N + 1$.

3.4 Convolution as matrix-matrix multiplication

We observe that we can also represent convolution as the product to two Toeplitz matrices:

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 2 & 1 & 0 & 0 & 0 \\ 3 & 2 & 1 & 0 & 0 \\ 0 & 3 & 2 & 1 & 0 \\ 0 & 0 & 3 & 2 & 1 \end{bmatrix} \begin{bmatrix} 4 & 0 & 0 & 0 & 0 \\ 5 & 4 & 0 & 0 & 0 \\ 6 & 5 & 4 & 0 & 0 \\ 0 & 6 & 5 & 4 & 0 \\ 0 & 0 & 6 & 5 & 4 \end{bmatrix} = \begin{bmatrix} 4 & 0 & 0 & 0 & 0 \\ 13 & 4 & 0 & 0 & 0 \\ 28 & 13 & 4 & 0 & 0 \\ 27 & 28 & 13 & 4 & 0 \\ 18 & 27 & 28 & 13 & 4 \end{bmatrix}. \quad (3.13)$$

Check the matrix multiplication; you will see this is correct. Again, the matrices have to be large enough to give room for the whole convolution to appear.

You could write out infinite matrices, but it's too hard to typeset at the moment!

3.5 Convolution by flipping and shifting

You can also obtain a convolution by flipping the order of one of the vectors, taking point-wise products, and sum. So, for instance, with the same \mathbf{x}, \mathbf{y} as in the previous section, we flip around the \mathbf{y} and write it underneath the \mathbf{x} vector:

$$\begin{array}{cccccc} \dots & 0 & 0 & 1 & 2 & 3 & 0 \dots \\ \dots & 6 & 5 & 4 & 0 & 0 & 0 \dots \end{array} \cdot \quad (3.14)$$

The pointwise product is

$$\dots 0 \quad 0 \quad 4 \quad 0 \quad 0 \quad 0 \dots \quad (3.15)$$

which sums up to 4, the first component of the convolution. To get the second component, we shift \mathbf{y} once, and line up the vectors as

$$\begin{array}{cccccc} \dots & 0 & 0 & 1 & 2 & 3 & 0 \dots \\ \dots & 0 & 6 & 5 & 4 & 0 & 0 \dots \end{array} \cdot \quad (3.16)$$

The pointwise product is

$$\dots 0 \quad 0 \quad 5 \quad 8 \quad 0 \quad 0 \dots \quad (3.17)$$

which sums up to 13, the second component of the convolution.

To get the third component, we shift \mathbf{y} again, and line up the vectors as

$$\begin{array}{cccccc} \dots & 0 & 0 & 1 & 2 & 3 & 0 \dots \\ \dots & 0 & 0 & 6 & 5 & 4 & 0 \dots \end{array} \cdot \quad (3.18)$$

The pointwise product is

$$\dots 0 \quad 0 \quad 6 \quad 10 \quad 12 \quad 0 \dots \quad (3.19)$$

which sums up to 28, the third component of the convolution. And so on. This works in general.

3.6 Convolution as a system

Fixing a vector \mathbf{h} , we define a system that acts on signals \mathbf{x} as

$$\mathbf{x} \rightarrow A(\mathbf{x}) = \mathbf{h} * \mathbf{x} = \mathbf{y}. \quad (3.20)$$

That is, for input signal \mathbf{x} , our system outputs a signal \mathbf{y} that is computed as $\mathbf{y} = \mathbf{h} * \mathbf{x}$.

From our description of convolution as matrices, just like in linear algebra, we certainly expect that this system is linear. That is,

$$\begin{aligned} A(\mathbf{x}_1 + \mathbf{x}_2) &= A(\mathbf{x}_1) + A(\mathbf{x}_2), \\ A(\alpha \mathbf{x}) &= \alpha A(\mathbf{x}). \end{aligned}$$

This is easy to verify from the formulas for convolution. Equivalently, the first equation follows since convolution distributes over addition, $\mathbf{h} * (\mathbf{x}_1 + \mathbf{x}_2) = \mathbf{h} * \mathbf{x}_1 + \mathbf{h} * \mathbf{x}_2$, and the second equation follows since multiplication by a scalar commutes with polynomial multiplication.

3.7 The shift operator

The shift operator, S , is an example of a system. It takes an input \mathbf{x} and outputs the same vectors, except shifted to the right by one unit. So for instance, with

$$\begin{aligned}\mathbf{x} &= (\dots, 0, 0, 1, 2, 3, 0, 0, \dots) \\ \text{then } S\mathbf{x} &= (\dots, 0, 0, 0, 1, 2, 3, 0, \dots).\end{aligned}$$

In general, with

$$\begin{aligned}\mathbf{x} &= (\dots, x_{-2}, x_{-1}, x_0, x_1, x_2, x_3, \dots) \\ \text{then } S\mathbf{x} &= (\dots, x_{-3}, x_{-2}, x_{-1}, x_0, x_1, x_2, \dots).\end{aligned}$$

In terms of vector components, we can see that

$$(S\mathbf{x})_n = x_{n-1}. \quad (3.21)$$

It will be useful to notice that the operator S can be expressed as a convolution. We write δ^1 to be the special signal which is zero in all its components, except at the $n = 1$ place, where it takes the value 1. That is,

$$\delta^1 = (\dots, 0, 0, 0, 1, 0, 0, \dots), \quad (3.22)$$

where the 1 in the vector is placed exactly at the $n = 1$ place. Then

$$\delta^1 * \mathbf{x} = \sum_k x_{n-k} \delta_k^1 = x_{n-1} \cdot 1 + \text{a bunch of zeros.} \quad (3.23)$$

Thus $\delta^1 * \mathbf{x} = S\mathbf{x}$, which shows convolution by δ^1 is the same as shifting by 1.

It is easy to check that convolution by δ^n is the same as shifting by n steps (to the right, when n is bigger than zero, to the left when n is less than zero). This is the same as applying the operator S to the signal n times, which we denote by S^n .

You can think of δ^n as the vector

$$\delta^n = (\dots, 0, 0, 0, 0, 0, 1, 0, 0, \dots), \quad (3.24)$$

where the 1 in the vector is placed exactly at the n -th spot.

You can think of the operator S as a matrix of the form

$$S = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}, \quad (3.25)$$

which is a Toeplitz matrix with ones just below the main diagonal. So, for instance, we apply this matrix to the vector $\mathbf{x} = (\dots, 0, 0, 1, 2, 3, 0, \dots)$ in a short form, to see

$$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ 2 \\ 3 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \\ 2 \\ 3 \\ 0 \end{bmatrix}. \quad (3.26)$$

That is, this matrix applied to \mathbf{x} just shifts it.

Of course, for longer signals \mathbf{x} , we need a bigger matrix to represent the shift S . But the idea is the same. S will always have ones just below the main diagonal.

3.8 Signals as sums of δ^n 's.

Look how we can write a 3-component vector as a linear combination of three basic vectors:

$$(1, 2, 3) = 1 \cdot (1, 0, 0) + 2 \cdot (0, 1, 0) + 3 \cdot (0, 0, 1). \quad (3.27)$$

For our infinite vectors, we can do something similar:

$$\begin{aligned} \mathbf{x} &= (\dots, 0, 0, 1, 2, 3, 0, 0, \dots) \\ &= 1 \cdot (\dots, 0, 0, 1, 0, 0, 0, 0, \dots) + 2 \cdot (\dots, 0, 0, 0, 1, 0, 0, 0, \dots) + 3 \cdot (\dots, 0, 0, 0, 0, 1, 0, 0, \dots) \\ &= 1 \cdot \delta^0 + 2 \cdot \delta^1 + 3 \cdot \delta^2. \end{aligned}$$

That is, the vector \mathbf{x} is written as a linear combination of those basic vectors δ^n that we saw in the last section.

In general, any signal \mathbf{x} can be written in the form

$$\mathbf{x} = \sum_n x_n \delta^n. \quad (3.28)$$

Notice in this formula, \mathbf{x} is a vector, each δ^n is a vector, but the x_n are just numbers.

We need this for the proof in the next chapter.

Chapter 4

Linear, shift-invariant systems

We want to find out what are the linear, shift-invariant systems. Linear means the system is relatively simple in that sums of input signals map to sums of output signals, while shift-invariant means if we delay the input signal, the resulting output signal is the same as before, only delayed. This is also a reasonable assumption for a system that does not change over time.

We express these conditions on a system A in the following equations:

$$\begin{aligned} A(\mathbf{x}_1 + \mathbf{x}_2) &= A(\mathbf{x}_1) + A(\mathbf{x}_2), \text{ for all signals } \mathbf{x}_1, \mathbf{x}_2, \\ A(\alpha \mathbf{x}) &= \alpha A(\mathbf{x}), \text{ for all signals } \mathbf{x}, \text{ scalars } \alpha, \\ A(S\mathbf{x}) &= SA(\mathbf{x}), \text{ for all signals } \mathbf{x}, \end{aligned}$$

where S is the shift operator.

Theorem 9 *Suppose A is a linear, shift-invariant system. Then there is a vector \mathbf{h} so that A is just convolution by \mathbf{h} . That is,*

$$A(\mathbf{x}) = \mathbf{h} * \mathbf{x}, \text{ for all signals } \mathbf{x}. \quad (4.1)$$

The proof will go like this. \mathbf{h} is just A applied to the delta vector δ^0 , the vector with zeros everywhere except at the $n = 0$ spot. We then use shift invariance to find that A acting on any δ^n is just that delta function, convolved with \mathbf{h} . We then use linearity to conclude A acts on any vector by convolving with \mathbf{h} .

The details are like this. Let $\mathbf{h} = A(\delta^0)$, which is a vector, since A acts on the given input vector δ^0 to produce some output vector, which we call \mathbf{h} . We note the shift operator S takes the vector δ^0 to δ^1 , so by shift-invariance, we see

$$A(\delta^1) = A(S\delta^0) = S(A(\delta^0)) = S\mathbf{h} = \mathbf{h} * \delta^1, \quad (4.2)$$

since by definition, $A(\delta^0)$ is \mathbf{h} , and applying S to \mathbf{h} is the same as convolving by δ^1 .

Repeating this argument, we see that $A(\delta^n) = A(S^n\delta^0) = S^n A(\delta^0) = S^n \mathbf{h} = \mathbf{h} * \delta^n$. So now we know that A , applied to any of the delta vectors, just gives \mathbf{h} convolved with the vector.

Now, any vector \mathbf{x} is a linear combination of the δ^n , so by linearity of A , we have

$$A(\mathbf{x}) = A\left(\sum x_n \delta^n\right) = \sum x_n A(\delta^n) = \sum x_n \mathbf{h} * \delta^n = \mathbf{h} * \left(\sum x_n \delta^n\right) = \mathbf{h} * \mathbf{x}, \quad (4.3)$$

where at the second last equality, we use the fact that convolution by \mathbf{h} is linear.

And that's it. If you want to worry about mathematical details, you should worry about whether these infinite sums converge. For our purposes, we can just assume the \mathbf{x} is always given by a finite sum. (All but finitely many of the x_n are zero.)

4.1 Impulse response of a LSI system, Z transform

The vector \mathbf{h} that appeared in the last section is call the *impulse response* of the system A. From an engineering point of view, it is the response of the system to getting a whack at time $t = 0$. Knowing the impulse response basically tells you everything you need to know about the linear, shift invariant system.

The Z transform of the system A is given as the Z transform of the impulse response \mathbf{h} , which is the polynomial

$$H(Z) = \sum_k h_k Z^k, \quad (4.4)$$

which we discussed before.

Chapter 5

Building practical LSI systems

In this chapter we discuss two standard ways of implementing LSI systems on a computer: The FIR and IIR systems. Their advantage is that the outputs of the systems can be computed with only a finite amount of calculation by the computer.

5.1 FIR systems, minimum phase, maximum phase

A Finite Impulse Response (FIR) system is a linear, shift-invariant system where the impulse response function \mathbf{h} only has finitely many non-zero coefficients h_n . These are particularly useful in computations, since only finite sums are needed to compute them.

As an example, let's take $\mathbf{h} = (6, 1, -1)$ (in our short vector notation). The LSI system $\mathbf{x} \rightarrow \mathbf{h} * \mathbf{x} = \mathbf{y}$ is given by the formula

$$y_n = 6x_n + x_{n-1} - x_{n-2}, \text{ for all } n. \quad (5.1)$$

See how there is only 3 terms in the sum given by the convolution with \mathbf{h} .

This is an example of a *causal* system: the value of y_n depends only on the value of x_n and earlier coefficients in \mathbf{x} . In general, a LSI system will be causal if the impulse response \mathbf{h} is zero on the negative integers.

For the example above, the Z transform is the polynomial

$$H(Z) = 6 + Z - Z^2. \quad (5.2)$$

Notice this polynomial factors, as $(6 + Z - Z^2) = (3 - Z)(2 + Z)$. The linear terms $(3 - Z)$, $(2 + Z)$ are called couplets. From the couplets, we can see the zeros of this polynomial are $Z = 3$ and $Z = -2$. These zeros have magnitude bigger than one, so they live outside the unit circle in the complex plane. Because of this, we say the system is minimum phase.

By the Fundamental Theorem of Algebra, any polynomial can be factored into couplets. The zeros can be identified as points on the complex plane. If the zeros are all outside the unit circle, we say the system is minimum phase. If the zeros are inside the unit circle, we say the system is maximum phase. If some zeros are inside the circle, and some are outside, we say the system is mixed phase.

Minimum and maximum phase will have something to do with delays in our systems, which we will see later.

5.2 Recursive (IIR) systems

When the impulse response \mathbf{h} of a LSI system has infinitely many non-zero coefficients, we say the system is an Infinite Impulse Response (IIR) system. From a computational point of view, this is a nuisance since the direct convolution formula

$$y_n = \sum_k x_{n-k} h_k \quad (5.3)$$

has an infinite sum in it, which our computer will choke on.

However, there is a nice little structure that allows us to compute some IIR systems with only finite sums. This structure feeds back the earlier y values into the computation, and hence is called a *recursive* system.

It is probably best understood with a simple example.

Suppose we have a formula to compute output \mathbf{y} in terms of \mathbf{x} , in the form:

$$y_n = x_n + \frac{1}{2}y_{n-1}. \quad (5.4)$$

This is a finite sum (only two terms), and assuming we know the earlier output coefficients $\dots, y_{n-3}, y_{n-2}, y_{n-1}$, we can always compute the next y_n .

Let's see what the impulse response is. With $\mathbf{x} = \delta^0$, we have that all the x_n equal zero, except for $x_0 = 1$. We can assume then that all the earlier y_n are zero, for all $n < 0$. Then, we find

$$\begin{aligned} y_0 &= x_0 + \frac{1}{2}y_{-1} = 1 + \frac{1}{2} \cdot 0 = 1 \\ y_1 &= x_1 + \frac{1}{2}y_0 = 0 + \frac{1}{2} \cdot 1 = \frac{1}{2} \\ y_2 &= x_2 + \frac{1}{2}y_1 = 0 + \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4} \\ &\dots \\ y_n &= \frac{1}{2^n}, \text{ for each } n \geq 0 \end{aligned}$$

Thus, the impulse response is the signal $\mathbf{h} = (\dots, 0, 0, 1, \frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \dots)$, which is an infinite impulse response. The Z transform of the system is thus

$$H(Z) = 1 + \frac{Z}{2} + \frac{Z^2}{4} + \frac{Z^3}{8} + \dots + \frac{Z^n}{2^n} + \dots, \quad (5.5)$$

which we recognize as a geometric series, so we can sum it to

$$H(Z) = \frac{1}{1 - Z/2}. \quad (5.6)$$

Notice the linear polynomial in the denominator as a root at $Z = 2$. This point is called a pole of the filter. Because this root is outside the unit circle, we would call this a minimum phase system.

5.3 Rational functions and IIR systems

Here is another way to compute the Z transform of an IIR system. Using the same example as in the last section, we write a simple recursive system as

$$y_n = x_n + \frac{1}{2}y_{n-1}. \quad (5.7)$$

Pulling the y 's onto the lefthand side, we have

$$y_n - \frac{1}{2}y_{n-1} = x_n, \quad (5.8)$$

which we can express as a convolution,

$$\mathbf{g} * \mathbf{y} = \mathbf{x}, \quad (5.9)$$

where $\mathbf{g} = (1, -1/2)$ is the vector of coefficients for the convolution on the y side of the equation. Taking Z transforms of everything in the last equation, we have

$$G(Z)Y(Z) = X(Z), \quad (5.10)$$

which we rewrite as

$$Y(Z) = \frac{1}{G(Z)}X(Z). \quad (5.11)$$

Notice that since \mathbf{g} is a short vector, its Z transform is easily computed, as $G(Z) = 1 - Z/2$. The fraction $\frac{1}{G(Z)} = \frac{1}{1-Z/2}$ is the system response of the LSI with impulse response \mathbf{h} as in the last section. Notice it agrees with the geometric series summation we did in the last section.

A general recursive system is written like this:

$$y_n = \sum_{k=0}^N f_k x_{n-k} - \sum_{k=1}^M g_k y_{n-k}, \quad (5.12)$$

where $\mathbf{f} = (f_0, f_1, \dots, f_N)$ and $\mathbf{g} = (1, g_1, g_2, \dots, g_M)$ are some fixed coefficients that will determine our LSI system. By carefully choosing $g_0 = 1$, we can rewrite the last equation as two convolutions, so

$$\sum_{k=0}^M g_k y_{n-k} = \sum_{k=0}^N f_k x_{n-k}, \quad (5.13)$$

or more succinctly in vector form as

$$\mathbf{g} * \mathbf{y} = \mathbf{f} * \mathbf{x}. \quad (5.14)$$

Taking Z transforms, we have

$$G(Z)Y(Z) = F(Z)X(Z), \quad (5.15)$$

which we write in input-output form as

$$Y(Z) = \frac{F(Z)}{G(Z)}X(Z). \quad (5.16)$$

The function

$$H(Z) = \frac{F(Z)}{G(Z)} = \frac{f_0 + f_1 Z + f_2 Z^2 + \dots + f_N Z^N}{g_0 + g_1 Z + g_2 Z^2 + \dots + g_M Z^M} \quad (5.17)$$

is the Z transform of the LSI systems. Notice it is the ratio of two polynomials in Z , which is an example of a rational function in complex analysis.

5.4 Zeros and poles

By the fundamental theorem of algebra, the polynomials in a rational function can always be factored into a product of linear terms, so

$$\begin{aligned} H(Z) &= \frac{F(Z)}{G(Z)} = \frac{f_0 + f_1Z + f_2Z^2 + \cdots + f_NZ^N}{g_0 + g_1Z + g_2Z^2 + \cdots + g_MZ^M} \\ &= \frac{f_N}{g_M} \frac{(Z - z_1)(Z - z_2) \cdots (Z - z_N)}{(Z - p_1)(Z - p_2) \cdots (Z - p_M)} \end{aligned}$$

where the complex numbers z_1, \dots, z_N are the roots of the polynomial on top, and p_1, \dots, p_M are the roots of the polynomial on the bottom. The z_k are called the *zeros* of the function $H(Z)$, and the p_k are called the *poles* of the function. This is because, as a complex analytic function, $H(Z)$ is equal to zero at the z_k , and blows up (division by zero) at the p_k .

5.5 Stability of recursive LSI systems

Fix a complex number α and define a recursive LSI system as

$$y_n = x_n + \alpha y_{n-1}. \quad (5.18)$$

Computing the impulse response as we did above for the simple example, we see that the impulse response is just

$$\mathbf{h} = (\dots, 0, 0, 1, \alpha, \alpha^2, \alpha^3, \dots, \alpha^n, \dots). \quad (5.19)$$

This sequence will go to zero if $|\alpha| < 1$, but will blow up with $|\alpha| > 1$. The first case is called *stable*, while the second case is called *unstable*. Notice that the reciprocal $1/\alpha$ is the single pole in this LSI system, since $G(Z) = 1 - \alpha Z$ has a root at $Z = 1/\alpha$. Thus the stability of this simple system depends on the location of the pole $1/\alpha$: outside the unit circle, the system is stable. Inside the unit circle, it is unstable.

This result generalizes. Given a rational function

$$H(Z) = \frac{F(Z)}{G(Z)}, \quad (5.20)$$

the corresponding LSI system is stable if all the poles are outside the unit circle. It is unstable if the poles are inside the unit circle.

Try not to worry about the case when the poles are exactly on the unit circle. Generally speaking, they tend to be unstable due to small numerical errors that build up on the computer.

5.6 Example: a smoothing system. FIR and IIR

Let's imagine you want to design a system that smooths out any given input signal. For instance, with a photo, you might want to blur out some details so that certain faces cannot be identified. Or, if you are building an electrical system, you might realize sudden jumps in current are bad for the electronic devices, so you want to smooth out any jumps. Or you are building a car, and sudden jerks and shocks in the motion will annoy the passengers, so you want to smooth those out.

Without explaining why, here are two smoothers. The first is an FIR, non-recursive system. We choose \mathbf{h} to be the vector

$$\mathbf{h} = (\dots, 0, 0, 1, 1, 1, 1, 1, 1, 1, 1, 1, 0, 0, \dots). \quad (5.21)$$

That is, \mathbf{h} has 10 ones in it, the rest are zero. (This is sometimes called a boxcar, because its graph looks like a chunky boxcar on a train.) The system $\mathbf{x} \rightarrow \mathbf{h} * \mathbf{x} = \mathbf{y}$ is given from the convolution as the term-by-term sum

$$y_n = x_n + x_{n-1} + x_{n-2} + \dots + x_{n-9}. \quad (5.22)$$

If we take a simple boxcar $\mathbf{x} = (\dots, 0, 0, 1, 1, 1, \dots, 1, 1, 1, 0, 0, \dots)$ as an input, it is easy to compute by hand what the output is: the ones will sum up as an increasing sequence, gets as big as 10, stays there for a while, and comes down. (You should try this yourself.) The output is thus

$$\mathbf{y} = (\dots, 0, 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 10, 10, \dots, 10, 10, 9, 8, 7, 6, 5, 4, 3, 2, 1, 0, 0, \dots). \quad (5.23)$$

This ramp \mathbf{y} is a bit smoother than the original boxcar input \mathbf{x} . Figure 5.1 shows the \mathbf{x} and the \mathbf{y} .




Figure 5.1: FIR smoother: A boxcar input \mathbf{x} and its smoothed output \mathbf{y} .

Now, we could choose a different \mathbf{h} to get a different kind of smoothing operation. But, let's now do something completely different. Let's build a recursive system of the form

$$y_n = x_n + 0.9y_{n-1}. \quad (5.24)$$

Now, with the same input $\mathbf{x} = (\dots, 0, 0, 1, 1, 1, \dots, 1, 1, 1, 0, 0, \dots)$, we can compute the output by hand, as

$$\begin{aligned} y_0 &= x_0 + 0.9y_{-1} = 1 + 0 = 1 \\ y_1 &= x_1 + 0.9y_0 = 1 + 0.9 * 1 = 1.9 \\ y_2 &= x_2 + 0.9y_1 = 1 + 0.9 * 1.9 = 2.71 \end{aligned}$$

$$\begin{aligned}
 y_3 &= x_3 + 0.9y_2 = 1 + 0.9 * 2.71 = 3.439 \\
 y_4 &= 4.0951 \\
 y_5 &= 5.2170 \\
 y_6 &= 5.6953 \\
 y_7 &= 6.1258 \\
 &\text{etc.}
 \end{aligned}$$

As you can see, the numbers aren't pretty. But as you work it out, you will see they go up for a while, level off, and then go down. We can plot this in MATLAB, and the input/output looks something like Figure 5.2.




Figure 5.2: IIR smoother: A boxcar input \mathbf{x} and its smoothed output \mathbf{y} .

Again, we see that the output \mathbf{y} is a smoothed out version of the input \mathbf{x} . However, the details are different from the FIR case in Figure 5.1. For instance, we don't have a linear ramp going up in the output, instead it is more of a smooth concave taper that creeps up to the value 10. And on the way down, the curve is a smooth, convex curve leveling off at zero.

Which is better? Well, it depends on what you are using these things for. Notice that the IIR system only uses one multiplication and one addition, to compute each y_n . This can be an advantage when speed is required. The FIR requires 10 additions, per sample y_n , which is a lot.

Much of the interest in signal processing is in designing good systems, FIR or IIR, which are fast to compute and do the required modifications on signals. We will get to this later.

Chapter 6

Frequency response of a LSI system

Given a linear, shift invariant system A , we know there is an impulse response vector \mathbf{h} that completely describe the system as a convolution,

$$x \mapsto A(\mathbf{x}) = \mathbf{h} * \mathbf{x}. \quad (6.1)$$

From this impulse response, we define the *frequency response* of the system as the function

$$H(\omega) = \sum_k h_k e^{-i\omega k}. \quad (6.2)$$

Now, remembering that the Z transform is given by $H(Z) = \sum_k h_k Z^k$, we see immediately that the frequency response is obtained from the Z transform by setting $Z = e^{-i\omega}$. That is, we have that the frequency response is given as

$$H(\omega) = H(Z), \text{ where } Z = e^{-i\omega}. \quad (6.3)$$

Now, there might be a little confusion here because $H(Z)$ can be thought of as a function of a complete variable, but $H(\omega)$ is a function of the real parameter ω . But this notation is standard in DSP, so we are kind of stuck with it.

You should notice that $H(\omega)$ is periodic in variable ω , since the exponential $Z = e^{-i\omega}$ is periodic. So we have that $H(\omega + 2\pi) = H(\omega)$, and thus $H(\omega)$ is completely determined by its values on the interval $[-\pi, \pi]$.

You should also be aware that some people (including me) like to use the definition

$$H(\omega) = \sum_k h_k e^{-2\pi i \omega k}, \quad (6.4)$$

in which case $H(\omega)$ is a periodic function, with period one. It is completely determined by its values on the interval $[-1/2, 1/2]$. Our textbook author Karl uses the other convention.

The function $H(\omega)$ tells us the frequency response of the system. Specifically, if we input a sinusoid \mathbf{x} of the form

$$x_n = e^{i\omega n} = \cos(\omega n) + i \sin(\omega n), \quad (6.5)$$

where ω is a fixed number, then the output from our system A is just a multiple of the original,

$$A(\mathbf{x}) = H(\omega)\mathbf{x}. \quad (6.6)$$

That is, the output $A(\mathbf{x})$ is just the original input \mathbf{x} , multiplied by the fixed number $H(\omega)$. Roughly speaking, this says “sine wave in, yields sine wave out.”

Except we are using complex sinusoids here.

To see why this happens, let's compute the convolution of $x_n = e^{i\omega n}$ with the impulse response \mathbf{h} , to see the result of the system $A(\mathbf{x}) = \mathbf{h} * \mathbf{x} = \mathbf{y}$. We have

$$\begin{aligned} y_n &= \sum_k x_{n-k} h_k \\ &= \sum_k e^{i\omega(n-k)} h_k \\ &= e^{i\omega n} \sum_k e^{i\omega k} h_k \\ &= x_n H(\omega). \end{aligned}$$

That is, in vector notation,

$$\mathbf{y} = H(\omega)\mathbf{x}, \quad (6.7)$$

so the output \mathbf{y} is some multiple of the input \mathbf{x} .

This is another way of saying that the complex sinusoids are eigenvectors for the LSI system A , with eigenvalues given as $H(\omega)$.

6.1 Frequency response of a simple smoother

Here is a simple example. Define an FIR system by the formula

$$y_n = \frac{1}{3}(x_{n-1} + x_n + x_{n+1}). \quad (6.8)$$

That is, we just average the sample x_n with its nearest neighbours. This will tend to smooth out a given input signal, like we saw in the example in Section 21.

The impulse response \mathbf{h} has $h_{-1} = h_0 = h_1 = 1/3$, with all the other coefficients zero. The Z transform of the system is

$$H(Z) = \frac{Z^{-1} + 1 + Z}{3}, \quad (6.9)$$

and the frequency response is obtained by setting $Z = e^{-i\omega}$, so

$$H(\omega) = \frac{e^{i\omega} + 1 + e^{-i\omega}}{3} = \frac{1}{3}(1 + 2\cos\omega). \quad (6.10)$$

A plot of this function is in Figure 6.1.

What should you notice. Well, first, the function $H(\omega)$ is real valued, which is nice, and perhaps unexpected since we started with complex exponentials. Second, the function is symmetric about $\omega = \pi$, which also may seem unexpected, but happens a lot in filter design. Third, the function is zero for certain values of ω , which means for certain sinusoid inputs, we can get a zero output. For instance, with $\omega = 2\pi/3$, we have $H(\omega) = 0$. This means if we set input

$$x_n = e^{2\pi i n/3} = \left(-\frac{1}{2} + \frac{\sqrt{3}}{2}i\right)^n \quad (6.11)$$

we get output $\mathbf{y} = 0$. (Think about your complex numbers. Any three consecutive x_n are the three different cube roots of unity, so the three of them sum to zero. Hence their three point average is zero. Hence the total output is zero.)




Figure 6.1: Plot of the frequency response $H(\omega)$ for a three point smoother.

You also should notice that $H(\omega)$ is negative for some values of ω , which indicates the sign of the input signal can flip. This is easy enough to see, by taking \mathbf{x} to be an alternating series of ± 1 , and compute the output as an alternating series of $\pm 1/3$, but with the opposite signs. That is, using \mathbf{x} like this, we compute the three point moving average to find \mathbf{y} as:

$$\begin{aligned} \mathbf{x} &= (\dots, 1, -1, 1, -1, 1, -1, \dots) \\ \mathbf{y} &= (\dots, -\frac{1}{3}, \frac{1}{3}, -\frac{1}{3}, \frac{1}{3}, -\frac{1}{3}, \frac{1}{3}, \dots) \end{aligned} \quad (6.12)$$

See how the sign flips? See also how the magnitude is $-1/3$, just like the value of $H(\omega)$ on the graph, at its lowest point.

The point is, the function $H(\omega)$ tells us a lot about how the system behaves. In fact, from $H(\omega)$, we can learn everything about the system.

6.2 A better three point smoother

Frankly, in the last example, it is weird to have negative values for $H(\omega)$, and weird for it to be zero at a frequency $\omega = 2\pi/3$. It seems it would be better for a smoother to pass the low frequencies, and attenuate the high frequencies, only hitting zero at the point $\omega = \pi$. That is, we might want a filter response like the curve in Figure 6.2. With some inspired thinking, we can get this using the cosine function, shifted up by one, so we write

$$H(\omega) = \frac{1}{2}(1 + \cos \omega) = \frac{1}{4}(2 + e^{-i\omega} + e^{i\omega}), \quad (6.13)$$

as our desired frequency response. Examining this carefully, we see that this is the frequency response of a system with Z transform

$$H(Z) = \frac{1}{4}(Z^{-1} + 2 + Z). \quad (6.14)$$




Figure 6.2: Plot of the frequency response $H(\omega)$ for a better three point smoother.

From this, we see the impulse response is

$$\mathbf{h} = (\dots, 0, 0, \frac{1}{4}, \frac{2}{4}, \frac{1}{4}, 0, 0, \dots). \quad (6.15)$$

Thus, we can define our better three point smoother using this \mathbf{h} , and find

$$y_n = \frac{x_{n-1} + 2x_n + x_{n+1}}{4}. \quad (6.16)$$

This way, the output y_n is a weighted average of input x_n and its nearest neighbours. It turns out, for many reasons, that this is a better smoother for our system.

6.3 Linear algebra notes: Eigenvalues and eigenvectors

This has been moved to an appendix at the end of the book.

6.4 A rant on radians

The book (and many practitioners) insist on parameterizing frequencies ω in measures of radians per unit time. So, for instance, with a sinusoidal signal in the form

$$f(t) = e^{i\omega t} \quad (6.17)$$

where t is measured in seconds, $\omega = 1$ corresponds to oscillations of one radian per second, while $\omega = 10$ corresponds to oscillations of ten radians per second. Figure 6.3 shows plots of these two sinusoids (the imaginary part), and it is hard to see what 1 rad, or 10 rad, means.




Figure 6.3: Sine waves at 1 rad per second (left) and 10 rad per second (right). These are not natural units!

Compare this to the more natural system of parameterizing frequencies ω in measures of cycles per second. With this parameterization, we write a sinusoidal signal in the form

$$f(t) = e^{2\pi i \omega t}. \quad (6.18)$$

Again with time measured in seconds, then $\omega = 1$ corresponds to oscillations of one cycle per second, while $\omega = 10$ corresponds to oscillations of ten cycles per second. Figure 6.4 shows plots of these two sinusoids (the imaginary part), and we clearly see in the plots the one cycle, and ten cycles, respectively.




Figure 6.4: Sine waves at 1 cycle per second (left) and 10 cycles per second (right). This is natural.

With this in mind, from now on in the course we will use natural parameterization of ω in cycles per unit of time, even though this conflicts with the book. (It turns out to have some mathematical benefits as well, in that the natural Fourier transform is automatically normalized, which we will see later.)

6.5 Symmetries of frequency response $H(\omega)$ in natural units.

With these natural units in mind, we define the frequency response of a system with impulse response $\mathbf{h} = (\dots, h_{-1}, h_0, h_1, \dots)$ as

$$H(\omega) = \sum_k h_k e^{-2\pi i k \omega}, \quad (6.19)$$

which is the eigenvalue for the (sampled signal) eigenvector \mathbf{x} at the corresponding frequency ω in cycles per unit time given by

$$\mathbf{x}_n = e^{2\pi i n \omega}. \quad (6.20)$$

You might like to check the algebra in Section 22 for this new parameterization. Everything works out.

What's different is that this new function $H(\omega)$ is 1-periodic; that is, it repeats itself as ω is incremented by units of one. Thus,

$$\begin{aligned} H(0) &= H(1) = H(2) = H(3) = \dots \\ H(0.2) &= H(1.2) = H(2.2) = H(3.2) = \dots \\ H(0.4) &= H(1.4) = H(2.4) = H(3.4) = \dots \\ &\text{etc.,} \end{aligned}$$

and in general, $H(\omega) = H(\omega + N)$ for all integers N , and all real numbers ω .

This implies that it is enough just to look at the values of $H(\omega)$ on the interval $[0, 1]$, since it just repeats itself after that. So for instance, with the three point smoother considered earlier, where $\mathbf{h} = (1/4, 1/2, 1/4)$, we have that

$$H(\omega) = \frac{1}{2}(1 + \cos 2\pi\omega), \quad (6.21)$$

which we see plotted in Figure 6.5. Observe how the values on the interval $[0, 1]$ (plotted in the solid




Figure 6.5: 1-periodicity of the frequency response curve $H(\omega)$ for a 3 point smoother.

line) just get repeated in each subsequent interval. This is the periodicity of $H(\omega)$, with period one.

It is common to plot only the values of $H(\omega)$ in the interval $[0, 1]$, as shown in Figure 6.6 (left). Also common is to plot it in the interval $[-0.5, 0.5]$, Figure 6.6 (center). And since these frequency response functions are often symmetric about 0, it is good enough to plot just the values in $[0, 0.5]$, Figure 6.6 (right).

You should get used to these three plots, as you will see all three in in different contexts. MATLAB often displays the one on $[0, 1]$, but also sometimes on $[-0.5, 0.5]$. Make sure you know what you are looking at.




Figure 6.6: The frequency response curve of a 3 point smoother. Three standard plots.

And what do these plots mean? Well, $\omega = 0$ corresponds to a sinusoid with zero frequency – one that does not oscillate at all. $H(0)$ gives the DC response of the system – what it does to constant signals. In our example $H(0) = 1$, so a DC signal is passed unchanged. The case $\omega = 0.5$ corresponds to a sinusoidal signal that is oscillating at $1/2$ of the sampling rate. For a system that samples at 10kHz, this would be a 5kHz signal. In our example, $H(1/2) = 0$, so our 3 point smoother will kill off this signals (a 5kHz signal is attenuated to zero amplitude). At an intermediate point, say $\omega = 0.25$, we find $H(0.25) = .5 * (1 + \cos(\pi/2)) = .5$, so a signal at $1/4$ the sampling rate will be attenuated by a factor of $1/2$. Eg. at 10kHz sampling, the 2.5kHz signal is attenuated by one half.

It is important to remember that ω is a measure of frequency relative to the sampling rate. So, for instance with a sampling rate of 10kHz, the $\omega = 0.5$ corresponds to a signal at frequency 5kHz. For a CD system, the sampling rate is 44.1kHz. In this case, the same $\omega = 0.5$ corresponds to a signal at 22.05kHz. Just remember that ω is relative to the sampling rate (NOT the Nyquist rate).

In our example of the three point smoother, we saw the symmetry $H(-\omega) = H(\omega)$; equivalently, H is even. In the assignment, we saw that for symmetric impulse response, the frequency response is a sum of cosines,

$$H(\omega) = h_0 + 2 \sum_k h_k \cos(2\pi k\omega), \quad (6.22)$$

and thus this $H(\omega)$ is also even. That is, $H(-\omega) = H(\omega)$.

This even property is not true in general. However, we have the following:

Theorem 10 Suppose the impulse response $\mathbf{h} = (\dots, h_{-1}, h_0, h_1, \dots)$ is real-valued. Then the frequency response function $H(\omega)$ satisfies a skew-symmetry,

$$H(-\omega) = \overline{H(\omega)}, \text{ for all } \omega, \quad (6.23)$$

where the bar above means complex conjugate.

This is an important theorem since many systems we design must have real coefficients in them, so we get the skew symmetry for free. It also implies the absolute value of the frequency response is an even function, so

$$|H(-\omega)| = |H(\omega)| \text{ for all } \omega, \quad (6.24)$$

since the absolute value is the same on any complex number and its conjugate.

The proof of the theorem is easy. We just observe the sum at $-\omega$ is

$$\begin{aligned}
 H(\omega) &= \sum_k h_k e^{-2\pi i k(-\omega)} \\
 &= \sum_k h_k \overline{e^{-2\pi i k \omega}}, \text{ the complex conjugate} \\
 &= \sum_k \overline{h_k e^{-2\pi i k \omega}}, \text{ since } h_k \text{ is real} \\
 &= \overline{\sum_k h_k e^{-2\pi i k \omega}}, \text{ by linearity} \\
 &= \overline{H(\omega)},
 \end{aligned}$$

as desired.

6.6 Magnitude and phase response

Since $H(\omega)$ is a complex number for each value of ω , it can be expanded in the form of a real amplitude $|H| = |H(\omega)| \geq 0$ and a phase $\phi = \phi(\omega)$, with

$$H(\omega) = |H(\omega)| e^{2\pi i \phi(\omega)}. \quad (6.25)$$

Both the amplitude and the phase have a specific physical meaning for action of the system on the sinusoids of frequency ω .

Figure 6.7 shows this. Here we see a sine wave of a particular fixed frequency ω input into a system, and we plot right next to it the output. We see the output is still basically a sine wave, but with a small height, and shifted to the right. The change in amplitude is given by $A = |H(\omega)|$, the amplitude spectrum, and the amount of shift is given by $\phi = \phi(\omega)$, the phase spectrum. We have to be careful, though: for a positive ϕ , the shift is to the left; for negative ϕ , it is to the right.

To see this, observe that if we start with a sampled sinusoid $x_n = e^{2\pi i \omega n}$, the

$$\begin{aligned}
 y_n &= H(\omega) x_n \\
 &= |H(\omega)| e^{2\pi i \phi(\omega)} x_n \\
 &= |H(\omega)| e^{2\pi i (\omega n + \phi(\omega))}
 \end{aligned}$$

which is just a complex sinusoid of the same frequency, but with amplitude $|H(\omega)|$, and shifted (to the left) by phase $\phi(\omega)$. Maybe it is easier to look at the, say, imaginary parts, with

$$\begin{aligned}
 \text{Im}(x_n) &= \sin(2\pi \omega n) \\
 \text{Im}(y_n) &= |H(\omega)| \sin(2\pi (\omega n + \phi(\omega)))
 \end{aligned}$$

So we see that the output is just a sine wave, but with amplitude $|H(\omega)|$ and phase shifted by $-\phi(\omega)$.

Think about what ϕ really measures. It is the shift, in units of a cycle at that frequency. So $\phi = 0.5$ corresponds to a shift by exactly half a cycle, at that frequency. Half a cycle, of course, changes a sine wave to a negative sine wave. A shift of .25 changes a sine wave to a cosine wave.

A system that delays by exactly n samples (eg $\mathbf{h} = \delta^n$), has a linear phase delay, $\phi(\omega) = -n\omega$. Waves at different frequencies get shifted by different number of cycles – which makes sense, since for higher frequencies, we have shorter cycles, so we must shift by more of them to get the delay of




Figure 6.7: Input/output response with the amplitude and phase effect.

n samples. Linear phase delay is often a desirable characteristic, as it does not change the shape of a given waveform — it just delays it.

In general, when one is designing filters (systems), we have to be concerned with specifying both amplitude and phase characteristics. At this point in the course, I will be happy if you just realize that phase delays happen, we can measure them, and that the phase of $H(\omega)$ identifies the phase delay at frequency ω .

Chapter 7

Interlude: Some practical examples

Here are a few concrete examples using the material learned so far.

7.1 MacRecorder

In 1985, your prof invented the MacRecorder, a digital microphone for Macintosh computers. The Mac, and the Macintosh, had an unusual sampling rate of 22257 samples per second, (This was tied to the video refresh rate of about 60.15 Hz, with effectively 370 lines per screen, where $6.15 * 370 = 22257$.)

What's the highest frequency of sound that this device could record?

Answer: it is half the sampling rate, or 11128 Hz. In practice, it was more like 10kHz.

Later versions of the Macintosh played sounds at a sampling rate of 22050 samples per second. (This was exactly half the sampling rate of CDs, which made it convenient to copy data from a CD to a Mac, and back, although at lower quality of sound.)

At this new rate, how would a musical note recorded in the old rate sound?

Answer: because the sampling rate is lower, the musical note played out would sound lower (i.e. it would sound flat.) More precisely, if we took the note "A above middle C" which is 440Hz, it would play out at the frequency

$$freq = 440 * \frac{22050}{22257} = 435.9Hz. \quad (7.1)$$

This is about one percent lower in frequency, which you can definitely hear. To go down a semitone (the next note on the scale, is to go down by about six percent. So this is quite a bit less than a full semitone.

How do you get rid of the high frequencies before sampling, to avoid aliasing?

Answer: I used a simple RC filter, which is a single resistor at the input, with a capacitor to ground. It acts like the shocks on a car, reducing the high frequency oscillations while passing the low frequency oscillations. The attenuation factor at frequency ω (in Hz) is given by

$$A = \frac{1}{\sqrt{(2\pi RC\omega)^2 + 1}}. \quad (7.2)$$

To attenuate by a factor of 1/2 at frequency $\omega = 10000$, we need

$$(2\pi RC\omega)^2 + 1 = 4 \quad (7.3)$$

or more simply

$$RC = \frac{\sqrt{3}}{2\pi * 10000} = 2.7 \times 10^{-5}. \quad (7.4)$$

A capacitor with $C = .1\mu F$ (microfarads) and $R = 270\Omega$ (ohms) will do the trick.

This is not a perfect filter. A higher order filter that removes more of the higher frequencies, while leaving the lower frequencies as is, can be designed using more electronic parts.

7.2 A 60Hz noise filter

Suppose you have seismic data sampled at 500 samples per second, which is contaminated with 60Hz noise, that was induced from a nearby power line. Design a simple filter to remove this noise.

Answer: Let's design a simple filter with a zero at 60Hz to eliminate the noise. The 60Hz frequency corresponds to the normalized frequency parameter $\omega_0 = 60/500 = .12$. The zero in the complex plane, on the unit circle, is at $Z_0 = e^{-2\pi i \omega_0} = e^{-.24\pi i}$. We can insert a zero at the negative frequency as well, at $\overline{Z_0}$, in order to get a filter with real coefficients.

Our filter then has two zeros, Z_0 and $\overline{Z_0}$, with Z-transform

$$H(Z) = (Z - Z_0)(Z - \overline{Z_0}) = Z^2 - (Z_0 + \overline{Z_0})Z + Z_0\overline{Z_0}. \quad (7.5)$$

Since Z_0 is on the unit circle, we have $Z_0\overline{Z_0} = 1$, while the middle term is $(Z_0 + \overline{Z_0}) = 2\cos(.24\pi) = 1.4579$. Thus the Z-transform can be written as

$$H(Z) = Z^2 - 1.4579 * Z + 1 \quad (7.6)$$

and the corresponding FIR filter has three terms, given by

$$y_n = x_n - 1.4579 * x_{n-1} + x_{n-2}. \quad (7.7)$$

Note: this is not a very good filter, which you can verify by plotting the frequency response. But it does do the job of eliminating the 60Hz noise. Later in the course we will talk about making better filters.

Chapter 8

Simple Filters

The systems we design (FIR, IIR) are specified by Z-transforms which are rational functions – one polynomial divided by another. From the Fundamental Theorem of Algebra, we know these polynomials can be factored as a product of terms of the form $(Z - Z_*)$, where Z_* is the root of the polynomial. That is,

$$H(Z) = \frac{\prod_{j=1} (Z - Z_j)}{\prod_{k=1} (Z - Z_k)}. \quad (8.1)$$

Many terms means we have a complicated system. The simplest filters, then, have a single factor on top, a single factor on bottom, or one each above and below.

8.1 A single zero

An FIR filter with a single zero has a Z-transform of the form

$$H(Z) = c(Z - Z_*), \quad (8.2)$$

where Z_* is some fixed point on the complex plane. The frequency response is given as

$$H(\omega) = c(e^{-2\pi i \omega} - Z_*), \quad (8.3)$$

and the amplitude and phase response are given by the amplitude and phase of $H(\omega)$.

Rather than computing these algebraically, let's just look at a few examples.

For $H(Z) = Z - 1$, we have a zero at $Z = 1$, which corresponds to a zero at the frequency $\omega = 0$. For $H(Z) = Z - i$, we have a zero at $Z = i$, which corresponds to a zero at the frequency $\omega = -0.25$. We plot the amplitude response in Figure 8.1, and see a nice smooth function that zeros out at the given frequency.

This would be a good filter for removing one specific frequency. Note we get a zero in the amplitude response when the root Z_* is on the unit circle in the complex plane.

If we move the zero off the complex plane, we don't get a zero, but a small notch at the corresponding frequency. Figure 8.2 compares the frequency response of the filter with $H(Z) = Z - 1$, and the filter $H(Z) = (2/3) * (Z - 2)$. The first has a nice zero at $\omega = 0$ whereas the second has a notch at $\omega = 0$, but not a full zero.

In general, if the zero is at $Z_* = \rho e^{2\pi i \omega_0}$, then there will be a notch at frequency $\omega = -\omega_0$, when $\rho \neq 1$. The closer that ρ is to one, the closer the zero is to the unit circle, and so the notch gets sharper and sharper. Until we get $\rho = 1$, and then the notch turns into a zero in the amplitude response.




Figure 8.1: Amplitude response for a filter with zero at $Z = 1$ ($\omega = 0$) and at $Z = i$ ($\omega = -0.25$)




Figure 8.2: Amplitude response for a filter with zero at $Z = 1$ and at $Z = 2$

We should note that a filter with a zero at $Z = 2$ has the same amplitude response as one with a zero at $Z = 1/2$. I could plot this, but you get exactly the same picture as the right side of Figure 8.2. What is different is the phase response. Figure 8.3 shows the phase response for the first, $H(Z) = Z - 2$, and for the second, $H(Z) = 2Z - 1$. Look how small the phase response is for the first one. That's why it is called minimum phase. Look how big the second one is. That is why it is called maximum phase.

Notice the correspondence: Min phase has the zero outside the unit circle (i.e. $Z = 2$), while Max phase has the zero inside the unit circle ($Z = 1/2$).




Figure 8.3: Phase response for a filter with zero at $Z = 2$ (min phase), and at $Z = 1/2$ (max phase).

You should check, algebraically, that the filter $H(Z) = Z - Z_*$ has the same amplitude response as the filter $H(Z) = \overline{Z_*}Z - 1$. But of course the phase response is different.

8.2 A single pole

An FIR filter with a single pole has a Z-transform of the form

$$H(Z) = \frac{c}{Z - Z_*}, \quad (8.4)$$

where Z_* is some fixed point on the complex plane. The frequency response is given as

$$H(\omega) = \frac{c}{e^{-2\pi i \omega} - Z_*}, \quad (8.5)$$

and the amplitude and phase response are given by the amplitude and phase of $H(\omega)$.

Again, rather than computing these algebraically, let's just look at a few examples.

For $H(Z) = 1/(Z - 1)$, we have a pole at $Z = 1$, which corresponds to a big spike at the frequency $\omega = 0$. For $H(Z) = i$, we have a pole at $Z = i$, which corresponds to a spike at the frequency $\omega = -0.25$. We plot the amplitude response in Figure 8.4, and see a nice smooth function that spikes at the given frequency.

This would be a good filter for amplifying one specific frequency. Note we get an infinite in the amplitude response when the root Z_* is on the unit circle in the complex plane.




Figure 8.4: Amplitude response for a filter with pole at $Z = 1$ ($\omega = 0$) and at $Z = -i$ ($\omega = .25$)



If we move the zero off the complex plane, we don't get an infinite spike, but a smaller spike at the corresponding frequency. Figure 8.5 compares the frequency response of the filter with $H(Z) = 1/(Z - 1)$, and the filter $H(Z) = 1/(Z - 2)$. The first has a big spike at $\omega = 0$ whereas the second has a only a large hump at $\omega = 0$, but not a full spike.

In general, if the pole is at $Z_* = \rho e^{2\pi i \omega_0}$, then there will be a large hump at frequency $\omega = -\omega_0$, when $\rho \neq 1$. The closer that ρ is to one, the closer the pole is to the unit circle, and so the hump gets sharper and sharper. Until we get $\rho = 1$, and then the hump turns into an infinite spike in the amplitude response.

We should note that a filter with a pole at $Z = 2$ has the same amplitude response as one with a pole at $Z = 1/2$. I could plot this, but you get exactly the same picture as the right side of Figure 8.5. What is different is the phase response. Figure 8.6 shows the phase response for the first, $H(Z) = 1/(Z - 2)$, and for the second, $H(Z) = 1/(2Z - 1)$. Look how small the phase response is for the first one. That's why it is called minimum phase. Look how big the second one is. That is why it is called maximum phase.

Notice the correspondence: Min phase has the pole outside the unit circle (i.e. $Z = 2$), while Max phase has the pole inside the unit circle ($Z = 1/2$).

As in the last section, the filter $H(Z) = 1/(Z - Z_*)$ has the same amplitude response as the filter $H(Z) = 1/(\overline{Z_*}Z - 1)$. But of course the phase response is different.

Figure 8.5: Amplitude response for a filter with pole at $Z = 1$ and at $Z = 2$ Figure 8.6: Phase response for a filter with pole at $Z = 2$ (min phase), and at $Z = 1/2$ (max phase).

8.3 All-pass: matched zero/pole filters

By taking a filter of the form

$$H(Z) = \frac{Z - Z_*}{1 - \overline{Z_*}Z} \quad (8.6)$$

you get a single zero at Z_* and a single pole at $1/\overline{Z_*}$. For instance, you could put a pole at $Z = 2i$ and a zero at $i/2$, with a filter of the form

$$H(Z) = \frac{2Z - i}{Z - 2i}. \quad (8.7)$$

By matching the pole and zero this way, we get a filter with a flat amplitude response. Such a filter is called an all-pass filters, since it passes all frequencies with equal amplitude. Why would we want this? In order to get a non-constant phase response, despite the flat amplitude response.

Figure 8.7 shows the amplitude and phase response for a filter with a zero at $Z = 1/2$ and a pole at $Z = 2$. Note the amplitude response is completely flat. The phase response is doing something interesting; the big change at $\omega = 0$ corresponds to the location of the zero-pole pair being match across the unit circle at $Z = 1$, where $\omega = 0$.

8.4 Stability

Note again, a filter with poles inside the unit circle will be unstable.




Figure 8.7: Amplitude and phase response for a filter with matched zero and pole pair.

We can see this easily in the single pole case, with

$$H(Z) = \frac{a}{Z - a}. \quad (8.8)$$

This corresponds to a recursive (IIR) system of the form

$$y_n = x_n + \frac{1}{a}y_{n-1}. \quad (8.9)$$

The impulse response of this system is

$$(1, \frac{1}{a}, \frac{1}{a^2}, \frac{1}{a^3}, \frac{1}{a^4}, \dots), \quad (8.10)$$

which grows without bound if $\frac{1}{|a|} > 1$. That is $|a| < 1$, so a single pole inside the unit circle gives an unstable system.

Chapter 9

Fourier transforms I

It turns out there are Fourier transforms defined in all sorts of situations. We start with the two we need immediately: FT on sampled signals (sequences) and on periodic functions on the real line.

9.1 Defining the FT on sampled signals and on 1-periodic functions

The Fourier transform of a sequence $\mathbf{x} = (\dots, x_{-1}, x_0, x_1, \dots)$ is given by the formula

$$\hat{x}(\omega) = \sum_{n=-\infty}^{\infty} x_n e^{-2\pi i n \omega}, \quad \omega \in \mathbb{R}, \quad (9.1)$$

where $\hat{x}(\omega)$ is a 1-periodic function of the real parameter ω . We've seen this formula before, as the frequency response of a LSI system with impulse response \mathbf{x} . Now we give it its proper name, the Fourier transform.

We can also define the Fourier transform of a 1-periodic function $f(t)$ via the formula

$$\hat{f}(n) = \int_{-1/2}^{1/2} f(\omega) e^{-2\pi i n t} dt, \quad n \in \mathbb{Z} \quad (9.2)$$

where $\hat{f}(n)$ is a function on the integers $n \in \mathbb{Z}$.

What you should notice is that the FT of a sequence is a 1-periodic function on the real line, and the FT of a 1-periodic function is a sequence. So you use the FT to go back and forth between sequences and functions.

9.2 Examples of FTs, on sequences

Here are some simple examples.

1. Take the sequence $\mathbf{x} = \delta^0$; that is, the sequence with $x_0 = 1$, and all the other coefficients are zero. Then

$$\hat{\mathbf{x}}(\omega) = \sum_n \mathbf{x}_n e^{-2\pi i n \omega} = x_0 e^{-2\pi i 0 \omega} = 1. \quad (9.3)$$

That is, the Fourier transform of the delta function δ^0 is the constant function 1.

2. Take the sequence $\mathbf{x} = \delta^3$; that is, the sequence with $x_3 = 1$, and all the other coefficients are zero. Then

$$\widehat{\mathbf{x}}(\omega) = \sum_n \mathbf{x}_n e^{-2\pi i n \omega} = x_3 e^{-2\pi i 3 \omega} = e^{-6\pi i \omega}. \quad (9.4)$$

That is, the Fourier transform of the delta function δ^3 is the complex exponential function $\omega \mapsto e^{-6\pi i \omega}$, which has exactly 3 cycles in the interval $[0, 1]$.

3. Take the sequence $\mathbf{x} = \delta^{-3} + \delta^3$; that is, the sequence with $x_{-3} = 1$, $x_3 = 1$, and all the other coefficients are zero. Then

$$\widehat{\mathbf{x}}(\omega) = e^{2\pi i 3 \omega} + e^{-2\pi i 3 \omega} = 2 \cos 6\pi \omega. \quad (9.5)$$

That is, the Fourier transform of the sum of these two delta functions is a cosine, $\omega \mapsto 2 \cos 6\pi \omega$, which has exactly 3 cycles in the interval $[0, 1]$.

4. Take the sequence $\mathbf{x} = \delta^{-3} - \delta^3$; that is, the sequence with $x_{-3} = 1$, $x_3 = -1$, and all the other coefficients are zero. Then

$$\widehat{\mathbf{x}}(\omega) = e^{2\pi i 3 \omega} - e^{-2\pi i 3 \omega} = 2i \sin 6\pi \omega. \quad (9.6)$$

That is, the Fourier transform of the difference of these two delta functions is a sine, $\omega \mapsto 2i \sin 6\pi \omega$, which has exactly 3 cycles in the interval $[0, 1]$.

5. Take the sequence \mathbf{x} which is one for all x_n in the range $-N \leq n \leq N$. Then

$$\widehat{\mathbf{x}}(\omega) = \sum_{n=-N}^N e^{2\pi i n \omega} = \frac{e^{-2\pi i N \omega} - e^{2\pi i (N+1) \omega}}{1 - e^{2\pi i \omega}} = \frac{\sin(\pi(2N+1)\omega)}{\sin(\pi\omega)}. \quad (9.7)$$

That is, this Fourier transform is a ratio of two sines. It is interesting enough that we should plot it, in Figure 9.1.




Figure 9.1: The Fourier transform of a long stretch of ones. ($N = 10$)

9.3 Examples of FTs, on 1-periodic functions

1. Take $f(\omega) = 1$ the constant one function. The FT is the sequence with

$$\hat{f}(n) = \int_{-1/2}^{1/2} f(t)e^{-2\pi int} dt = \int_{-1/2}^{1/2} e^{-2\pi int} dt = 0, \text{ for } n \neq 0, \quad (9.8)$$

and is 1 when $n = 0$. Thus $\hat{f} = \delta^0$, the delta function supported at zero.

2. Take $f(\omega) = e^{2\pi i 3\omega}$ the complex exponential with 3 cycles in the unit interval. The FT is the sequence with

$$\hat{f}(n) = \int_{-1/2}^{1/2} f(t)e^{-2\pi int} dt = \int_{-1/2}^{1/2} e^{2\pi i(3-n)t} dt = 0, \text{ for } n \neq 3, \quad (9.9)$$

and is 1 when $n = 3$. Thus $\hat{f} = \delta^3$, the delta function supported at three.

3. Take $f(\omega) = 2 \cos 6\pi\omega = e^{2\pi i 3\omega} + e^{-2\pi i 3\omega}$ the cosine with 3 cycles in the unit interval. The FT is the sequence with

$$\hat{f}(n) = \int_{-1/2}^{1/2} f(t)e^{-2\pi int} dt = \int_{-1/2}^{1/2} e^{2\pi i(3-n)t} + e^{2\pi i(3+n)t} dt = 0, \text{ for } n \neq 3, -3, \quad (9.10)$$

and is 1 when $n = 3$ or -3 . Thus $\hat{f} = \delta^{-3} + \delta^3$, the sum of two delta functions.

4. Take $f(\omega) = 2i \sin 6\pi\omega = e^{2\pi i 3\omega} - e^{-2\pi i 3\omega}$ the sine with 3 cycles in the unit interval. The FT is the sequence with

$$\hat{f}(n) = \int_{-1/2}^{1/2} f(t)e^{-2\pi int} dt = \int_{-1/2}^{1/2} e^{2\pi i(3-n)t} - e^{2\pi i(3+n)t} dt = 0, \text{ for } n \neq 3, -3, \quad (9.11)$$

and is 1 when $n = 3$, and is -1 when $n = -3$. Thus $\hat{f} = \delta^3 - \delta^{-3}$, the difference of two delta functions. You might notice this is almost like the difference of delta functions in example 4 of the last section, except the difference is reversed.

5. Take $f(\omega) = 1$ on the interval $[-a, a]$ and zero elsewhere. (This is known as a boxcar, or a step function.) The FT is the sequence

$$\hat{f}(n) = \int_{-1/2}^{1/2} f(t)e^{-2\pi int} dt = \int_{-a}^a f(t)e^{-2\pi int} dt = \frac{e^{2\pi ina} - e^{-2\pi ina}}{2\pi in} = \frac{\sin 2\pi na}{\pi n}. \quad (9.12)$$

This is also worth plotting. It is the set of uniform samples from the sinc function, as seen in Figure 9.2.

9.4 Properties of the Fourier transform

Here are some important properties of the Fourier transform on sequences:

1. $\widehat{\hat{x}}(n) = x(-n)$.
2. The FT has an inverse, given by a similar formula.
3. $\widehat{\mathbf{x} * \mathbf{y}} = \hat{\mathbf{x}} \cdot \hat{\mathbf{y}}$.




Figure 9.2: The Fourier transform of a boxcar, supported on interval $[-1/3, 1/3]$. The FT is a time series sequence, that is a set of uniform samples of the sinc function.

$$4. \widehat{\mathbf{x} \cdot \mathbf{y}} = \widehat{\mathbf{x}} * \widehat{\mathbf{y}}.$$

$$5. \|\mathbf{x}\|_2 = \|\widehat{\mathbf{x}}\|_2.$$

$$6. \langle \mathbf{x}, \mathbf{y} \rangle = \langle \widehat{\mathbf{x}}, \widehat{\mathbf{y}} \rangle.$$

In words, we can summarize as the following:

1. Repeating the FT twice returns the original signal, but in reverse order.
2. The inverse FT is computed using a sum/integral with exponentials, but with a $+$ in the exponential argument.
3. The FT maps a convolution to a product.
4. The FT maps a product to a convolution.
5. The FT preserves energy (or length).
6. The FT preserves inner products.

9.5 Some worked examples

1. Let $\mathbf{x} = (0, 1, 2, 3)$, that is, the sequence with $x_1 = 1, x_2 = 2, x_3 = 3$, and all other coefficients zero. Then we compute the first Fourier transform

$$\widehat{\mathbf{x}}(\omega) = \sum_n x_n e^{-2\pi i n \omega} = e^{-2\pi i \omega} + 2e^{-4\pi i \omega} + 3e^{-6\pi i \omega}, \quad (9.13)$$

which is a 1-periodic function, and then compute the second Fourier transform as

$$\widehat{\widehat{\mathbf{x}}}(n) = \int_{-1/2}^{1/2} \widehat{\mathbf{x}}(\omega) e^{-2\pi i n \omega} d\omega \quad (9.14)$$

$$= \int_{-1/2}^{1/2} (e^{-2\pi i \omega} + 2e^{-4\pi i \omega} + 3e^{-6\pi i \omega}) e^{-2\pi i n \omega} d\omega \quad (9.15)$$

$$= \int_{-1/2}^{1/2} (e^{-2\pi i(1+n)\omega} + 2e^{-2\pi i(2+n)\omega} + 3e^{-2\pi i(3+n)\omega}) d\omega, \quad (9.16)$$

which gives a non-zero result only when one of the factors $n+1, n+2, n+3$ is zero. Thus we find that $\widehat{\widehat{\mathbf{x}}}(-1) = 1, \widehat{\widehat{\mathbf{x}}}(-2) = 2, \widehat{\widehat{\mathbf{x}}}(-3) = 3$, which is exactly the reverse of the coefficients in the original sequence \mathbf{x} . This example verifies property 1.

2. We can fix the problem of reverse coefficients in the last example by defined the inverse Fourier transform using a $+$ in the exponential. For instance, we define the inverse Fourier transform of a 1-periodic function $f(\omega)$ as

$$\check{f}(n) = \int_{-1/2}^{1/2} f(\omega) e^{+2\pi i n \omega} d\omega, \quad (9.17)$$

where we see in the exponential we have a plus sign rather than a negative sign. This way, when we apply the inverse transform to the function $\widehat{\mathbf{x}}$ of the last example, we have

$$\check{\widehat{\mathbf{x}}}(n) = \int_{-1/2}^{1/2} \widehat{\mathbf{x}}(\omega) e^{2\pi i n \omega} d\omega \quad (9.18)$$

$$= \int_{-1/2}^{1/2} (e^{-2\pi i \omega} + 2e^{-4\pi i \omega} + 3e^{-6\pi i \omega}) e^{2\pi i n \omega} d\omega \quad (9.19)$$

$$= \int_{-1/2}^{1/2} (e^{-2\pi i(1-n)\omega} + 2e^{-2\pi i(2-n)\omega} + 3e^{-2\pi i(3-n)\omega}) d\omega, \quad (9.20)$$

which is zero except when one of the factors $1-n, 2-n, 3-n$ is zero. Thus we find that $\check{\widehat{\mathbf{x}}}(1) = 1 = x_1, \check{\widehat{\mathbf{x}}}(2) = 2 = x_2, \check{\widehat{\mathbf{x}}}(3) = 3 = x_3$. In other words, $\check{\widehat{\mathbf{x}}} = \mathbf{x}$, so we have recovered the original signal. This example verifies property 2.

3. Let's do a simple convolution example. Take $\mathbf{x} = (1, 2), \mathbf{y} = (1, 3)$ and so we easily compute the convolution $\mathbf{x} * \mathbf{y} = (1, 5, 6)$. The Fourier transforms are

$$\widehat{\mathbf{x}}(\omega) = 1 + 2e^{-2\pi i \omega} \quad (9.21)$$

$$\widehat{\mathbf{y}}(\omega) = 1 + 3e^{-2\pi i \omega} \quad (9.22)$$

$$\widehat{\mathbf{x} * \mathbf{y}}(\omega) = 1 + 5e^{-2\pi i \omega} + 6e^{-4\pi i \omega} \quad (9.23)$$

and it is easy to verify that the product of the first two functions equals the third, as

$$(1 + 2e^{-2\pi i \omega})(1 + 3e^{-2\pi i \omega}) = 1 + 5e^{-2\pi i \omega} + 6e^{-4\pi i \omega}. \quad (9.24)$$

This example verifies property 3.

This is not surprising, since we've already seen that the Z-transform of a convolution gives the product of Z-transforms; we just included the substitution $Z = e^{-2\pi i \omega}$.

4. Let's repeat the above example, using a pointwise product of sequences. Take $\mathbf{x} = (1, 2), \mathbf{y} = (1, 3)$ and the pointwise product is $\mathbf{x} \cdot \mathbf{y} = (1, 6)$. The Fourier transforms are

$$\widehat{\mathbf{x}}(\omega) = 1 + 2e^{-2\pi i \omega} \quad (9.25)$$

$$\widehat{\mathbf{y}}(\omega) = 1 + 3e^{-2\pi i \omega} \quad (9.26)$$

$$\widehat{\mathbf{x} \cdot \mathbf{y}}(\omega) = 1 + 6e^{-2\pi i \omega}. \quad (9.27)$$

The convolution for 1-periodic functions is defined in the obvious way, so

$$(\widehat{\mathbf{x}} * \widehat{\mathbf{y}})(\omega) = \int \widehat{\mathbf{x}}(\omega - s) \widehat{\mathbf{y}}(s) ds \quad (9.28)$$

$$= \int \left(1 + 2e^{-2\pi i(\omega-s)}\right) (1 + 3e^{-2\pi i s}) ds \quad (9.29)$$

$$= \int 1 + 2e^{-2\pi i \omega} e^{2\pi i s} + 3e^{-2\pi i s} + 6e^{-2\pi i \omega} ds \quad (9.30)$$

$$= 1 + 0 + 0 + 6e^{-2\pi i \omega} \quad (9.31)$$

$$= \widehat{\mathbf{x} \cdot \mathbf{y}}(\omega), \quad (9.32)$$

where we see in the last integral that the terms with an exponential depending on s just integrate to zero, and the other ones give the part depending on ω only. So we have that the FT of the pointwise product is the convolution of the two FTs. This verifies property 4.

5. Let's take the vector $\mathbf{x} = (2, 3, 6)$, which has length $\|\mathbf{x}\| = \sqrt{2^2 + 3^2 + 6^2} = \sqrt{49} = 7$. The Fourier transform of this signal is the 1-periodic function

$$\widehat{\mathbf{x}} = 2 + 3e^{-2\pi i \omega} + 6e^{-4\pi i \omega}. \quad (9.33)$$

To get the “length” of this function, we start by computing its energy, which is the integral of the square of the function, so

$$\text{Energy} = \int |\widehat{\mathbf{x}}(\omega)|^2 d\omega = \int \widehat{\mathbf{x}}(\omega) \overline{\widehat{\mathbf{x}}(\omega)} d\omega \quad (9.34)$$

$$= \int (2 + 3e^{-2\pi i \omega} + 6e^{-4\pi i \omega}) (2 + 3e^{2\pi i \omega} + 6e^{4\pi i \omega}) d\omega \quad (9.35)$$

$$= \int 4 + 9 + 36 + \text{cross terms involving } e^{2\pi i n \omega} d\omega \quad (9.36)$$

$$= 49, \quad (9.37)$$

where the cross terms all integrate to zero, since we know those complex exponentials all integrate to zero. So the energy is 49, and the square root is the “length” of the function, which is 7. So it matches the length of the initial vectors $\mathbf{x} = (2, 3, 6)$. This verifies property 5, that the FT preserves length.

6. Let's take vectors $\mathbf{x} = (2, 3, 6)$ and $\mathbf{y} = (3, 2, 1)$, with Fourier transforms

$$\widehat{\mathbf{x}}(\omega) = 2 + 3e^{-2\pi i \omega} + 6e^{-4\pi i \omega} \quad (9.38)$$

$$\widehat{\mathbf{y}}(\omega) = 3 + 2e^{-2\pi i \omega} + 1e^{-4\pi i \omega}. \quad (9.39)$$

The inner product (dot product) of the two vectors is just $\langle \mathbf{x}, \mathbf{y} \rangle = 2 * 3 + 3 * 2 + 6 * 1 = 18$, while the inner product of the two functions is

$$\langle \widehat{\mathbf{x}}, \widehat{\mathbf{y}} \rangle = \int \widehat{\mathbf{x}}(\omega) \overline{\widehat{\mathbf{y}}(\omega)} d\omega \quad (9.40)$$

$$= \int (2 + 3e^{-2\pi i \omega} + 6e^{-4\pi i \omega}) (3 + 2e^{2\pi i \omega} + 1e^{4\pi i \omega}) d\omega \quad (9.41)$$

$$= \int 6 + 6 + 6 + \text{cross terms involving } e^{2\pi i n \omega} d\omega \quad (9.42)$$

$$= 18, \quad (9.43)$$

since the cross terms all integrate to zero. So this inner product is 18, which matches the inner product of the two vectors. This example verifies property 6, that the FT preserves inner products.

9.6 Convolution, etc, on 1-periodic functions

In the examples in the last section, we stumbled across some operations on 1-periodic functions which have not been defined yet. So we define them here.

Convolution of two functions $f(\omega)$ and $g(\omega)$ is defined as the integral

$$(f * g)(\omega) = \int_{-1/2}^{1/2} f(\omega - s)g(s) ds. \quad (9.44)$$

Note the similarity to our definition of convolution of vectors,

$$(\mathbf{x} * \mathbf{y})_n = \sum_k \mathbf{x}_{n-k} \mathbf{y}_k. \quad (9.45)$$

The key difference is that the sum is replaced by an integral, and rather than summing over variable k , we integrate over variable s .

Because f, g are 1-periodic, the result of the convolution is also a 1-periodic function. It is interesting to note that the interval $[-1/2, 1/2]$ can be replaced with any interval of length one, and the result is the same. For instance, in many books, you will see the defining integral on the interval $[0, 1]$. This gives exactly the same result. It is also worth noting that $f * g = g * f$, so the defining integral can also be written as

$$(f * g)(\omega) = \int_{-1/2}^{1/2} f(s)g(\omega - s) ds. \quad (9.46)$$

The inner product of two functions $f(\omega)$ and $g(\omega)$ is defined as the integral

$$\langle f, g \rangle = \int_{-1/2}^{1/2} f(\omega) \overline{g(\omega)} d\omega. \quad (9.47)$$

Note the complex conjugate indicated on the function g ; this makes the inner product anti-linear in the second component, and linear in the first component. It is analogous to the inner product of vectors, given by

$$\langle \mathbf{x}, \mathbf{y} \rangle = \sum_n x_n \overline{y_n}. \quad (9.48)$$

Remember, the inner product of two vectors is a *number* which measures the angle between the two vectors. Of course, here we are measuring angles in an infinite dimensional space, which you may find hard to visualize.

The energy of a function is defined using the inner product, and hence is also given by an integral. We define it as

$$\text{Energy} = \langle f, f \rangle = \int_{-1/2}^{1/2} |f(\omega)|^2 d\omega. \quad (9.49)$$

The “length” or L2 norm of the function is defined as the square root of the energy, so

$$\|f\| = \sqrt{\langle f, f \rangle} = \sqrt{\int_{-1/2}^{1/2} |f(\omega)|^2 d\omega}. \quad (9.50)$$

Again, notice that the norm of a function is a number (not a function), and roughly speaking it is a measure of how far away from zero is the function.

9.7 Verifying the six properties

The above examples show that it is plausible that the six properties of the FT, as listed in Section 9.4, hold. We really should show this in general. However, let's not go crazy! We've just verify a few.

Property 1 says if we repeat the FT twice, we get the same signal back, reversed. To check, start with signal \mathbf{x} . It's FT is

$$\hat{\mathbf{x}}(\omega) = \sum_n x_n e^{-2\pi i n \omega}. \quad (9.51)$$

Taking the FT a second time, we get an integral form

$$\hat{\hat{\mathbf{x}}}(m) = \int \hat{\mathbf{x}}(\omega) e^{-2\pi i m \omega} d\omega \quad (9.52)$$

$$= \int \sum_n x_n e^{-2\pi i n \omega} e^{-2\pi i m \omega} d\omega, \text{ now reverse order of int, sum} \quad (9.53)$$

$$= \sum_n x_n \int e^{-2\pi i (n+m)\omega} d\omega \quad (9.54)$$

and we notice the integral is always zero, except when $n + m = 0$, in which case we get a one. So the sum drops out to a single nonzero term, at $n = -m$, and we have

$$\hat{\hat{\mathbf{x}}}(m) = \mathbf{x}_{-m}, \quad (9.55)$$

which is exactly what property 1 says. That is, we recover the original signal, but time-reversed.

Property 3 says if we take the FT of a convolution, we get the product of two FTs. Let's check, using sequences \mathbf{x}, \mathbf{y} and $\mathbf{z} = \mathbf{x} * \mathbf{y}$. Then the FT of the convolution is

$$\hat{\mathbf{z}}(\omega) = \sum_n z_n e^{-2\pi i n \omega} \quad (9.56)$$

$$= \sum_n \left(\sum_k x_{n-k} y_k \right) e^{-2\pi i n \omega}, \text{ then split up the exponential} \quad (9.57)$$

$$= \sum_{n,k} x_{n-k} e^{-2\pi i (n-k)\omega} y_k e^{-2\pi i k \omega}, \text{ then change variables on the sums} \quad (9.58)$$

$$= \sum_{n',k} x_{n'} e^{-2\pi i (n')\omega} y_k e^{-2\pi i k \omega}, \quad (9.59)$$

$$= \hat{\mathbf{x}}(\omega) \hat{\mathbf{y}}(\omega), \quad (9.60)$$

which is the product of the two FTs, as desired.

Property 4 says the FT of a pointwise product gives the convolution of two FTs. Again we check, using \mathbf{x}, \mathbf{y} and $\mathbf{z} = \mathbf{x} \cdot \mathbf{y}$. Then the FT of the product is

$$\hat{\mathbf{z}}(\omega) = \sum_n z_n e^{-2\pi i n \omega} \quad (9.61)$$

$$= \sum_n x_n y_n e^{-2\pi i n \omega}, \text{ then express } x_n \text{ using the inverse FT} \quad (9.62)$$

$$= \sum_n \left(\int \hat{\mathbf{x}}(t) e^{2\pi i n t} dt \right) y_n e^{-2\pi i n \omega}, \text{ then switch the sum, int} \quad (9.63)$$

$$= \int \hat{\mathbf{x}}(t) \sum_n y_n e^{-2\pi i n(\omega - t)}, \quad (9.64)$$

$$= \int \hat{\mathbf{x}}(t) \hat{\mathbf{y}}(\omega - t), \quad (9.65)$$

which is the convolution of the two FTs as 1-periodic functions, which is as desired.

Property 6 says that the FT preserves inner products. Again to check, we use two sequences \mathbf{x} and \mathbf{y} . Then the inner product of their FTs is given by

$$\langle \hat{\mathbf{x}}, \hat{\mathbf{y}} \rangle = \int \hat{\mathbf{x}}(\omega) \overline{\hat{\mathbf{y}}(\omega)} d\omega \quad (9.66)$$

$$= \int \sum_n x_n e^{-2\pi i n\omega} \overline{\sum_m y_m e^{-2\pi i m\omega}} d\omega, \quad (9.67)$$

$$= \int \sum_{m,n} x_n \overline{y_m} e^{-2\pi i n\omega} \overline{e^{-2\pi i m\omega}} d\omega, \text{ then take the sums out} \quad (9.68)$$

$$= \sum_{m,n} x_n \overline{y_m} \int e^{-2\pi i(n-m)\omega} d\omega \quad (9.69)$$

and we notice this integral is zero, except when $n = m$, in which case we get a one. So the double sum drops down to a single sum (with $n = m$) and we have

$$\langle \hat{\mathbf{x}}, \hat{\mathbf{y}} \rangle = \sum_n x_n \overline{y_n} \quad (9.70)$$

$$= \langle \mathbf{x}, \mathbf{y} \rangle, \quad (9.71)$$

which is the inner product of the two sequences.

Property 5 follows immediately from property 6, since the length (or norms) are defined using the inner product.

9.8 What does the FT mean?

The Fourier transform of a signal \mathbf{x} tells you the frequency content of that signal. That is, we can think of the signal \mathbf{x} as being made up of a sum of sines and cosines at various frequencies. By taking the Fourier transform, we can tell what frequencies are present, and at what amplitude.

Specifically, if value of $\hat{\mathbf{x}}(\omega)$ is large at some frequency $\omega = \omega_0$, then the signal \mathbf{x} has some sine or cosine at that frequency. For instance, in Figure 9.3, we have a short signal that is the sum of two sine waves, $f(t) = \sin(2\pi 10t) + .3 * \sin(2\pi 50t)$, sampled at 500 samples per second. The second half of the figure shows the Fourier transform. The peak at $\omega = .02$ indicates the presence of the sine wave with frequency $10Hz = .02 * 500$, and amplitude 1. The peak at $\omega = .1$ indicates the presence of the sine wave with frequency $50Hz = .1 * 500$ and amplitude .3. The peaks at the negative frequencies are just a symmetry property of the FT. Or, you can think of the sine waves as being the sum of two exponentials each, one with a positive frequency, one with a negative frequency.

A more typical signal is shown in Figure 9.4, top, which is a little spike that might be generated by a hammer blow to a metal rod, or another impulsive source. A “boom.” Its Fourier transform is shown in the bottom of the figure, it indicates that this signal has quite a bit of low frequency energy, and not so much high frequency content.




Figure 9.3: Plot of signal made up of sine waves, and the FT showing the frequency content.

What is very interesting about the FT is that we can recover the signal from the frequency content alone. In fact, Figure 9.4 was created by choosing a shape for the Fourier transform (the lower part of the diagram), and then applying the inverse Fourier transform to get the corresponding signal.

9.9 How do we apply the FT?

The FT converts convolutions to pointwise multiplication. Multiplication is much faster than convolution, so we can implement convolution much faster by working in the frequency domain.

As an example, suppose we have a random signal that we want to smooth out. Essentially, this means we have to remove the high frequencies. To do this, we take the FT of the signal, set the amplitudes of the high frequencies to zero, and then convert back to signal domain. We demonstrate this step by step in Figure 9.5. In the first frame, we see a noisy signal. The second frame shows the Fourier transform, which shows the noise spread across many frequencies. The third frame shows the Fourier transform, with the high frequency components set to zero. The fourth frame shows the resulting signal, obtained by inverting the partially zeroed Fourier transform.

More generally, we can decide to modify the frequency content by multiplying the FT of the signal with some fixed function $g(\omega)$. With this technique, we can apply very sophisticated filters to the signal.




Figure 9.4: Plot of spiky signal, and the FT showing the frequency content.




Figure 9.5: Smoothing out a noisy signal by zeroing out the high frequencies.

Chapter 10

Fourier transforms II

The Fourier transform of a function $f(t)$ defined on the real line is given by

$$\widehat{f}(\omega) = \int_{-\infty}^{\infty} f(t) e^{-2\pi i \omega t} dt, \quad \omega \in \mathbb{R}, \quad (10.1)$$

which gives $\widehat{f}(\omega)$ as a function on the real line. The inverse Fourier transform is defined as

$$\check{f}(\omega) = \int_{-\infty}^{\infty} f(t) e^{+2\pi i \omega t} dt, \quad \omega \in \mathbb{R}, \quad (10.2)$$

which is also a function on the real line.

For a wide class of functions, it is true that $\check{\check{f}} = f$; that is, applying the FT followed by the inverse FT gives you back the original function.

We can also define the convolution of two functions on the real line, as

$$(f * g)(t) = \int_{-\infty}^{\infty} f(t-s)g(s) ds. \quad (10.3)$$

The inner product of two such functions is defined as

$$\langle f, g \rangle = \int_{-\infty}^{\infty} f(t) \overline{g(t)} dt. \quad (10.4)$$

The energy of a signal is defined as

$$Energy = \langle f, f \rangle = \int_{-\infty}^{\infty} |f(t)|^2 dt, \quad (10.5)$$

and the length, or L2 norm is defined as the square root of the energy, so

$$\|f\| = \sqrt{\langle f, f \rangle}. \quad (10.6)$$

With these definitions, the six properties of FT again hold:

1. $\widehat{\widehat{f}}(t) = f(-t)$.
2. The FT has an inverse, given by a similar formula.
3. $\widehat{f * g} = \widehat{f} \cdot \widehat{g}$.

4. $\widehat{f \cdot g} = \widehat{f} * \widehat{g}$.
5. $\|f\| = \|\widehat{f}\|$.
6. $\langle f, g \rangle = \langle \widehat{f}, \widehat{g} \rangle$.

The proofs of these properties are very similar to those for the FT on sequences and 1-periodic function. Try them!

You should also try computing the FT of some simple functions. A good one to try is the Gaussian,

$$f(t) = e^{-t^2}. \quad (10.7)$$

Another one to try is the boxcar function, where $f(t) = 1$ on some interval $[-a, a]$ and zero elsewhere. You might look back at the earlier example we did on the 1-periodic boxcar.

To tell you the truth, though, while the Fourier transform on the real line is very useful in theory, in practice we never can compute the FT of a signal on the real line, since we can never measure the real signal $f(t)$ for all times t . So we are left with using the discrete Fourier transform, discussed in the next chapter.

Chapter 11

Fourier transforms III

The discrete Fourier transform is the practical workhorse for all our signal processing needs. It is a special case of the Fourier transform applied to sequences.

Recall that the Fourier transform of a sequence $\mathbf{x} = (\dots, x_{-1}, x_0, x_1, \dots)$ is defined as the function

$$\hat{\mathbf{x}}(\omega) = \sum_{n=-\infty}^{\infty} x_n e^{-2\pi i n \omega}, \text{ for all real } \omega, \quad (11.1)$$

which is a 1-periodic function on the real line. When the signal is a finite sequence (which often happens in practice), of the form $\mathbf{x} = (x_0, x_1, \dots, x_{N-1})$, then the Fourier transform reduces to a finite sum

$$\hat{\mathbf{x}}(\omega) = \sum_{n=0}^{N-1} x_n e^{-2\pi i n \omega}, \text{ for all real } \omega, \quad (11.2)$$

which again is a 1-periodic function on the real line. Now, it turns out that if we sample this function at N uniformly spaced points in the interval, $\omega = k/N, k = 0, 1, \dots, N-1$, we obtain an invertible linear transform from vectors of length N to vectors of length N . That is, we define

$$\hat{\mathbf{x}}_k = \hat{\mathbf{x}}\left(\frac{k}{N}\right) = \sum_{n=0}^{N-1} x_n e^{-2\pi i n k / N}, \quad 0 \leq k \leq N-1, \quad (11.3)$$

as the linear map which takes the sequence $\mathbf{x} = (x_0, x_1, \dots, x_{N-1})$ to the new sequence

$$\hat{\mathbf{x}} = (\hat{\mathbf{x}}_0, \hat{\mathbf{x}}_1, \hat{\mathbf{x}}_2, \dots, \hat{\mathbf{x}}_{N-1}). \quad (11.4)$$

It is convenient to write the Fourier transform in the form

$$\hat{\mathbf{x}}_k = \sum_{n=0}^{N-1} x_n u^{-nk}, \quad 0 \leq k \leq N-1, \quad (11.5)$$

where $u = e^{2\pi i / N}$ is the first N -th root of unity.

11.1 The discrete FT as a matrix

As a simple example, set $N = 4$, and we have the formula

$$\hat{\mathbf{x}}_k = \sum_{n=0}^3 x_n e^{-2\pi i n k / 4}, \quad 0 \leq k \leq 3, \quad (11.6)$$

where of course $e^{-2\pi ink/4} = (e^{\pi i/2})^{-nk} = (i)^{nk}$. So we can rewrite the transform in matrix form, as

$$\begin{bmatrix} \hat{\mathbf{x}}_0 \\ \hat{\mathbf{x}}_1 \\ \hat{\mathbf{x}}_2 \\ \hat{\mathbf{x}}_3 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -i & -1 & i \\ 1 & -1 & 1 & -1 \\ 1 & i & -1 & -i \end{bmatrix} \begin{bmatrix} x_0 \\ x_1 \\ x_2 \\ x_3 \end{bmatrix} \quad (11.7)$$

In general, the n -th order discrete FT can be expressed as a matrix, where the entries are complex numbers of the form u^{-nk} , with $u = e^{2\pi i/N}$ the first N -th root of unity. As in the 4x4 example, we can write the transform in matrix form as

$$\begin{bmatrix} \hat{\mathbf{x}}_0 \\ \hat{\mathbf{x}}_1 \\ \hat{\mathbf{x}}_2 \\ \vdots \\ \hat{\mathbf{x}}_{N-1} \end{bmatrix} = \begin{bmatrix} \dots & u^{-kn} & \dots \end{bmatrix} \begin{bmatrix} x_0 \\ x_1 \\ x_2 \\ \vdots \\ x_{N-1} \end{bmatrix}, \quad (11.8)$$

where we use the u^{-kn} to indicate the matrix entries in the $N \times N$ transform.

So, a quick way to think of the discrete FT is that it is just a special matrix transform on vectors.

11.2 What does the DFT mean?

Recall the DFT is just a special case of the FT applied to sequences, so it simply gives a frequency representation of a signal. We should be careful about noting what the indices mean, however.

For instance, if the particular coefficient $\hat{\mathbf{x}}_k$ is large, this means there is a sinusoid present with normalized frequency $\omega = k/N$. And the normalized frequency means something relative to the sample rate.

As a concrete example, suppose you have a signal $f(t)$ which is sampled at 500 samples per second, and you collect 256 samples into a finite vector $\mathbf{x} = (x_0, x_1, \dots, x_{255})$. You take the discrete Fourier transform to obtain vector $\hat{\mathbf{x}} = (\hat{\mathbf{x}}_0, \hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_{255})$. Now, suppose $\hat{\mathbf{x}}_{10}$ is large: what frequency is present in the signal?

Well, the normalized frequency is $10/256 = .04$, and converting to Hertz from the sampling rate, the frequency present is $F = .04 * 500 = 20\text{Hz}$. So there is a 20Hz sinusoid somewhere in your signal.

Be careful about frequency wrap around. For instance, if $\hat{\mathbf{x}}_{200}$ is large, this corresponds to a normalized frequency of $\omega = 200/256 = .78$, but we should map this to a frequency in the range $[-1/2, 1/2]$. Thus we use $\omega' = .78 - 1 = -.22$ and the corresponding real frequency is $-.22 * 500 = -110\text{Hz}$. (And in practice, we ignore the negative.)

11.3 The six FT properties, for the discrete FT

Here are some important properties of the Discrete Fourier transform, that we want on sequences of length N :

1. $\hat{\hat{x}}_n = x_{-n}$.
2. The FT has an inverse, given by a similar formula.

3. $\widehat{\mathbf{x} * \mathbf{y}} = \widehat{\mathbf{x}} \cdot \widehat{\mathbf{y}}$.
4. $\widehat{\mathbf{x} \cdot \mathbf{y}} = \widehat{\mathbf{x}} * \widehat{\mathbf{y}}$.
5. $\|\mathbf{x}\| = \|\widehat{\mathbf{x}}\|$.
6. $\langle \mathbf{x}, \mathbf{y} \rangle = \langle \widehat{\mathbf{x}}, \widehat{\mathbf{y}} \rangle$.

Unfortunately, this does not quite work with the way we defined the discrete Fourier transform above. First, in property 1, we have to define indices modulo N . Second, it turns out in many of the formulas we are often off by a factor of \sqrt{N} or even N . We also need to use circular convolution, which we will define in a section below. Here are the actual properties that work. (We use \otimes to indicate circular convolution.)

1. $\widehat{\widehat{x}}_0 = Nx_0$, and $\widehat{\widehat{x}}_n = Nx_{N-n}$ for $n > 0$.
2. The FT has an inverse, $\check{\mathbf{x}}_k = \frac{1}{N} \sum_{n=0}^{N-1} x_n u^{+nk}$, $0 \leq k \leq N-1$,
3. $\widehat{\mathbf{x} \otimes \mathbf{y}} = \widehat{\mathbf{x}} \cdot \widehat{\mathbf{y}}$.
4. $N\widehat{\mathbf{x} \cdot \mathbf{y}} = \widehat{\mathbf{x}} \otimes \widehat{\mathbf{y}}$.
5. $\sqrt{N}\|\mathbf{x}\| = \|\widehat{\mathbf{x}}\|$.
6. $N\langle \mathbf{x}, \mathbf{y} \rangle = \langle \widehat{\mathbf{x}}, \widehat{\mathbf{y}} \rangle$.

Maybe the following example explains property 1 better. Repeating the FT twice on the vector $[1, 2, 3, 4]$, we get

$$FT(FT([1, 2, 3, 4])) = 4 \cdot [1, 4, 3, 2] = [4, 16, 12, 8]. \quad (11.9)$$

That is, the first entry stays where it belongs, and the other guys get reversed. And then they all get multiplied by 4.

11.3.1 The normalized DFT

To get rid of some of those annoying factors of N and \sqrt{N} , it is common to define a normalized DFT as

$$\widehat{\mathbf{x}}_k = \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} x_n u^{-nk}, \quad 0 \leq k \leq N-1, \quad (11.10)$$

and the normalized inverse transform as

$$\check{\mathbf{x}}_k = \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} x_n u^{+nk}, \quad 0 \leq k \leq N-1, \quad (11.11)$$

again with $u = e^{2\pi i/N}$. Then, the six properties work out like this:

1. $\widehat{\widehat{x}}_0 = x_0$, and $\widehat{\widehat{x}}_n = x_{N-n}$ for $n > 0$.
2. The FT has an inverse, $\check{\mathbf{x}}_k = \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} x_n u^{+nk}$, $0 \leq k \leq N-1$,
3. $\widehat{\mathbf{x} \otimes \mathbf{y}} = \sqrt{N} \widehat{\mathbf{x}} \cdot \widehat{\mathbf{y}}$.

$$4. \sqrt{N} \widehat{\mathbf{x} \cdot \mathbf{y}} = \widehat{\mathbf{x}} \otimes \widehat{\mathbf{y}}.$$

$$5. ||\mathbf{x}|| = ||\widehat{\mathbf{x}}||.$$

$$6. \langle \mathbf{x}, \mathbf{y} \rangle = \langle \widehat{\mathbf{x}}, \widehat{\mathbf{y}} \rangle.$$

Note that now the annoying factor \sqrt{N} only appears in properties 4 and 5.

Because this normalized transform preserves the inner product, it is called a *unitary* operator. The matrix of the operator have the property that its inverse is the same as its conjugate transpose:

$$A^{-1} = A^*. \quad (11.12)$$

This would be interesting to you if you like linear algebra.

Here's a little calculation to see we get the inverse property: Write matrices A, B with entries

$$A_{jk} = \frac{1}{\sqrt{N}} u^{-jk}, B_{jk} = \frac{1}{\sqrt{N}} u^{+jk}. \quad (11.13)$$

Here, A is the matrix of the DFT and B is its conjugate transpose. Then the product of the two matrices has entries

$$(AB)_{jk} = \sum_n A_{jn} B_{nk} = \frac{1}{N} \sum_n u^{-jn} u^{+nk} = \frac{1}{N} \sum_n u^{+n(k-j)} \quad (11.14)$$

and as we saw in the appendix, these powers of roots of unity sum up to zero, except in the case when the power $k - j$ is zero. Thus the matrix AB has zeros everywhere except on the diagonal ($k = j$), where we get the value $1 = \sum_n 1 = N^{-1} N$. So AB is the identity matrix, and thus B is the inverse matrix of A . That is, B is the inverse Fourier transform.

11.4 MATLAB quirks

MATLAB normalizes its DFT and the inverse in a unique way, which you should be aware of. The FT is given by

$$\widehat{x}_k = \sum_{j=0}^{N-1} x_j u^{-jk} \quad (11.15)$$

and the inverse FT is given by

$$\check{x}_k = \frac{1}{N} \sum_{j=0}^{N-1} x_j u^{+jk}, \quad (11.16)$$

where u is the N -th root of unity, $u = e^{2\pi i/N}$.

So rather than sharing a normalization factor of \sqrt{N} for both the forward and inverse transform, MATLAB puts nothing on the forward transform, and the whole N on the inverse.

Also, MATLAB indexes its vectors from 1 to N , rather than from 0 to $N - 1$. So you have to keep track of that shift by one when thinking about frequencies and so forth.

11.5 Circular convolution

Circular convolution is a little different than the convolution we learnt for sequences. But not a lot different.

The main point is that when you convolve two finite sequences, the beginnings and the end need to wrap around, in order to obtain the FT property that convolutions map to products.

The circular convolution is defined by the same formula as we have for infinite sequences,

$$z_n = \sum_{k=0}^{N-1} x_{n-k} y_k, \quad (11.17)$$

except when the index $n - k$ becomes negative, we replace it with $N + n - k \geq 0$.

Here is an example. Take sequences of length 4, (x_0, x_1, x_2, x_3) and (y_0, y_1, y_2, y_3) . Then we have the result $\mathbf{z} = \mathbf{x} \otimes \mathbf{y}$ in circular convolution as

$$z_0 = x_0 y_0 + x_3 y_1 + x_2 y_2 + x_1 y_3, \quad (11.18)$$

$$z_1 = x_1 y_0 + x_0 y_1 + x_3 y_2 + x_2 y_3, \quad (11.19)$$

$$z_2 = x_2 y_0 + x_1 y_1 + x_0 y_2 + x_3 y_3, \quad (11.20)$$

$$z_3 = x_3 y_0 + x_2 y_1 + x_1 y_2 + x_0 y_3. \quad (11.21)$$

Weird, but that's the way it is.

For you math gurus, the indices are simply computed modulo N . Another way to see these sums is as follows: for z_2 , it is computed as the sum of all terms $x_j y_k$ where either $j + k = 2$ or $j + k = 2 + N$. And similar for other z_n . Check out the last example to see this.

Or mentally think of the sequences as lying in two loops: line up the terms in the loop, multiply pointwise, and sum. Shift the loops relative to each other, to get the different z_n .

We now can check that the DFT of a circular convolution gives the product of the two DFTs. (That is what property 3 in our list.)

Given two vectors \mathbf{x}, \mathbf{y} of length N , we write the circular convolution as

$$(\mathbf{x} \otimes \mathbf{y})_n = \sum_{k=0}^{N-1} x_{n-k} y_k, \quad (11.22)$$

where we understand the convention that x_{n-k} has index $n - k$ computed modulo N . That is, when index $n - k$ becomes negative, we replace it with $N + n - k \geq 0$.

Now,

$$\left(\widehat{\mathbf{x} \otimes \mathbf{y}} \right)_m = \sum_n (\mathbf{x} \otimes \mathbf{y})_n u^{-mn} \quad (11.23)$$

$$= \sum_n \left(\sum_k \mathbf{x}_{n-k} \mathbf{y}_k \right) u^{-mn} \quad (11.24)$$

$$= \sum_{n,k} \mathbf{x}_{n-k} u^{-m(n-k)} \mathbf{y}_k u^{-mk}, \text{ change variables, } j = n - k \quad (11.25)$$

$$= \sum_{j,k} \mathbf{x}_j u^{-mj} \mathbf{y}_k u^{-mk} \quad (11.26)$$

$$= \left(\sum_j \mathbf{x}_j u^{-mj} \right) \left(\sum_k \mathbf{y}_k u^{-mk} \right) \quad (11.27)$$

$$= \widehat{\mathbf{x}}_m \widehat{\mathbf{y}}_m. \quad (11.28)$$

Thus, we have that the DFT of the convolution $\mathbf{x} \otimes \mathbf{y}$ is the product of the two DFTs, $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$. This verifies property 3.

11.6 Zero padding

This circular convolution is kind of a pain. We are forced to use it, because of the way the DFT works on products. But it causes this strange “wrap around” effect, where the signal near the beginning $n = 0$ starts to interfere with the signal near the end $n = N$. This leads to artifacts in your signal processing algorithms.

One way to avoid this is with zero padding. We tack on some zeros to the end of our vectors, and then a circular convolution behaves like the regular convolution for infinite sequences.

For example, let's let

$$\mathbf{x} = (x_0, x_1, x_2, x_3, 0, 0, 0, 0) \text{ and } \mathbf{y} = (y_0, y_1, y_2, y_3, 0, 0, 0, 0). \quad (11.29)$$

That is, we take two vectors of length four, and tack on four zeros to each, to get vectors of length eight. We can compute the circular convolution

$$\mathbf{z} = \mathbf{x} \otimes \mathbf{y} = (x_0y_0, x_1y_0 + x_0y_1, x_2y_0 + x_1y_1 + x_0y_2, \dots, x_3y_2 + x_2y_3, x_3y_3, 0), \quad (11.30)$$

which agrees exactly with the regular convolution $\mathbf{x} * \mathbf{y}$. So by tacking on the zeros, we avoid the wrap around problem. (Technically, we have zeros wrapping around, but they don't add up to anything.)

Note each time you convolve, you will have to tack on more zeros. For some applications, this can be prohibitive, so you just deal with the annoyance of wrap around.

11.7 The Fast Fourier Transform

The DFT is a linear operator, and as such, to compute the DFT of an N -vector should take about N^2 arithmetic operations. When N is large, like a million, the cost of computing N^2 multiplications and additions is ridiculous (a million squared is a trillion operations).

There is an algorithm, called the fast fourier transform, that reduces this to about $N \log_2 N$ operations. For N equal a million, this is only twenty million operations, which is a huge savings. That is, twenty million is much smaller than a trillion.

The algorithm was invented by Cooley and Tukey in the 1960s, and is based on the observation that we can rewrite the basic transform equation

$$\hat{\mathbf{x}}_j = \sum_k \mathbf{x}_k u^{-jk} \quad (11.31)$$

as a sum over the even and odd indices k , so with

$$\hat{\mathbf{x}}_j = \sum_{k \text{ even}} \mathbf{x}_k u^{-jk} + \sum_{k \text{ odd}} \mathbf{x}_k u^{-jk} \quad (11.32)$$

$$= \sum_{r=0}^{N/2-1} \mathbf{x}_{2r} u^{-2rj} + \sum_{r=0}^{N/2-1} \mathbf{x}_{2r+1} u^{-j(2r+1)} \quad (11.33)$$

$$= \sum_r \mathbf{x}_{2r} (u^2)^{-rj} + u^{-j} \sum_r \mathbf{x}_{2r+1} (u^2)^{-jr} \quad (11.34)$$

we see the order N DFT can be computed as the sum of 2 order $N/2$ DFTs. Now this is a speed up of the algorithm since both of the $N/2$ DFTs take $(N/2)^2$ operations each, for a total of $N^2/2$. That is, a speed up by a factor of 2.

You can repeat this operation on the order $N/2$ DFT, writing each of them as a sum of two order $N/4$ DFTs. Now, if N started out as a power of two, you can repeat this over and over again, getting each DFT computed as the sum of two smaller operations.

A careful count of operations shows that this divide and conquer algorithm reduces the operations count for order N^2 to order $N \log_2 N$.

Now, in principle, you will never have to write FFT code yourself. Some jerk somewhere else has written code that is way better than what one of us mere mortals could do. (Although I had to do this back in the 80's.) In fact, there is a web page out there somewhere, where you can make a request for code sample of the FFT, in just about any computer language, and it will automatically generate highly optimized code for whatever special purpose you need it for. The stuff has been highly studied, and there is great code out there to do the FFT for you.

MATLAB has very good FFT code in it.

Chapter 12

General Fourier transforms

It turns out that the Fourier transform is a rather general notion that converts a function on one space into a function on another space. For different spaces, there are different transforms, but they all look fairly similar – involving a sum or integral, and some complex exponentials.

The spaces involved are called locally compact abelian groups. They include all the Euclidean spaces \mathbb{R}^n , multi-integers \mathbb{Z}^N , tori \mathbb{T}^n , finite cyclic groups \mathbb{Z}/N , as well as product of these things, and some more exotic examples. We won't have much chance to look at these general Fourier transforms, but let me assure you, they are very interesting. The six properties all hold in these general situations.

In physics, and geophysics, the Fourier transform on the plane \mathbb{R}^2 , on space \mathbb{R}^3 and space-time \mathbb{R}^4 are all important. In signal processing, we are often interested in the integer spaces \mathbb{Z}^2 and \mathbb{Z}^3 , as well as products of the finite spaces \mathbb{Z}/N .

Just so you can say you've seen it, the 2D Fourier transform of a function $f(x, y)$ on the real plane \mathbb{R}^2 is given by

$$\widehat{f}(\xi, \eta) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) e^{-2\pi i(\xi x + \eta y)} dx dy, \text{ for all } \xi, \eta \in \mathbb{R}. \quad (12.1)$$

Notice the resulting function is also a function on the plane, but in the dual variables ξ, η . The 3D transform is similar, except now you get a function of three variables, a triple integral, and the exponential function includes a factor $\xi x + \eta y + \zeta z$. The inverse transform simply replaces the $-2\pi i$ with $+2\pi i$.

The Fourier transform of a 2D sequence \mathbf{x}_{mn} is defined by the sum

$$\widehat{\mathbf{x}}(\xi, \eta) = \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} e^{-2\pi i(\xi m + \eta n)} \quad (12.2)$$

which is a 1-periodic function in both variables ξ, η . This transform is very useful for processing digital images, where the 2D sequence \mathbf{x}_{mn} represents intensity values for pixels in a picture. The inverse transform involves a double integral over the square $[0, 1] \times [0, 1]$.

Chapter 13

Symmetries of the Fourier transform

You should know this symmetry. When the signal \mathbf{x} is a real sequence, the Fourier transform $\hat{\mathbf{x}}$ is a Hermetian function. That is, $\hat{\mathbf{x}}(-\omega) = \overline{\hat{\mathbf{x}}(\omega)}$. That is, when we replace the argument ω with the negative, the function gives its complex conjugate.

Because of this, we have a symmetry in the amplitude, $|\hat{\mathbf{x}}(-\omega)| = |\hat{\mathbf{x}}(\omega)|$.

You will see this symmetry all the time when you plot the amplitude of the FT of a real signal, where the peaks at positive frequencies match peaks at the negative frequencies.

As a simple example, take a signal that is the sum of two sine waves, at two different frequencies, two different amplitudes. In Figure 13.1, we see the plot of the signal, $f(t) = \sin(2\pi 10t) + .3 * \sin(2\pi 50t)$ and its corresponding Fourier transform. The peaks at $\omega = .02$ and $\omega = .1$ are matched by identical peaks in the negative frequencies. We always see symmetry like this for a real signal.




Figure 13.1: A real signal made up of sine waves, and the FT showing symmetry in frequency content.

Mathematically, we see this symmetry for a real signal because, when we take a complex con-

jugate of the Fourier transform, and apply it to the defining sum, we have

$$\overline{\widehat{\mathbf{x}}(\omega)} = \overline{\sum x_n e^{-2\pi i n \omega}} \quad (13.1)$$

$$= \sum \overline{x_n e^{-2\pi i n \omega}} \quad (13.2)$$

$$= \sum x_n e^{+2\pi i n \omega} \text{ since } x_n \text{ is real} \quad (13.3)$$

$$= \sum x_n e^{-2\pi i n (-\omega)} \quad (13.4)$$

$$= \widehat{\mathbf{x}}(-\omega). \quad (13.5)$$

Thus, the symmetry holds.

There are other symmetries to be aware of: when the signal is imaginary. When it is even. When it is odd. See Karl for a nice picture of various possibilities. These are easy enough to verify on your own.

But the most important one is the one about real signals.

Chapter 14

The uncertainty principle

Fact: if a signal is concentrated in the time domain, then its energy is spread out in the frequency domain. And conversely: if the energy is concentrated in the frequency domain, then the signal must be spread out in the time domain.

As a simple example, in Figure 14.1 we show the Gaussian signal $f(t) = e^{-t^2}$ and its corresponding Fourier transform, which is $\hat{f}(\omega) = \sqrt{\pi}e^{-\pi^2\omega^2}$. We note the signal is rather wide (broad), and the Fourier transform is narrow. Compare this with the example in Figure 14.2 we show the Gaussian signal $f(t) = e^{-9t^2}$ and its corresponding Fourier transform, which is $\hat{f}(\omega) = \sqrt{\pi/9}e^{-\pi^2\omega^2/9}$. Here, the signal is broad, and the Fourier transform is narrow.




Figure 14.1: A wide Gaussian signal, and its Fourier transform, which is narrow.

For another example, in Figure 14.3 we show a boxcar signal $f(t) = 1$ on $[-2, 2]$ and its corresponding Fourier transform, which is a sinc function. We note the signal is rather wide (broad), and the Fourier transform is narrow. Compare this with the example in Figure 14.4, where we show a narrow boxcar signal supported on the interval $[-1/3, 1/3]$, and its corresponding Fourier transform, which is a wider sinc function.




Figure 14.2: A narrow Gaussian signal, and its Fourier transform, which is wide.

This is the essence of the uncertainty principle. It can be made more general, applying to arbitrary signals – examples more general than the simple Gaussians or boxcars.

14.1 Boring math calculations: FT of Gaussians and boxcars

Don't get distracted by the following mathematical computations. The important part of the uncertainty principle is the wide/narrow pairing in time/frequency. This section just computes the details to support the observations above.

We can quickly compute the Fourier transform of a Gaussian. First note that by completing the square, we can write

$$-t^2 - 2\pi i\omega t = -(t + \pi i\omega)^2 - \pi^2\omega^2, \quad (14.1)$$

and so the (real line) Fourier transform of the Gaussian $f(t) = e^{-t^2}$ is given by

$$\hat{f}(\omega) = \int_{-\infty}^{\infty} e^{-t^2} e^{-2\pi i\omega t} dt \quad (14.2)$$

$$= \int_{-\infty}^{\infty} e^{-t^2 - 2\pi i\omega t} dt, \text{ complete the square} \quad (14.3)$$

$$= \int_{-\infty}^{\infty} e^{-(t+\pi i\omega)^2 - \pi^2\omega^2} dt, \text{ then split up the exp} \quad (14.4)$$

$$= \int_{-\infty}^{\infty} e^{-(t+\pi i\omega)^2} e^{-\pi^2\omega^2} dt, \text{ then pull out one exp} \quad (14.5)$$

$$= e^{-\pi^2\omega^2} \int_{-\infty}^{\infty} e^{-(t+\pi i\omega)^2} dt, \text{ change variables } x = t + \pi i\omega \quad (14.6)$$

$$= e^{-\pi^2\omega^2} \int_{-\infty}^{\infty} e^{-x^2} dx, \text{ and integrate} \quad (14.7)$$




Figure 14.3: A wide boxcar signal, and its Fourier transform, which is narrow.

$$= e^{-\pi^2 \omega^2} \sqrt{\pi}. \quad (14.8)$$

That is, the Fourier transform of $f(t) = e^{-t^2}$ is $\hat{f}(\omega) = \sqrt{\pi} e^{-\pi^2 \omega^2}$, which is also a Gaussian, with a different width.

For more general Gaussians, $f(t) = e^{-(t/a)^2}$, the Fourier transform is $\hat{f}(\omega) = \frac{\sqrt{\pi}}{a} e^{-\pi^2 a^2 \omega^2}$. You can verify this using a simple change of variables. Note that the width of the signal is proportional to the parameter a , and the width of the Fourier transform is proportional to $1/a$. So, when a is large, the signal is wide and the FT is narrow. And vice versa, as shown in Figures 14.1 and 14.2.

For a boxcar supported on $[-a, a]$, the Fourier transform is

$$\hat{f}(\omega) = \int_{-a}^a e^{-2\pi i \omega t} dt = \frac{\sin(2\pi a \omega)}{\pi \omega}. \quad (14.9)$$

Again, if a is large, the boxcar is wide, and the sinc function is narrow. And vice versa, as shown in Figures 14.3 and 14.4.

If you think these signals are too simple, consider a modulated Gaussian signal $f(t) = e^{-(t/a)^2} e^{2\pi i \omega_0 t}$. Its Fourier transform is also a Gaussian, but unmodulated, and centered at frequency ω_0 in the frequency domain. This signal represents an oscillating sinusoid that starts out with zero amplitude, grows to full strength at time $t = 0$, and then decays to zero again. This is a useful model of a physical oscillating signal.

You might also consider a modulated boxcar. This represents an oscillating signal that is abruptly turned on, then turned off. Its FT is a sinc function, centered at frequency ω_0 .




Figure 14.4: A narrow boxcar signal, and its Fourier transform, which is wide.

Chapter 15

Calculus and the Fourier transform

The Fourier transforms of a function on the real line and its derivative are closely related: namely, if $\widehat{f}(\omega)$ is the FT of a function $f(t)$, then the FT of the derivative $f'(t)$ is given by

$$\widehat{(f')}(\omega) = 2\pi i \omega \widehat{f}(\omega), \text{ for all } \omega \in \mathbb{R}. \quad (15.1)$$

That is, we just get the FT of the function, multiplied by $2\pi i \omega$.

Similarly, for a 1-periodic function, the FT of the derivation is given by the sequence

$$\widehat{(f')}(n) = 2\pi i n \widehat{f}(n), \text{ for all } n \in \mathbb{Z}. \quad (15.2)$$

What does this say? First, it says that the derivative acts like a multiplier in the Fourier domain. Specifically, it is multiplication by the linear function $\omega \mapsto 2\pi i \omega$. Since this function is large in magnitude, for large frequencies ω , it means that derivatives will boost (amplify) high frequency signal, while attenuating low frequencies. Derivation is sort of the ultimate in high pass filtering.

This can be a problem in real systems, and in mathematical calculations. Noise tends to be distributed across all frequencies; differentiation will boost the high frequency noise, which really can mess up your system.

Second, this tells us that differentiation is a linear time invariant system, since it acts as a multiplication in the frequency domain. Finding its impulse response is tricky – it turns out to be a distribution, which is a generalized function. We won't discuss these in this class.

Third, it suggests we have a way now to solve certain differential equations. For instance, suppose we wish to find a solution to the ODE

$$y'' - y = e^{-t^2}, \quad (15.3)$$

where $y = y(t)$ is some unknown function on the real line. Applying the Fourier transform to this equation, we get

$$(2\pi i \omega)^2 \widehat{y} - \widehat{y} = e^{-\pi^2 \omega^2}, \quad (15.4)$$

where we use the fact that the FT of the Gaussian e^{-t^2} is $e^{-\pi^2 \omega^2}$. Simplifying this last equation algebraically, we have

$$\widehat{y}(\omega) = -\frac{e^{-\pi^2 \omega^2}}{4\pi^2 \omega^2 + 1}. \quad (15.5)$$

Thus the solution $y(t)$ will be the inverse Fourier transform of this Gaussian divided by $4\pi^2 \omega^2 + 1$. That is,

$$y(t) = \int_{-\infty}^{\infty} -\frac{e^{-\pi^2 \omega^2}}{4\pi^2 \omega^2 + 1} e^{+2\pi i \omega t} d\omega. \quad (15.6)$$

We may have to solve this integral numerically, but at least in principle we have a way of computing the solution to the differential equation.

Finally, we note that while the function $\widehat{f}(\omega)$ might be well-behaved, multiplying it by $2\pi i\omega$ might cause it to grow quickly as ω grows to infinity. This could cause the derivative to behave badly, perhaps even so badly that it cannot exist. We will see a student presentation in class of precisely such a “bad” function: Weierstrauss’ example of a continuous, no-where differentiable function.

15.1 More math: derivatives

Let’s see where these derivative formulas come from.

For a 1-periodic function $f(t)$, we have from the inverse formula for sequences that

$$f(t) = \sum_n \widehat{f}(n) e^{+2\pi i n t}, \quad (15.7)$$

and so if we differentiate term by term, we have

$$f'(t) = \sum_n \widehat{f}(n) \frac{d}{dt} e^{+2\pi i n t} = \sum_n \widehat{f}(n) 2\pi i n e^{+2\pi i n t}. \quad (15.8)$$

On the other hand, we know by definition that

$$f'(t) = \sum_n \widehat{(f')}(n) e^{+2\pi i n t}, \quad (15.9)$$

so, we can read off the Fourier coefficients for the derivative from the factors in the sum, thus

$$\widehat{(f')}(n) = 2\pi i n \widehat{f}(n). \quad (15.10)$$

For a function $f(t)$ on the real line, we have from the inverse formula that

$$f(t) = \int \widehat{f}(\omega) e^{+2\pi i \omega t} d\omega, \quad (15.11)$$

and so if we differentiate under the integral sign, we have

$$f'(t) = \int \widehat{f}(\omega) \frac{d}{dt} e^{+2\pi i \omega t} d\omega = \int \widehat{f}(\omega) 2\pi i \omega e^{+2\pi i \omega t} d\omega. \quad (15.12)$$

Again, we know by definition that

$$f'(t) = \int \widehat{(f')}(\omega) e^{+2\pi i \omega t} d\omega, \quad (15.13)$$

so we can read off the Fourier coefficients for the derivative from the factors in the integrand, and thus

$$\widehat{(f')}(\omega) = 2\pi i \omega \widehat{f}(\omega). \quad (15.14)$$

To be a good mathematician, I should really justify this term-by-term differentiation, or the derivative under the integral sign. Let’s not. Just let me assure you that if we are really interested, we could do it.

Chapter 16

Special Filters

We know that we can filter a signal by convolving it with some other signal, which is the general form of a LSI system. In the Fourier domain, the convolution turns into multiplication.

So this gives us a new way to filter signals. Fourier transform it, multiply it by some fixed function that has the frequency response we want, then inverse Fourier transform the results to get the filtered signal.




Figure 16.1: A signal, filtered with a low-pass brick wall, to get the smoothed output.

Symbolically, we take our signal \mathbf{x} , Fourier transform it to $\hat{\mathbf{x}}$, then multiply it by the filter response function h to get $h \cdot \hat{\mathbf{x}}$. Then we apply the inverse Fourier transform. Or, in MATLAB syntax, we have

```
y = ifft( h.* fft(x))
```

Since the FFT is fast, this is a fast way to do a filter.

For instance, given a slowly varying signal with some high frequency noise, we can filter it by multiplying the FT with a “brick wall” type of multiplier, that is equal to 1 on the low frequencies, and 0 on the high frequencies. The result is shown in Figure 16.1.

To do a high pass filter, you change the brick wall function to be equal to zero on the low frequencies, and equal to one on the high frequencies.

Remember that the function h should respect the symmetries of the Fourier transform, if you want to get a real-valued signal as an output. So for instance, in MATLAB, we want $h(2)$, to equal $h(N)$, $h(3)$ to equal $h(N-1)$, $h(4)$ to equal $h(N-2)$, etc. As an example, here is the lowpass brickwall that I used in the example above:

```
h = [ones(1,K+1),zeros(1,N-K-K-1) ones(1,K)];
```

Here, N is the length of the vector, and K counts how many ones to put in the pass band. See how the ones are both at the beginning and at the end of the h vector.

Here is another example of a special filter. We do a phase rotation of a signal, by setting the Fourier multiplier h to equal $e^{2\pi is}$ on the positive frequencies, and $e^{-2\pi is}$ on the negative frequencies. The parameter s controls how much of a rotation we get. The value $s = .25$ is a rotation by a quarter cycle, and turns peaks into zero crossings, and vice versa. The value $s = .5$ is a rotation by half a cycle, and just flips a waveform upside down. The value $s = 1.0$ is a rotation by a full cycle, which of course means no change: we have walked around the circle and returned to the beginning. Figure 16.2 shows the result of various phase rotation on a nice Rickard wavelet.

The code for this phase rotation is listed here.

```
% PhaseRot.m
% Computes a nice animation of a rotating wavelet
% We use a simple allpass filter by multiplying by e^{+/- theta}
% in the positive and negative frequency space
N=10001;
N2=501;
t = linspace(-5,5,N);
x = (.5-t.^2).*exp(-t.^2);
h = zeros(1,N);
for s=0:.01:1
    h(2:N2) = exp(2*pi*i*s);
    h((N2+1):N) = exp(-2*pi*i*s);
    plot(t,real(ifft(h.*fft(x))))
    xlabel(s)
    ylim([- .5, .5])
    pause(.25)
end
```




Figure 16.2: A wave being phase rotated through one complete cycle.

Chapter 17

FIR, IIR filters

So, if we can just filter in the Fourier domain, why do we ever worry about FIR and IIR filters?

The main reason is that to have to do the FFT of a signal, you have to have the whole signal in the computer first. There are many situations where this is not the case. For instance, you might have a long signal that is too big to store in the computer. This often happens with massive amounts of audio data, video data, or even raw seismic data.

Or you might have signals coming at you continuously and you can't wait to receive the whole thing before processing it. For instance, in a telephone conversation that is being processed, you have to receive and process the speech as it arrives, and send it off to the receiver as the sound is happening.

For these reasons, we might need to use FIR or IIR filters to process the signal as it comes in.

This raises the question of how to design an FIR or IIR filter with desired characteristics.

We don't have time to go into the details. But a quick idea is this. Say you want your filter to have a frequency response given by function h . Then the filtered response is the inverse FT of $h \cdot \hat{\mathbf{x}}$. By the convolution theorem, this is just $\check{h} * \mathbf{x}$, the convolution of the signal \mathbf{x} with the inverse transform of h . So now it is just a matter of using \check{h} to design the coefficients in your FIR filter.

(There are problems with this, mainly that \check{h} is rarely a finite sequence. And we need finite sequences to get an FIR. There are tricks to get good approximations to this.)

Chapter 18

Laplace transforms

Oops, no time for this. The basic idea is to transform a function $f(t)$ on the positive real line as

$$\hat{f}(s) = \int_0^{\infty} f(t)e^{-st} dt. \quad (18.1)$$

Because of the exponential decay, this integral usually converges for s sufficiently large (and positive). Turns out this is good for solving certain differential equations, especially those that show up in electrical engineering.

Chapter 19

Wavelet transforms

Fourier theory is concerned with the *analysis* and *synthesis* of signals using *sinusoids* of various frequencies. By contrast, wavelet theory is concerned with the *analysis* and *synthesis* of signals using *wavelets* of various *translations* and *dilations*. A wavelet is a small wave, and some examples are shown in Figure 19.2.




Figure 19.1: Four different examples of wavelets, or small waves.

Now, if this idea of analysis and synthesis is too abstract, let's think of food instead of signals. Nutritional theory allows us to analyze any food to determine what it is made up of. For instance, a chocolate bar could be analyzed, and it is determined that it is made up of 20 grams of sugar, 10 grams of milk protein, 30 grams of cocoa, and so on. To synthesize a chocolate bar, we take these different component (sugar, milk, cocoa, etc) and mix them together in the appropriate amounts, to create a chocolate bar.

The same thing with signals. To analysis a signal is to determine what components (eg sine waves) are present in the signal. To synthesize a signal is to reconstruct it as a sum of various

components (eg. a linear combination of sine waves). In Fourier theory, we analyze a signal f on an interval by computing its Fourier coefficients

$$\hat{f}(n) = \int f(t)e^{-2\pi int} dt, \quad (19.1)$$

and then we can reconstruct this same signal as a sum of sinusoids

$$f(t) = \sum_n \hat{f}(n)e^{2\pi int}. \quad (19.2)$$

19.1 Dilations and translations of a wavelet

For wavelet theory, we start with a single wavelet function $\phi(t)$, such as any one of the examples in Figure 19.2. We can translate the wavelet $\phi(t)$ by shifting the argument to $\phi(t-b)$, where b is a fixed constant. We dilate the wavelet by scaling the argument as $\phi(t/a)$, for some constant $a \neq 0$. It is useful to rescale the height of the function $\phi(t/a)$ so that the energy (integral of its square) stays the same. We combine both the translation and scaling of the wavelet to define a the function

$$\phi_{a,b}(t) \frac{1}{\sqrt{|a|}} \phi\left(\frac{t-b}{a}\right). \quad (19.3)$$

It is this functions $\phi_{a,b}(t)$ that we will use to analyze and synthesis any given signal $f(t)$.




Figure 19.2: A wavelet $\phi(t)$ centered at zero, its translate $\phi(t-3)$, and the translate/dilate $\frac{1}{\sqrt{2}}\phi(\frac{t-7}{2})$. Note the dilated one is twice as wide, and 70% of the height of the others.

19.2 Analysis: wavelet coefficients

Given a function $f(t)$ on the real line, and a fixed wavelet $\phi(t)$, we *analyze* the signal by defining the wavelet coefficients $c_{a,b}$ as the numbers obtained from the inner product of function f with wavelet $\phi_{a,b}$. That is, we define

$$c_{a,b} = \langle f, \phi_{a,b} \rangle = \int_{-\infty}^{\infty} f(t) \overline{\phi_{a,b}(t)} dt, \quad (19.4)$$

where this can be computed for any real numbers a, b with $a \neq 0$.

These coefficients $c_{a,b}$ tell us a lot about the function $f(t)$. The index a is scale, while the index b is position. The number $c_{a,b}$ tells us how the function $f(t)$ behaves at a scale of a and a position of b . The b part is easy: it just says we have local information about the function near time $t = b$. The a part is a bit harder to understand. For large a , like $a = 100$, the coefficient tells us if the function has “large scale features” that look a lot like the original wavelet ϕ spread out by a factor of 100. For small a , like $a = 0.01$, the coefficient tells us if the function has “small scale features” that look a lot like the original wavelet ϕ , shrunk down to a narrow range. Roughly speaking, large scale corresponds to slowly changing, low frequency behaviour, while small scale corresponds to rapidly oscillating, high frequency behaviour. That is, scale roughly indicates something like 1 over frequency.

That’s it for analysis. Notice that the wavelet analysis describes local information about the signal, which is quite different from the Fourier analysis we have seen. For instance, a simple boxcar has Fourier coefficients that are non-zero almost all frequencies; thus a small boxcar (non-zero on a narrow interval) influences the Fourier coefficients out to infinity. For the wavelet coefficients of a boxcar, only those coefficients $c_{a,b}$ with b near position of the boxcar will be non-zero. So the boxcar only influences a narrow range of the coefficients $c_{a,b}$.

As an example of analysis, we take a function $f(t) = \sin(200\pi t^3)$ which is a simple chirp (an oscillating signal that starts with a slow oscillations, and ramps up to a fast oscillation). Taking the wavelet transform, we compute the values of $c_{a,b}$ for a range of values a, b and plot. The result is shown in Figure 19.3. The display shows how the signal changes in time: at the left, it starts with low frequencies, then ramps up to the right with high frequency. In this case, we used a Morlet wavelet to compute the wavelet coefficients. Other wavelets will show similar behaviour, but the details will be different.

I want to stress that you can pretty much choose any function as your analyzing wavelet $\phi(t)$. Whether this gives you anything useful depends on the choice of ϕ , which is why so much work is done in applications to find useful wavelets. Often a wavelet is designed specially for a particular application. When you start using wavelets, the first thing to check out is what wavelets are other people using in applications that are related to yours.

19.3 Synthesis: reconstructing the signal

Synthesis is harder than the analysis. We ask the question: is it possible to rebuild the signal $f(t)$ by a sum, or integral involving the coefficients $c_{a,b}$ and the functions $\phi_{a,b}(t)$. The answer is yes, provided the wavelet ϕ was carefully designed.

In the orthonormal wavelet case, we have a reconstruction formula given by

$$f(t) = \sum_{a=2^n, b=m2^n} c_{a,b} \phi_{a,b}(t), \quad (19.5)$$




Figure 19.3: A plot of the wavelet coefficients for a chirp signal. See how it clearly indicates a signal which starts on the left at a low frequency (high scale a) and ends up at the right with a high frequency (low scale a).

where the sum is taken over all scales of the form $a = 2^n$ a power of two, and $b = m2^m$ an integer multiple of that same power. Here, the design problem is to find a wavelet ϕ so that these translates and dilates are all orthogonal.

It's not hard to see that the Haar wavelet works. But so do the Daubechies wavelets, which is an amazing piece of mathematics.

In the continuous wavelet case, we have a reconstruction formula given by

$$f(t) = \frac{1}{C} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} c_{a,b} \phi_{a,b}(t) \frac{da db}{a^2}, \quad (19.6)$$

where the normalizing constant C is given by

$$C = \int_{-\infty}^{\infty} \frac{|\hat{\phi}(\omega)|^2}{|\omega|} d\omega. \quad (19.7)$$

Here, we need to know this second integral is finite. This puts a condition on the wavelet that its Fourier transform $\hat{\phi}$ vanishes at $\omega = 0$. Turns out that we have many choices for a wavelet ϕ that satisfies this condition.

With more time, we would go into more details.

19.4 Fast transforms

The discrete wavelet transform takes a vector of length N and computes a new vector, also of length N , consisting of certain wavelet coefficients $c_{a,b}$. It is a linear transformation, just as the discrete

Fourier transform is. However, it is much faster, on the order of N operations, rather than $N \log N$ for the FFT. This makes the wavelet transform advantageous in many situations.

The details of how this is done is best left to another course. The key idea is using two pairs of matched, FIR filters that do all the necessary computations in computing the wavelet coefficients, and reconstruction the signal from those coefficients. The mathematics involved in designing those FIR filters is elegant, and quite intriguing. But, we are out of time! There are many software sources available to do these transforms efficiently (including MATLAB).

19.5 Multiresolution analysis

Oh dream on. We have no time for this here. Sorry.

19.6 Application: data compression

There is this amazing phenomena with the wavelet analysis of certain data. Often, many of the wavelet coefficients $c_{a,b}$ are zero, or close to zero. So you can in principle forget about them, and reconstruct your signal using only the larger non-zero coefficients. This is not a perfect reconstruction. However, it is close enough to be useful in many applications.

For instance, in a digital image, your raw data may include 16 megapixels, which means at least 16 million numerical values to store on your computer. By computing the wavelet transform, the computer may discover that of the 16 million wavelet coefficients, only about 5% of them are significantly different from zero. So we store these 5%, and forget the rest. When it comes time to reconstruct the image (say to display on the screen), we only use those 5% non-zero coefficients, and build a very good approximation to the original image.

So, instead of storing 16 megs of data, we get away with storing about 800,000 numbers. This is a significant savings!

The MATLAB wavelet toolbox contains code to compute these kinds of compressions.

APPENDIX 1: Collected wisdom

In case we are missing the forest for the trees, here are some basic facts about signal processing that you should always keep in mind. (I will add to this as the course goes along.)

Signals:

- Signals are functions (of time, space, etc).
- Many useful signals are sums of sines and cosines.
- A sine wave or cosine wave is specified by a frequency and an amplitude. Negative frequencies give basically the same signal as a positive frequency.
- Sines, cosines are the same function, just shifted in time.
- Unless we know the start time, sine waves and cosine waves at the same frequency are basically the same thing. Same for any shifts of these waves.
- The complex sinusoid $e^{2\pi i \omega t}$ makes the algebra of sines and cosines easy. But in real life, we never really see complex valued waves.

Physical signals:

- We hear sounds in the range 20Hz to 20,000 Hz.
- Seismic waves are measured in the range 4Hz to 150 Hz (approx.) As technology improves, we expand this range.
- Radio waves are in the range of kilohertz (AM), megahertz (FM), gigahertz (cellphones).
- Sound waves are acoustic waves (variations in air pressure). Seismic waves are elastic waves (solid motion). Radio waves are electromagnetic waves.

Sampled signals:

- For practical reasons, we sample signals.
- Sample rate determines the highest frequency we can represent

$$\frac{1}{2}(\text{sample rate}) = \text{Nyquist rate.} \quad (19.8)$$

- Frequencies higher than that are aliased, and cause errors.
- In real systems, we use electronics to eliminate those higher frequencies, before sampling.

Systems:

- Basic mode is “Signal in \rightarrow Signal out.”
- Convolution gives a LSI system, and vice versa.
- Specify a LSI system by its impulse response ***h***.
- Practical LSI system given by finite ***h*** or ratio of two such finite ones.

APPENDIX 2: Eigenvalues and eigenvectors

Recall a matrix A has an eigenvector $\mathbf{x} \neq 0$ and an eigenvalue λ if they satisfy

$$A\mathbf{x} = \lambda\mathbf{x}. \quad (19.9)$$

That is, the matrix applied to vector \mathbf{x} just returns the same vector \mathbf{x} , multiplied by the number λ .

As an example, with the matrix $A = \begin{bmatrix} 2 & 4 \\ 4 & 2 \end{bmatrix}$, we find two eigenvectors $[1, 1]$ and $[1, -1]$ by observing

$$\begin{aligned} \begin{bmatrix} 2 & 4 \\ 4 & 2 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \end{bmatrix} &= 6 \begin{bmatrix} 1 \\ 1 \end{bmatrix} \\ \begin{bmatrix} 2 & 4 \\ 4 & 2 \end{bmatrix} \begin{bmatrix} 1 \\ -1 \end{bmatrix} &= -2 \begin{bmatrix} 1 \\ -1 \end{bmatrix}. \end{aligned}$$

So 6 and -2 are the eigenvalues of the matrix. Notice that $6 - 2 = 4$, the trace of the matrix (sum of the diagonal elements), while $6 * (-2) = -12$, the determinate of the matrix. This is a general property of the eigenvalues: they tell us a lot about the matrix.

In the frequency response, the only difference is that we are working in infinite dimensions. The eigenvectors are signals \mathbf{x} that are complex exponentials at frequency ω . The corresponding eigenvalues are the complex numbers $H(\omega)$. That is, when \mathbf{x} is a complex exponential, we have the eigenvector equation

$$\mathbf{h} * \mathbf{x} = \lambda\mathbf{x}, \quad (19.10)$$

where the eigenvalue λ is the (complex) number $\lambda = H(\omega)$, given by the frequency response evaluated at ω .

APPENDIX 3: Sums and integrals of exponentials

We will use the following facts over and over again in signal processing:

$$\int_{-1/2}^{1/2} e^{-2\pi i n \omega} d\omega = \begin{cases} 1 & n = 0, \\ 0 & n = \pm 1, \pm 2, \pm 3, \dots, \end{cases} \quad (19.11)$$

and

$$\sum_{k=0}^{N-1} e^{-2\pi i n k/N} = \begin{cases} N & n = 0 \\ 0 & n = 1, 2, \dots, N-1. \end{cases} \quad (19.12)$$

The integral result is easy to see. When $n = 0$, the exponential has a zero in the argument, so we have $e^0 = 1$, and the integral of one over an interval of length one is just one. When $n \neq 0$, the exponential can be written as the sum of a cosine and i times a sine, each of which has exactly n cycles in the interval $[-1/2, 1/2]$. When you integrate, the positive and negative parts of each cycle cancels, giving a total integral of zero. You could also do this directly by computing the antiderivative of the exponential, and noticing it too is periodic.

Note there is nothing special about the interval $[-1/2, 1/2]$. We could replace it with any interval of length one.

The sum result is also easy to see. When $n = 0$, the exponential has a zero in the argument, so we are simply summing N copies of $e^0 = 1$, which adds up to N . When $n = 1, 2, \dots, N-1$, we are summing up powers of the non-trivial root of unity $z = e^{-2\pi i n/N}$. These powers z^k are uniformly spread around the unit circle, so like uniformly spread out vectors, they sum to zero.

Or, if you prefer a more explicit calculation,

$$\sum_k z^k = z^0 + z^1 + z^2 + \dots + z^{N-1} = \frac{z^N - 1}{z - 1} = 0, \quad (19.13)$$

since $z = e^{-2\pi i n/N}$ is an N -th root of unity, so $z^N = 1$.

APPENDIX 4: Music and frequencies

Humans can hear sounds in the range of 20Hz to 20,000 Hz. (Hz = Hertz = cycles per second). Basically any periodic wave that repeats itself that many times a second will be heard by the human ear as a musical tone. The frequency (cycles per second) corresponds to pitch of the note. The low (deep) sounds have the low frequency (eg 25) and the high (shrill) sounds have the high frequency

A piano plays notes in the range of about 25 Hz to 4200 Hz. The note “A above middle C” is defined as a pitch of 440 Hz. All the other pitches are defined relative to that frequency. The equal tempered scale uses powers of the number $\alpha = \sqrt[12]{2}$ to define other frequencies as

$$freq = 440 * \alpha^n, \quad (19.14)$$

where n is the number of semitones above “A”, or below it for n negative.

Doubling the frequency ($n = 12$ semitones up) corresponds to moving up an octave in the musical scale. Halving the frequency ($n = -12$ semitones down) moves down an octave.

Bach was involved in the notion of an equal tempered scale, which uses the power of the 12th root of 2 to determine frequencies. You may be familiar with his keyboard composition called “The Well-tempered Klavier.” This was motivated by the desire to get all instruments in an orchestra to sound in tune, not matter what key they were playing in.

Before equal tempering, notes were based on the idea that frequencies that were related as the ratio of simple fractions often sound good together. This is deeply connected to the notion of harmonics in musical tones, but it seems rather mathematical.

For instance, the frequencies 440Hz, 550Hz, 660Hz all sound good together, and their ratios are

$$\frac{550}{440} = \frac{5}{4} \quad \frac{660}{440} = \frac{3}{2}, \quad (19.15)$$

which are simple fraction with small integers in the fractions. These three notes correspond to the three notes in a major triad (A, C#, E). It is interesting to note that the middle note, 550Hz, would actually be as high as 554Hz in the equal tempered scale. To those of you with good ears, the 554Hz tone sounds a little sharp. We get this 554Hz value by going up four semitones in the equal tempered scale, so $440 * 2^{4/12} = 554.37$.

You can experiment with these ideas in MATLAB by setting up some simple signals and playing them out.

```
Fs = 10000; % the sampling rate
dt = 1/Fs; % the time step
T1 = 0:dt:1; % one second of time, in steps of dt
A = sin(2*pi*440*T1); % the note A above middle C
Cs = sin(2*pi*550*T1); % the note C#
E = sin(2*pi*660*T1); % the note E
sound([A,Cs,E],Fs); % play the notes one after the other
sound(.3*(A+Cs+E),Fs); % play the notes all together as one
```

If you are getting tired of hearing pretty sine waves, try raising it to some (odd) power, to get some richer harmonics in your music.

```

Fs = 10000;
dt = 1/Fs;
T1 = 0:dt:1;
A = sin(2*pi*440*T1).^5; % the note A above middle C, with harmonics
Cs = sin(2*pi*550*T1).^5; % the note C#
E = sin(2*pi*660*T1).^5; % the note E
sound([A,Cs,E],Fs);
sound(.3*(A+Cs+E),Fs);

```

To get the major scale, using fractions instead of powers of two, you can use the ratios

$$\frac{1}{1}, \frac{9}{8}, \frac{5}{4}, \frac{4}{3}, \frac{3}{2}, \frac{27}{16}, \frac{15}{8}, \frac{2}{1}. \quad (19.16)$$

In the A major scale, this corresponds to frequencies

$$440, 495, 550, 586.7, 660, 742.5, 825, 880 \text{ Hz}. \quad (19.17)$$

Some might argue that other fractions are better. For instance, we might replace the $5/4$ with $81/64 = (1/4) * (3/2)^4$. If you know something about music theory, you will see this is connected to the circle of fifths idea.

Now, just because its fun, try this. It is a sweep through the frequency range

```

Fs = 10000;
dt = 1/Fs;
T1 = 0:dt:1;
sweep = sin(2*pi*440*T1.^3); % a sweep
sound(sweep,Fs);

```