Fundamentals of Signal Processing

Minh Do

Fundamentals of Signal Processing

Course Authors:

Minh Do

Contributing Authors:

Richard Baraniuk

Hyeokho Choi

Minh Do

Benjamin Fite

Anders Gjendemsjo

Michael Haag

Don Johnson

Douglas Jones

Rob Nowak

Ricardo Radaelli-Sanchez

Justin Romberg

Clayton Scott

Ivan Selesnick

Melissa Selik

Online:

http://cnx.org/content/col10360/1.3/

ii

Table of Contents

iv

0 Introduction to Fundamentals of Signal Processing

0.1 What is Digital Signal Processing?

To understand what is *Digital Signal Processing (DSP)* let's examine what does each of its words mean. " *Signal*" is any physical quantity that carries information. " *Processing*" is a series of steps or operations to achieve a particular end. It is easy to see that *Signal Processing* is used everywhere to extract information from signals or to convert information-carrying signals from one form to another. For example, our brain and ears take input speech signals, and then process and convert them into meaningful words. Finally, the word " *Digital*" in Digital Signal Processing means that the process is done by computers, microprocessors, or logic circuits.

The field DSP has expanded significantly over that last few decades as a result of rapid developments in computer technology and integrated-circuit fabrication. Consequently, DSP has played an increasingly important role in a wide range of disciplines in science and technology. Research and development in DSP are driving advancements in many high-tech areas including telecommunications, multimedia, medical and scientific imaging, and human-computer interaction.

To illustrate the digital revolution and the impact of DSP, consider the development of digital cameras. Traditional film cameras mainly rely on physical properties of the optical lens, where higher quality requires bigger and larger system, to obtain good images. When digital cameras were first introduced, their quality were inferior compared to film cameras. But as microprocessors become more powerful, more sophisticated DSP algorithms have been developed for digital cameras to correct optical defects and improve the final image quality. Thanks to these developments, the quality of consumer-grade digital cameras has now surpassed the equivalence in film cameras. As further developments for digital cameras attached to cell phones (cameraphones), where due to small size requirements of the lenses, these cameras rely on DSP power to provide good images. Essentially, digital camera technology uses computational power to overcome physical limitations. We can find the similar trend happens in many other applications of DSP such as digital communications, digital imaging, digital television, and so on.

In summary, DSP has foundations on Mathematics, Physics, and Computer Science, and can provide the key enabling technology in numerous applications.

0.2 Overview of Key Concepts in Digital Signal Processing

The two main characters in DSP are *signals* and *systems*. A *signal* is defined as any physical quantity that varies with one or more independent variables such as time (one-dimensional signal), or space (2-D or 3-D signal). Signals exist in several types. In the real-world, most of signals are *continuous-time* or *analog signals* that have values continuously at every value of time. To be processed by a computer, a continuous-time signal has to be first *sampled* in time into a *discrete-time signal* so that its values at a discrete set of time instants can be stored in computer memory locations. Furthermore, in order to be processed by logic circuits, these signal values have to be *quantized* in to a set of discrete values, and the final result is called a *digital signal*. When the quantization effect is ignored, the terms discrete-time signal and digital signal can be used interchangeability.

In signal processing, a *system* is defined as a process whose input and output are signals. An important class of systems is the class of *linear time-invariant* (or *shift-invariant*) *systems*. These systems have a remarkable property is that each of them can be completely characterized by an *impulse response function* (sometimes is also called as *point spread function*), and the system is defined by a *convolution* (also referred to as a *filtering*) operation. Thus, a linear time-invariant system is equivalent to a (linear) *filter*. Linear time-invariant

systems are classified into two types, those that have *finite-duration impulse response (FIR)* and those that have an *infinite-duration impulse response (IIR)*.

A signal can be viewed as a *vector* in a *vector space*. Thus, *linear algebra* provides a powerful framework to study signals and linear systems. In particular, given a vector space, each signal can be represented (or expanded) as a *linear combination of elementary signals*. The most important *signal expansions* are provided by the *Fourier transforms*. The Fourier transforms, as with general transforms, are often used effectively to transform a problem from one domain to another domain where it is much easier to solve or analyze. The two domains of a Fourier transform have physical meaning and are called the *time domain* and the *frequency domain*.

Sampling, or the conversion of *continuous-domain real-life signals* to *discrete numbers* that can be processed by computers, is the essential bridge between the analog and the digital worlds. It is important to understand the connections between signals and systems in the real world and inside a computer. These connections are convenient to analyze in the frequency domain. Moreover, many signals and systems are specified by their *frequency characteristics*.

Because any *linear time-invariant system* can be characterized as a *filter*, the design of such systems boils down to the design the associated filters. Typically, in the *filter design* process, we determine the coefficients of an FIR or IIR filter that closely approximates the desired *frequency response* specifications. Together with Fourier transforms, the *z-transform* provides an effective tool to analyze and design digital filters.

In many applications, signals are conveniently described via *statistical models* as *random signals*. It is remarkable that optimum linear filters (in the sense of *minimum mean-square error*), so called *Wiener filters*, can be determined using only *second-order statistics* (*autocorrelation* and *crosscorrelation* functions) of a *stationary process*. When these statistics cannot be specified beforehand or change over time, we can employ *adaptive filters*, where the filter coefficients are adapted to the signal statistics. The most popular algorithm to adaptively adjust the filter coefficients is the *least-mean square (LMS)* algorithm.

Chapter 1

Foundations

1.1 Signals Represent Information

Whether analog or digital, information is represented by the fundamental quantity in electrical engineering: the **signal**. Stated in mathematical terms, *a signal is merely a function*. Analog signals are continuous-valued; digital signals are discrete-valued. The independent variable of the signal could be time (speech, for example), space (images), or the integers (denoting the sequencing of letters and numbers in the football score).

1.1.1 Analog Signals

Analog signals are usually signals defined over continuous independent variable(s). Speech[1] is produced by your vocal cords exciting acoustic resonances in your vocal tract. The result is pressure waves propagating in the air, and the speech signal thus corresponds to a function having independent variables of space and time and a value corresponding to air pressure: $s(\mathbf{x}, t)$ (Here we use vector notation $\mathbf{x}$ to denote spatial coordinates). When you record someone talking, you are evaluating the speech signal at a particular spatial location, $\mathbf{x}_0$ say. An example of the resulting waveform $s(\mathbf{x}_0, t)$ is shown in this figure (Figure 1.1).

Photographs are static, and are continuous-valued signals defined over space. Black-and-white images have only one value at each point in space, which amounts to its optical reflection properties. In Figure 1.2, an image is shown, demonstrating that it (and all other images as well) are functions of two independent spatial variables.

Color images have values that express how reflectivity depends on the optical spectrum. Painters long ago found that mixing together combinations of the so-called primary colors–red, yellow and blue–can produce very realistic color images. Thus, images today are usually thought of as having three values at every point in space, but a different set of colors is used: How much of red, *green* and blue is present. Mathematically, color pictures are multivalued–vector-valued–signals: $s(\mathbf{x}) = (r(\mathbf{x}), g(\mathbf{x}), b(\mathbf{x}))^T$.

Interesting cases abound where the analog signal depends not on a continuous variable, such as time, but on a discrete variable. For example, temperature readings taken every hour have continuous–analog–values, but the signal's independent variable is (essentially) the integers.

---

[1]http://cnx.org/content/m0049/latest/

Figure 1.1: A speech signal's amplitude relates to tiny air pressure variations. Shown is a recording of the vowel "e" (as in "speech").

1.1.2 Digital Signals

The word "digital" means discrete-valued and implies the signal has an integer-valued independent variable. Digital information includes numbers and symbols (characters typed on the keyboard, for example). Computers rely on the digital representation of information to manipulate and transform information. Symbols do not have a numeric value, and each is represented by a unique number. The ASCII character code has the upper- and lowercase characters, the numbers, punctuation marks, and various other symbols represented by a seven-bit integer. For example, the ASCII code represents the letter a as the number 97 and the letter A as 65. Figure 1.3 shows the international convention on associating characters with integers.

1.2 Introduction to Systems

Signals are manipulated by systems. Mathematically, we represent what a system does by the notation $y(t) = S(x(t))$, with x representing the input signal and y the output signal.

This notation mimics the mathematical symbology of a function: A system's input is analogous to an independent variable and its output the dependent variable. For the mathematically inclined, a system is a **functional**: a function of a function (signals are functions).

Simple systems can be connected together–one system's output becomes another's input– to accomplish some overall design. Interconnection topologies can be quite complicated, but usually consist of weaves of three basic interconnection forms.

Lena

(a)

(b)

Figure 1.2: On the left is the classic *Lena* image, which is used ubiquitously as a test image. It contains straight and curved lines, complicated texture, and a face. On the right is a perspective display of the Lena image as a signal: a function of two spatial variables. The colors merely help show what signal values are about the same size. In this image, signal values range between 0 and 255; why is that?

Ascii Table

| num-ber	char-ac-ter	num-ber	char-ac-ter	num-ber	char-ac-ter	num-ber	char-ac-ter	num-ber	char-ac-ter	num-ber	char-ac-ter	num-ber	char-ac-ter	num-ber	char-ac-ter
00	nul	01	soh	02	stx	03	etx	04	eot	05	enq	06	ack	07	bel
08	bs	09	ht	0A	nl	0B	vt	0C	np	0D	cr	0E	so	0F	si
10	dle	11	dc1	12	dc2	13	dc3	14	dc4	15	nak	16	syn	17	etb
18	car	19	em	1A	sub	1B	esc	1C	fs	1D	gs	1E	rs	1F	us
20	sp	21	!	22	"	23		24	$	25	%	26	&	27	'
28	(	29	)	2A	*	2B	+	2C	,	2D	-	2E	.	2F	/
30	0	31	1	32	2	33	3	34	4	35	5	36	6	37	7
38	8	39	9	3A	:	3B	;	3C	<	3D	=	3E	>	3F	?
40	@	41	A	42	B	43	C	44	D	45	E	46	F	47	G
48	H	49	I	4A	J	4B	K	4C	L	4D	M	4E	N	4F	0
50	P	51	Q	52	R	53	S	54	T	55	U	56	V	57	W
58	X	59	Y	5A	Z	5B	[	5C	\	5D	]	5E	^	5F	_
60	'	61	a	62	b	63	c	64	d	65	e	66	f	67	g
68	h	69	i	6A	j	6B	k	6C	l	6D	m	6E	n	6F	o
70	p	71	q	72	r	73	s	74	t	75	u	76	v	77	w
78	x	79	y	7A	z	7B	{	7C	|	7D	}	7E	~	7F	del

Figure 1.3:   The ASCII translation table shows how standard keyboard characters are represented by integers. This table displays the so-called 7-bit code (how many characters in a seven-bit code?); extended ASCII has an 8-bit code. The numeric codes are represented in hexadecimal (base-16) notation. The mnemonic characters correspond to control characters, some of which may be familiar (like cr for carriage return) and some not (bel means a "bell").

Definition of a system

Figure 1.4:   The system depicted has input $x(t)$ and output $y(t)$. Mathematically, systems operate on function(s) to produce other function(s). In many ways, systems are like functions, rules that yield a value for the dependent variable (our output signal) for each value of its independent variable (its input signal). The notation $y(t) = S(x(t))$ corresponds to this block diagram. We term $S(\cdot)$ the input-output relation for the system.

Figure 1.5: The most rudimentary ways of interconnecting systems are shown in the figures in this section. This is the cascade configuration.

Figure 1.6: The parallel configuration.

1.2.1 Cascade Interconnection

The simplest form is when one system's output is connected only to another's input. Mathematically, $w(t) = S_1(x(t))$, and $y(t) = S_2(w(t))$, with the information contained in $x(t)$ processed by the first, then the second system. In some cases, the ordering of the systems matter, in others it does not. For example, in the fundamental model of communication [2] the ordering most certainly matters.

1.2.2 Parallel Interconnection

A signal $x(t)$ is routed to two (or more) systems, with this signal appearing as the input to all systems simultaneously and with equal strength. Block diagrams have the convention that signals going to more than one system are not split into pieces along the way. Two or more systems operate on $x(t)$ and their outputs are added together to create the output $y(t)$. Thus, $y(t) = S_1(x(t)) + S_2(x(t))$, and the information in $x(t)$ is processed separately by both systems.

1.2.3 Feedback Interconnection

The subtlest interconnection configuration has a system's output also contributing to its input. Engineers would say the output is "fed back" to the input through system 2, hence the terminology. The mathematical statement of the feedback interconnection (Figure 1.7)

---

[2]http://cnx.org/content/m0002/latest/#commsys

Figure 1.7:   The feedback configuration.

is that the feed-forward system produces the output: $y(t) = S_1(e(t))$. The input $e(t)$ equals the input signal minus the output of some other system's output to $y(t)$: $e(t) = x(t) - S_2(y(t))$. Feedback systems are omnipresent in control problems, with the error signal used to adjust the output to achieve some condition defined by the input (controlling) signal. For example, in a car's cruise control system, $x(t)$ is a constant representing what speed you want, and $y(t)$ is the car's speed as measured by a speedometer. In this application, system 2 is the identity system (output equals input).

1.3 Discrete-Time Signals and Systems

Mathematically, analog signals are functions having as their independent variables continuous quantities, such as space and time. Discrete-time signals are functions defined on the integers; they are sequences. As with analog signals, we seek ways of decomposing discrete-time signals into simpler components. Because this approach leading to a better understanding of signal structure, we can exploit that structure to represent information (create ways of representing information with signals) and to extract information (retrieve the information thus represented). For symbolic-valued signals, the approach is different: We develop a common representation of all symbolic-valued signals so that we can embody the information they contain in a unified way. From an information representation perspective, the most important issue becomes, for both real-valued and symbolic-valued signals, efficiency: what is the most parsimonious and compact way to represent information so that it can be extracted later.

1.3.1 Real- and Complex-valued Signals

A discrete-time signal is represented symbolically as $s(n)$, where $n = \{\ldots, -1, 0, 1, \ldots\}$.

We usually draw discrete-time signals as stem plots to emphasize the fact they are functions defined only on the integers. We can delay a discrete-time signal by an integer just as with analog ones. A delayed unit sample has the expression $\delta(n-m)$, and equals one when $n = m$.

1.3.2 Complex Exponentials

The most important signal is, of course, the **complex exponential sequence**.

$$s(n) = e^{j2\pi f n} \tag{1.1}$$

Figure 1.8: The discrete-time cosine signal is plotted as a stem plot. Can you find the formula for this signal?

Figure 1.9: The unit sample.

1.3.3 Sinusoids

Discrete-time sinusoids have the obvious form $s(n) = A\cos(2\pi f n + \phi)$. As opposed to analog complex exponentials and sinusoids that can have their frequencies be any real value, frequencies of their discrete-time counterparts yield unique waveforms *only* when f lies in the interval $\left(-\left(\frac{1}{2}\right), \frac{1}{2}\right]$. This property can be easily understood by noting that adding an integer to the frequency of the discrete-time complex exponential has no effect on the signal's value.

$$\begin{aligned} e^{j2\pi(f+m)n} &= e^{j2\pi f n}e^{j2\pi m n} \\ &= e^{j2\pi f n} \end{aligned} \tag{1.2}$$

This derivation follows because the complex exponential evaluated at an integer multiple of 2π equals one.

1.3.4 Unit Sample

The second-most important discrete-time signal is the **unit sample**, which is defined to be

$$\delta(n) = \begin{cases} 1 \text{ if } n = 0 \\ 0 \text{ otherwise} \end{cases} \tag{1.3}$$

Examination of a discrete-time signal's plot, like that of the cosine signal shown in Figure 1.8, reveals that all signals consist of a sequence of delayed and scaled unit samples.

Because the value of a sequence at each integer m is denoted by $s(m)$ and the unit sample delayed to occur at m is written $\delta(n-m)$, we can decompose *any* signal as a sum of unit samples delayed to the appropriate location and scaled by the signal value.

$$s(n) = \sum_{m=-\infty}^{\infty} \left(s(m)\,\delta(n-m) \right) \tag{1.4}$$

This kind of decomposition is unique to discrete-time signals, and will prove useful subsequently.

1.3.5 Symbolic Signals

An interesting aspect of discrete-time signals is that their values do not need to be real numbers. We do have real-valued discrete-time signals like the sinusoid, but we also have signals that denote the sequence of characters typed on the keyboard. Such characters certainly aren't real numbers, and as a collection of possible signal values, they have little mathematical structure other than that they are members of a set. More formally, each element of the *symbolic-valued* signal $s(n)$ takes on one of the values $\{a_1, \ldots, a_K\}$ which comprise the **alphabet** A. This technical terminology does not mean we restrict symbols to being members of the English or Greek alphabet. They could represent keyboard characters, bytes (8-bit quantities), integers that convey daily temperature. Whether controlled by software or not, discrete-time systems are ultimately constructed from digital circuits, which consist *entirely* of analog circuit elements. Furthermore, the transmission and reception of discrete-time signals, like e-mail, is accomplished with analog signals and systems. Understanding how discrete-time and analog signals and systems intertwine is perhaps the main goal of this course.

1.3.6 Discrete-Time Systems

Discrete-time systems can act on discrete-time signals in ways similar to those found in analog signals and systems. Because of the role of software in discrete-time systems, many more different systems can be envisioned and "constructed" with programs than can be with analog signals. In fact, a special class of analog signals can be converted into discrete-time signals, processed with software, and converted back into an analog signal, all without the incursion of error. For such signals, systems can be easily produced in software, with equivalent analog realizations difficult, if not impossible, to design.

1.4 Linear Time-Invariant Systems

A discrete-time signal $s(n)$ is **delayed** by n_0 samples when we write $s(n-n_0)$, with $n_0 > 0$. Choosing n_0 to be negative advances the signal along the integers. As opposed to analog delays[3], discrete-time delays can *only* be integer valued. In the frequency domain, delaying a signal corresponds to a linear phase shift of the signal's discrete-time Fourier transform: $\left(s(n-n_0) \leftrightarrow e^{-(j2\pi f n_0)} S\left(e^{j2\pi f}\right) \right)$.

    Linear discrete-time systems have the superposition property.

Superposition

$$S(a_1 x_1(n) + a_2 x_2(n)) = a_1 S(x_1(n)) + a_2 S(x_2(n)) \tag{1.5}$$

---

[3]http://cnx.org/content/m0006/latest/#delay

A discrete-time system is called **shift-invariant** (analogous to time-invariant analog systems[4]) if delaying the input delays the corresponding output.

Shift-Invariant

$$\text{If}\ \ S\left(x\left(n\right)\right)=y\left(n\right),\ \text{Then}\ \ S\left(x\left(n-n_0\right)\right)=y\left(n-n_0\right) \tag{1.6}$$

We use the term shift-invariant to emphasize that delays can only have integer values in discrete-time, while in analog signals, delays can be arbitrarily valued.

We want to concentrate on systems that are both linear and shift-invariant. It will be these that allow us the full power of frequency-domain analysis and implementations. Because we have no physical constraints in "constructing" such systems, we need only a mathematical specification. In analog systems, the differential equation specifies the input-output relationship in the time-domain. The corresponding discrete-time specification is the **difference equation**.

The Difference Equation

$$y\left(n\right)=a_1y\left(n-1\right)+\cdots+a_py\left(n-p\right)+b_0x\left(n\right)+b_1x\left(n-1\right)+\cdots+b_qx\left(n-q\right) \tag{1.7}$$

Here, the output signal $y\left(n\right)$ is related to its *past* values $y\left(n-l\right)$, $l=\{1,\ldots,p\}$, and to the current and past values of the input signal $x\left(n\right)$. The system's characteristics are determined by the choices for the number of coefficients p and q and the coefficients' values $\{a_1,\ldots,a_p\}$ and $\{b_0,b_1,\ldots,b_q\}$.

> ASIDE: There is an asymmetry in the coefficients: where is a_0 ? This coefficient would multiply the $y\left(n\right)$ term in the difference equation (Equation 1.7). We have essentially divided the equation by it, which does not change the input-output relationship. We have thus created the convention that a_0 is always one.

As opposed to differential equations, which only provide an *implicit* description of a system (we must somehow solve the differential equation), difference equations provide an *explicit* way of computing the output for any input. We simply express the difference equation by a program that calculates each output from the previous output values, and the current and previous inputs.

1.5 Discrete-Time Convolution

1.5.1 Overview

Convolution is a concept that extends to all systems that are both **linear and time-invariant**[5] (**LTI**). The idea of **discrete-time convolution** is exactly the same as that of continuous-time convolution[6]. For this reason, it may be useful to look at both versions to help your understanding of this extremely important concept. Recall that convolution is a very powerful tool in determining a system's output from knowledge of an arbitrary input and the system's impulse response. It will also be helpful to see convolution graphically with your own eyes and to play around with it some, so experiment with the applets[7] available on the internet. These resources will offer different approaches to this crucial concept.

---

[4]http://cnx.org/content/m0007/latest/#timeinv
[5]http://cnx.org/content/m10084/latest/
[6]http://cnx.org/content/m10085/latest/
[7]http://www.jhu.edu/~signals

1.5.2 Convolution Sum

As mentioned above, the convolution sum provides a concise, mathematical way to express the output of an LTI system based on an arbitrary discrete-time input signal and the system's response. The **convolution sum** is expressed as

$$y[n] = \sum_{k=-\infty}^{\infty} (x[k]\,h[n-k]) \tag{1.8}$$

As with continuous-time, convolution is represented by the symbol *, and can be written as

$$y[n] = x[n] * h[n] \tag{1.9}$$

By making a simple change of variables into the convolution sum, $k = n - k$, we can easily show that convolution is **commutative**:

$$x[n] * h[n] = h[n] * x[n] \tag{1.10}$$

For more information on the characteristics of convolution, read about the Properties of Convolution[8].

1.5.3 Derivation

We know that any discrete-time signal can be represented by a summation of scaled and shifted discrete-time impulses. Since we are assuming the system to be linear and time-invariant, it would seem to reason that an input signal comprised of the sum of scaled and shifted impulses would give rise to an output comprised of a sum of scaled and shifted impulse responses. This is exactly what occurs in **convolution**. Below we present a more rigorous and mathematical look at the derivation:

Letting $\mathscr{H}$ be a DT LTI system, we start with the following equation and work our way down the convolution sum!

$$
\begin{aligned}
y[n] &= \mathscr{H}[x[n]] \\
&= \mathscr{H}\left[\sum_{k=-\infty}^{\infty} (x[k]\,\delta[n-k])\right] \\
&= \sum_{k=-\infty}^{\infty} (\mathscr{H}[x[k]\,\delta[n-k]]) \\
&= \sum_{k=-\infty}^{\infty} (x[k]\,\mathscr{H}[\delta[n-k]]) \\
&= \sum_{k=-\infty}^{\infty} (x[k]\,h[n-k])
\end{aligned} \tag{1.11}
$$

Let us take a quick look at the steps taken in the above derivation. After our initial equation, we using the DT sifting property[9] to rewrite the function, $x[n]$, as a sum of the function times the unit impulse. Next, we can move around the $\mathscr{H}$ operator and the summation because $\mathscr{H}\{\}$ is a linear, DT system. Because of this linearity and the fact that $x[k]$ is a constant, we can pull the previous mentioned constant out and simply multiply it by $\mathscr{H}\{\}$. Finally, we use the fact that $\mathscr{H}\{\}$ is time invariant in order to reach our final state - the convolution sum!

A quick graphical example may help in demonstrating why convolution works.

1.5.4 Convolution Through Time (A Graphical Approach)

In this section we will develop a second graphical interpretation of discrete-time convolution. We will begin this by writing the convolution sum allowing x to be a causal, length-m signal

---

[8] http://cnx.org/content/m10088/latest/
[9] http://cnx.org/content/m10059/latest/#sifting

Figure 1.10: A single impulse input yields the system's impulse response.
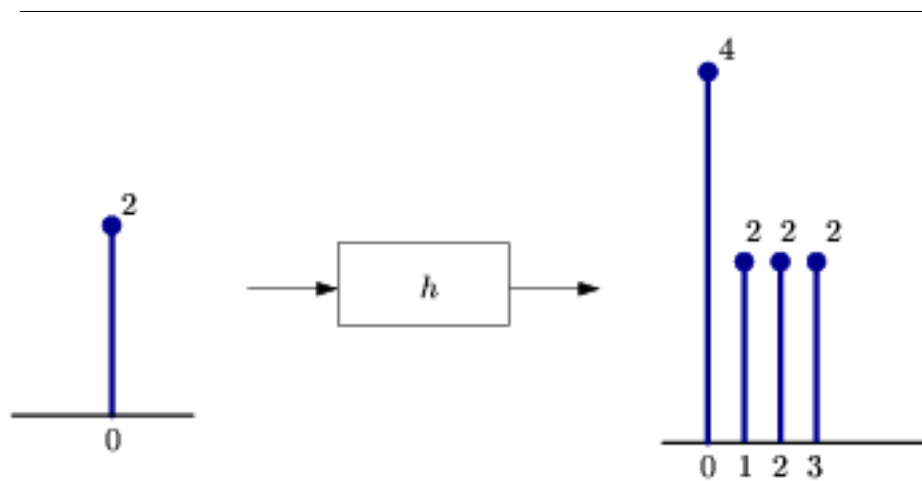
Figure 1.11: A scaled impulse input yields a scaled response, due to the scaling property of the system's linearity.

Figure 1.12: We now use the time-invariance property of the system to show that a delayed input results in an output of the same shape, only delayed by the same amount as the input.

Figure 1.13: We now use the additivity portion of the linearity property of the system to complete the picture. Since any discrete-time signal is just a sum of scaled and shifted discrete-time impulses, we can find the output from knowing the input and the impulse response.

Figure 1.14: This is the end result that we are looking to find.

and h to be a causal, length-k, LTI system. This gives us the finite summation,

$$y[n] = \sum_{l=0}^{m-1} (x[l] h[n-l]) \tag{1.12}$$

Notice that for any given n we have a sum of the products of x_l and a time-delayed h_{-l}. This is to say that we multiply the terms of x by the terms of a time-reversed h and add them up.

Going back to the previous example:

What we are doing in the above demonstration is reversing the impulse response in time and "walking it across" the input signal. Clearly, this yields the same result as scaling, shifting and summing impulse responses.

This approach of time-reversing, and sliding across is a common approach to presenting convolution, since it demonstrates how convolution builds up an output through time.

1.6 Review of Linear Algebra

Vector spaces are the principal object of study in linear algebra. A vector space is always defined with respect to a field of scalars.

1.6.1 Fields

A field is a set F equipped with two operations, addition and mulitplication, and containing two special members 0 and 1 ($0 \neq 1$), such that for all $\{a, b, c\} \in F$

1. (a) $a + b \in F$

Figure 1.15: Here we reverse the impulse response, h , and begin its traverse at time 0.
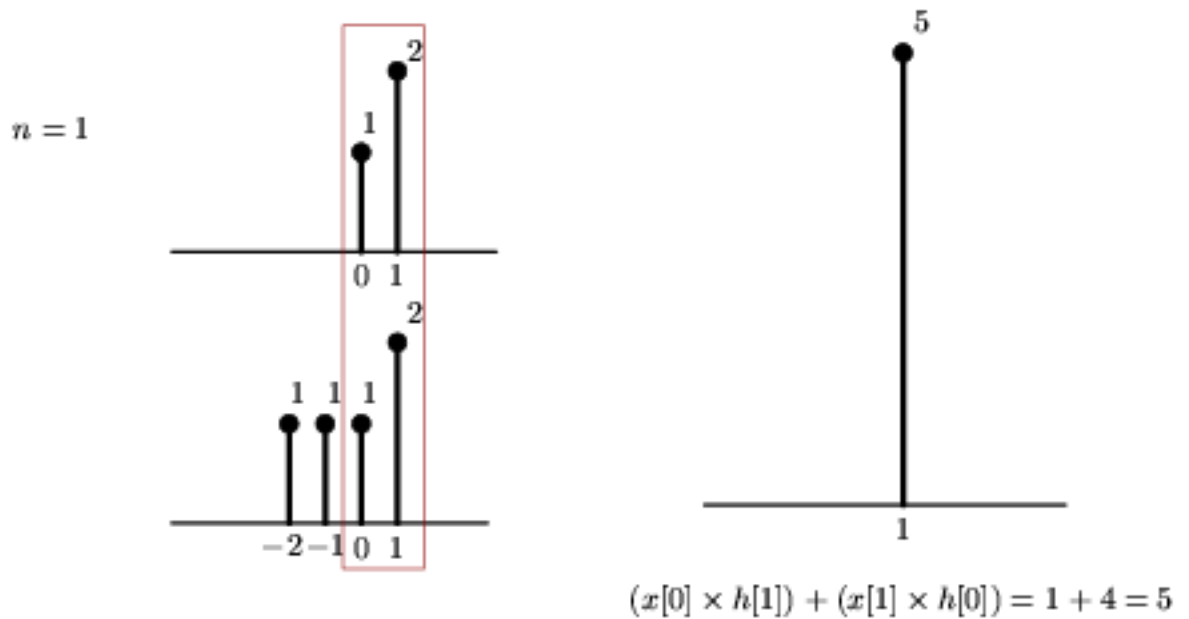
Figure 1.16: We continue the traverse. See that at time 1 , we are multiplying two elements of the input signal by two elements of the impulse response.

$n = 2$

$$(x[0] \times h[2]) + (x[1] \times h[1]) = 1 + 2 = 3$$
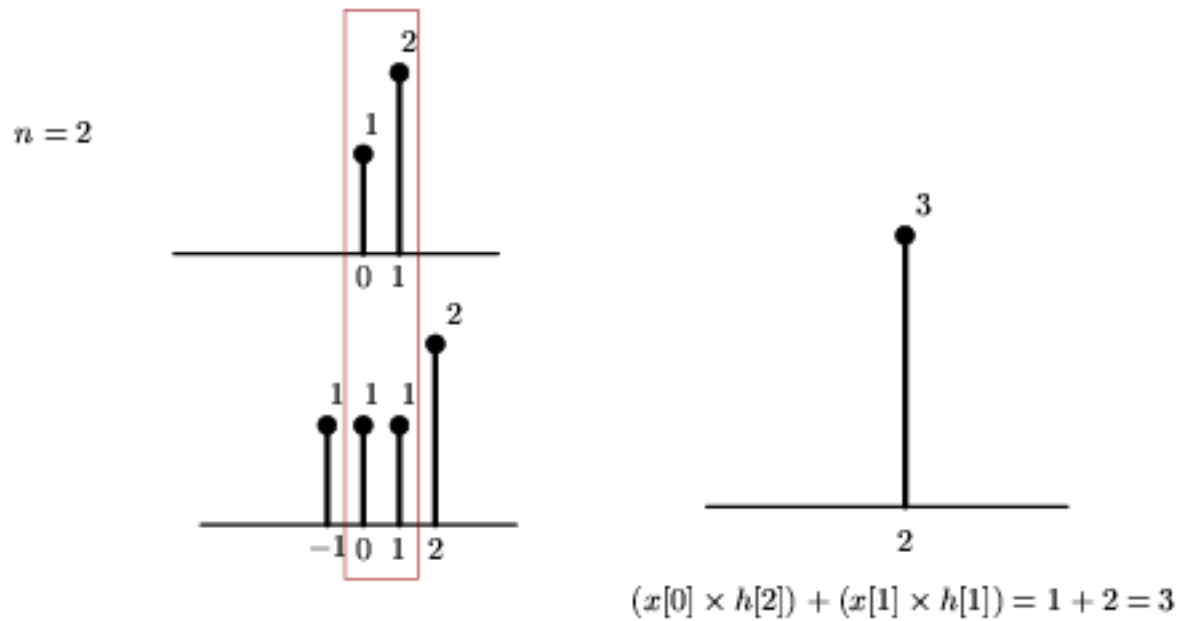
Figure 1.17

    (b) $a + b = b + a$

    (c) $(a + b) + c = a + (b + c)$

    (d) $a + 0 = a$

    (e) there exists $-a$ such that $a + (-a) = 0$

2. (a) $ab \in F$

    (b) $ab = ba$

    (c) $(ab)\,c = a\,(bc)$

    (d) $a \cdot 1 = a$

    (e) there exists a^{-1} such that $aa^{-1} = 1$

3. $a\,(b + c) = ab + ac$

More concisely

1. F is an abelian group under addition

2. F is an abelian group under multiplication

3. multiplication distributes over addition

1.6.1.1 Examples

$\mathbb{Q}$, R, C

Figure 1.18:   If we follow this through to one more step, $n = 4$, then we can see that we produce the same output as we saw in the initial example.

1.6.2 Vector Spaces

Let F be a field, and V a set. We say V *is a vector space over F* if there exist two operations, defined for all $a \in F$, $\mathbf{u} \in V$ and $\mathbf{v} \in V$:

- vector addition: $(\mathbf{u}, \mathbf{v}) \to \mathbf{u} + \mathbf{v} \in V$

- scalar multiplication: $(a,\mathbf{v}) \to a\mathbf{v} \in V$

and if there exists an element denoted $\mathbf{0} \in V$, such that the following hold for all $a \in F$, $b \in F$, and $\mathbf{u} \in V$, $\mathbf{v} \in V$, and $\mathbf{w} \in V$

1.   (a) $\mathbf{u} + (\mathbf{v} + \mathbf{w}) = (\mathbf{u} + \mathbf{v}) + \mathbf{w}$

    (b) $\mathbf{u} + \mathbf{v} = \mathbf{v} + \mathbf{u}$

    (c) $\mathbf{u} + \mathbf{0} = \mathbf{u}$

    (d)   there exists $-\mathbf{u}$ such that $\mathbf{u} + (-\mathbf{u}) = \mathbf{0}$

2.   (a) $a\left(\mathbf{u} + \mathbf{v}\right) = a\mathbf{u} + a\mathbf{v}$

    (b) $\left(a + b\right)\mathbf{u} = a\mathbf{u} + b\mathbf{u}$

    (c) $\left(ab\right)\mathbf{u} = a\left(b\mathbf{u}\right)$

    (d) $1 \cdot \mathbf{u} = \mathbf{u}$

More concisely,

1. V is an abelian group under plus

2. Natural properties of scalar multiplication

1.6.2.1 Examples

- $\mathbb{R}^N$ is a vector space over R

- C^N is a vector space over C

- C^N is a vector space over R

- $\mathbb{R}^N$ is *not* a vector space over C

The elements of V are called **vectors**.

1.6.3 Euclidean Space

Throughout this course we will think of a signal as a vector

$$\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_N \end{pmatrix} = \begin{pmatrix} x_1 & x_2 & \ldots & x_N \end{pmatrix}^T$$

The samples $\{x_i\}$ could be samples from a finite duration, continuous time signal, for example.

    A signal will belong to one of two vector spaces:

Figure 1.19: S is any line through the origin.

1.6.3.1 Real Euclidean space

$\mathbf{x} \in \mathbb{R}^N$ (over R)

1.6.3.2 Complex Euclidean space

$\mathbf{x} \in C^N$ (over C)

1.6.4 Subspaces

Let V be a vector space over F.

A subset $S \subseteq V$ is called a **subspace** of V if S is a vector space over F in its own right.

Example 1.1:

$V = \mathbb{R}^2$, $F = \mathbb{R}$, $S =$ any line though the origin.

Are there other subspaces?

Theorem 1.1:

$S \subseteq V$ is a subspace if and only if for all $a \in F$ and $b \in F$ and for all $\mathbf{s} \in S$ and $\mathbf{t} \in S$, $a\mathbf{s} + b\mathbf{t} \in S$

1.6.5 Linear Independence

Let $\mathbf{u}_1, \ldots, \mathbf{u}_k \in V$.

We say that these vectors are **linearly dependent** if there exist scalars $a_1, \ldots, a_k \in F$ such that

$$\sum_{i=1}^{k} (a_i u_i) = \mathbf{0} \tag{1.13}$$

and at least one $a_i \neq 0$.

If Equation 1.13 only holds for the case $a_1 = \cdots = a_k = 0$, we say that the vectors are **linearly independent**.

Example 1.2:

Figure 1.20: $< S >$ is the xy-plane.

$$1 \begin{pmatrix} 1 \\ -1 \\ 2 \end{pmatrix} - 2 \begin{pmatrix} -2 \\ 3 \\ 0 \end{pmatrix} + 1 \begin{pmatrix} -5 \\ 7 \\ -2 \end{pmatrix} = \mathbf{0}$$

so these vectors are linearly dependent in $\mathbb{R}^3$.

1.6.6 Spanning Sets

Consider the subset $S = \{v_1, v_2, \ldots, v_k\}$. Define the **span** of S

$$< S > \equiv \mathrm{span}\,(S) \equiv \left\{ \sum_{i=1}^{k} (a_i v_i) \, | a_i \in F \right\}$$

Fact: $< S >$ is a subspace of V.

Example 1.3:

$V = \mathbb{R}^3$, $F = \mathbb{R}$, $S = \{v_1, v_2\}$, $v_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$, $v_2 = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \Rightarrow\, < S >=$ xy-plane.

1.6.6.1 Aside

If S is infinite, the notions of linear independence and span are easily generalized:

We say S is linearly independent if, for every finite collection $u_1, \ldots, u_k \in S$, (k arbitrary) we have

$$\sum_{i=1}^{k} (a_i u_i) = \mathbf{0} \Rightarrow a_i = 0 \ \ ,$$

The span of S is

$$< S >= \left\{ \sum_{i=1}^{k} (a_i u_i) \, | a_i \in F \ \ u_i \in S \ \ k < \infty \right\}$$

NOTE: In both definitions, we only consider **finite** sums.

1.6.7 Bases

A set $B \subseteq V$ is called a **basis** for V over F if and only if

1. B is linearly independent

2. $< B >= V$

Bases are of fundamental importance in signal processing. They allow us to decompose a signal into building blocks (basis vectors) that are often more easily understood.

Example 1.4:

$V = $ (real or complex) Euclidean space, $\mathbb{R}^N$ or C^N.

$$B = \{e_1, \ldots, e_N\} \equiv \text{standard basis}$$

$$e_i = \begin{pmatrix} 0 \\ \vdots \\ 1 \\ \vdots \\ 0 \end{pmatrix}$$

where the 1 is in the i^{th} position.

Example 1.5:

$V = C^N$ over C.

$$B = \{u_1, \ldots, u_N\}$$

which is the DFT basis.

$$u_k = \begin{pmatrix} 1 \\ e^{-\left(j2\pi \frac{k}{N}\right)} \\ \vdots \\ e^{-\left(j2\pi \frac{k}{N}(N-1)\right)} \end{pmatrix}$$

where $j = \sqrt{-1}$.

1.6.7.1 Key Fact

If B is a basis for V, then every $\mathbf{v} \in V$ can be written uniquely (up to order of terms) in the form

$$\mathbf{v} = \sum_{i=1}^{N} (a_i v_i)$$

where $a_i \in F$ and $v_i \in B$.

1.6.7.2 Other Facts

- If S is a linearly independent set, then S can be extended to a basis.

- If $< S >= V$, then S contains a basis.

1.6.8 Dimension

Let V be a vector space with basis B. The dimension of V, denoted $\dim(V)$, is the cardinality of B.

Theorem 1.2:

Every vector space has a basis.

Theorem 1.3:

Every basis for a vector space has the same cardinality.

$\Rightarrow \dim(V)$ is **well-defined**.
If $\dim(V) < \infty$, we say V is **finite dimensional**.

1.6.8.1 Examples

vector space	field of scalars	dimension
$\mathbb{R}^N$	$\mathbb{R}$	
C^N	C	
C^N	$\mathbb{R}$	

Every subspace is a vector space, and therefore has its own dimension.

Example 1.6:

Suppose $S = \{u_1, \ldots, u_k\} \subseteq V$ is a linearly independent set. Then

$$\dim(<S>) =$$

Facts

- If S is a subspace of V, then $\dim(S) \le \dim(V)$.

- If $\dim(S) = \dim(V) < \infty$, then $S = V$.

1.6.9 Direct Sums

Let V be a vector space, and let $S \subseteq V$ and $T \subseteq V$ be subspaces.

We say V is the **direct sum** of S and T, written $V = (S \oplus T)$, if and only if for every $\mathbf{v} \in V$, there exist unique $\mathbf{s} \in S$ and $\mathbf{t} \in T$ such that $\mathbf{v} = \mathbf{s} + \mathbf{t}$.

If $V = (S \oplus T)$, then T is called a **complement** of S.

Example 1.7:

$$V = C' = left\{f : \mathbb{R} \to \mathbb{R} | f \text{is continuous} right\}$$

$$S = \text{even funcitons in} C'$$

$$T = \text{odd funcitons in} C'$$

$$f(t) = \frac{1}{2}(f(t) + f(-t)) + \frac{1}{2}(f(t) - f(-t))$$

If $f = g + h = g' + h'$, $g \in S$ and $g' \in S$, $h \in T$ and $h' \in T$, then $g - g' = h' - h$ is odd and even, which implies $g = g'$ and $h = h'$.

1.6.9.1 Facts

1. Every subspace has a complement

2. $V = (S \oplus T)$ if and only if

   (a) $S \bigcap T = \{\mathbf{0}\}$
   (b) $< S, T > = V$

3. If $V = (S \oplus T)$, and $\dim (V) < \infty$, then $\dim (V) = \dim (S) + \dim (T)$

1.6.9.2 Proofs

Invoke a basis.

1.6.10 Norms

Let V be a vector space over F. A norm is a mapping $(V \to F)$, denoted by $\| \cdot \|$, such that forall $\mathbf{u} \in V$, $\mathbf{v} \in V$, and $\lambda \in F$

1. $\| \mathbf{u} \| > 0$ if $\mathbf{u} \neq \mathbf{0}$

2. $\| \lambda \mathbf{u} \| = |\lambda| \| \mathbf{u} \|$

3. $\| \mathbf{u} + \mathbf{v} \| \leq \| \mathbf{u} \| + \| \mathbf{v} \|$

1.6.10.1 Examples

Euclidean norms:
   $\mathbf{x} \in \mathbb{R}^N$:

$$\| \mathbf{x} \| = \left(\sum_{i=1}^{N} \left(x_i{}^2 \right) \right)^{\frac{1}{2}}$$

$\mathbf{x} \in C^N$:

$$\| \mathbf{x} \| = \left(\sum_{i=1}^{N} \left((|x_i|)^2 \right) \right)^{\frac{1}{2}}$$

1.6.10.2 Induced Metric

Every norm induces a metric on V

$$d (\mathbf{u}, \mathbf{v}) \equiv \| \mathbf{u} - \mathbf{v} \|$$

which leads to a notion of "distance" between vectors.

1.6.11 Inner products

Let V be a vector space over F, $F = \mathbb{R}$ or C. An inner product is a mapping $V \times V \to F$, denoted $\cdot \ \cdot$, such that

1. $\mathbf{v} \cdot \mathbf{v} \geq 0$, and $(\mathbf{v} \cdot \mathbf{v} = 0 \Leftrightarrow \mathbf{v} = 0)$

2. $\mathbf{u} \cdot \mathbf{v} = \mathbf{v} \cdot \mathbf{u}^*$

3. $a\mathbf{u} + b\mathbf{v} \cdot \mathbf{w} = a\mathbf{u} \cdot \mathbf{w} + b\mathbf{v} \cdot \mathbf{w}$

1.6.11.1 Examples

$\mathbb{R}^N$ over R:

$$\mathbf{x} \cdot \mathbf{y} = \mathbf{x}^T \mathbf{y} = \sum_{i=1}^{N} (x_i y_i)$$

C^N over C:

$$\mathbf{x} \cdot \mathbf{y} = \mathbf{x}^H \mathbf{y} = \sum_{i=1}^{N} (x_i{}^* y_i)$$

If $\mathbf{x} = (x_1, \ldots, x_N)^T \in C$, then

$$\mathbf{x}^H \equiv \begin{pmatrix} x_1{}^* \\ \vdots \\ x_N{}^* \end{pmatrix}^T$$

is called the "Hermitian," or "conjugate transpose" of $\mathbf{x}$.

1.6.12 Triangle Inequality

If we define $\| \mathbf{u} \| = \mathbf{u} \cdot \mathbf{u}$, then

$$\| \mathbf{u} + \mathbf{v} \| \leq \| \mathbf{u} \| + \| \mathbf{v} \|$$

Hence, every inner product induces a norm.

1.6.13 Cauchy-Schwarz Inequality

For all $\mathbf{u} \in V$, $\mathbf{v} \in V$,

$$|\mathbf{u} \cdot \mathbf{v}| \leq \| \mathbf{u} \| \| \mathbf{v} \|$$

In inner product spaces, we have a notion of the angle between two vectors:

$$\angle (\mathbf{u}, \mathbf{v}) = \arccos \left(\frac{\mathbf{u} \cdot \mathbf{v}}{\| \mathbf{u} \| \| \mathbf{v} \|} \right) \in [0, 2\pi)$$

1.6.14 Orthogonality

$\mathbf{u}$ and $\mathbf{v}$ are **orthogonal** if

$$\mathbf{u} \cdot \mathbf{v} = 0$$

Notation: $(\mathbf{u} \perp \mathbf{v})$.

If in addition $\| \mathbf{u} \| = \| \mathbf{v} \| = 1$, we say $\mathbf{u}$ and $\mathbf{v}$ are **orthonormal**.

In an orthogonal (orthonormal) *set*, each pair of vectors is orthogonal (orthonormal).

1.6.15 Orthonormal Bases

An Orthonormal basis is a basis $\{v_i\}$ such that

$$v_i \cdot v_i = \delta_{ij} = \begin{cases} 1 \text{ if } i = j \\ 0 \text{ if } i \neq j \end{cases}$$

Figure 1.21: Orthogonal vectors in $\mathbb{R}^2$.

Example 1.8:

The standard basis for $\mathbb{R}^N$ or C^N

Example 1.9:

The normalized DFT basis

$$u_k = \frac{1}{\sqrt{N}} \begin{pmatrix} 1 \\ e^{-\left(j2\pi\frac{k}{N}\right)} \\ \vdots \\ e^{-\left(j2\pi\frac{k}{N}(N-1)\right)} \end{pmatrix}$$

1.6.16 Expansion Coefficients

If the representation of $\mathbf{v}$ with respect to $\{v_i\}$ is

$$\mathbf{v} = \sum (a_i v_i)$$

then

$$a_i = v_i \ \cdot \ \mathbf{v}$$

1.6.17 Gram-Schmidt

Every inner product space has an orthonormal basis. Any (countable) basis can be made orthogonal by the Gram-Schmidt orthogonalization process.

1.6.18 Orthogonal Compliments

Let $S \subseteq V$ be a subspace. The **orthogonal compliment** S is

$$S^\perp = \{\, \mathbf{u} \, | \mathbf{u} \in V \ \ \mathbf{u} \ \cdot \ \mathbf{v} = 0 \ \ \mathbf{v} \in S \ \ , \ \ \}$$

$S^\perp$ is easily seen to be a subspace.
  If $\dim(v) < \infty$, then $V = \left(S \oplus S^\perp\right)$.

ASIDE: If $\dim(v) = \infty$, then in order to have $V = \left(S \oplus S^\perp\right)$ we require V to be a **Hilbert Space**.

1.6.19 Linear Transformations

Loosely speaking, a linear transformation is a mapping from one vector space to another that *preserves* vector space operations.

More precisely, let V, W be vector spaces over the same field F. A **linear transformation** is a mapping $T : V \to W$ such that

$$T(a\mathbf{u} + b\mathbf{v}) = aT(\mathbf{u}) + bT(\mathbf{v})$$

for all $a \in F$, $b \in F$ and $\mathbf{u} \in V$, $\mathbf{v} \in V$.

In this class we will be concerned with linear transformations between (real or complex) *Euclidean spaces*, or subspaces thereof.

1.6.20 Image

$$\text{image}(T) = \{\, \mathbf{w} \mid \mathbf{w} \in W \ \ T(\mathbf{v}) = \mathbf{w} \text{for some} \mathbf{v} \,\}$$

1.6.21 Nullspace

Also known as the kernel:

$$\ker(T) = \{\, \mathbf{v} \mid \mathbf{v} \in V \ \ T(\mathbf{v}) = \mathbf{0} \,\}$$

Both the image and the nullspace are easily seen to be subspaces.

1.6.22 Rank

$$\text{rank}(T) = \dim(\text{image}(T))$$

1.6.23 Nullity

$$\text{null}(T) = \dim(\ker(T))$$

1.6.24 Rank plus nullity theorem

$$\text{rank}(T) + \text{null}(T) = \dim(V)$$

Figure 1.22

1.6.25 Matrices

Every linear transformation T has a **matrix representation**. If $T : E^N \to E^M$, $E = \mathbb{R}$ or C, then T is represented by an $M \times N$ matrix

$$A = \begin{pmatrix} a_{11} & \dots & a_{1N} \\ \vdots & \ddots & \vdots \\ a_{M1} & \dots & a_{MN} \end{pmatrix}$$

where $(a_{1i}, \dots, a_{Mi})^T = T(e_i)$ and $e_i = (0, \dots, 1, \dots, 0)^T$ is the i^{th} **standard basis** vector.

> ASIDE: A linear transformation can be represented with respect to any bases of E^N and E^M, leading to a different A. We will always represent a linear transformation using the standard bases.

1.6.26 Column span

$$\text{colspan}\,(A) = <A> = \text{image}\,(A)$$

1.6.27 Duality

If $A : \mathbb{R}^N \to \mathbb{R}^M$, then

$$\ker^{\perp}(A) = \text{image}\left(A^T\right)$$

If $A : C^N \to C^M$, then

$$\ker^{\perp}(A) = \text{image}\left(A^H\right)$$

1.6.28 Inverses

The linear transformation/matrix A is **invertible** if and only if there exists a matrix B such that $AB = BA = I$ (identity).

Only *square* matrices can be invertible.

Theorem 1.4:

Let $A : F^N \to F^N$ be linear, $F = \mathbb{R}$ or C. The following are equivalent:

1. A is invertible (nonsingular)
2. $\text{rank}(A) = N$
3. $\text{null}(A) = 0$
4. $\det A \neq 0$
5. The columns of A form a basis.

If $A^{-1} = A^T$ (or A^H in the complex case), we say A is **orthogonal** (or **unitary**).

1.7 Hilbert Spaces

1.7.1 Hilbert Spaces

A vector space S with a valid inner product[10] defined on it is called an **inner product space**, which is also a **normed linear space**. A **Hilbert space** is an inner product space that is complete with respect to the norm defined using the inner product. Hilbert spaces are named after David Hilbert[11], who developed this idea through his studies of integral equations. We define our valid norm using the inner product as:

$$\| \mathbf{x} \| = \sqrt{\mathbf{x} \cdot \mathbf{x}} \tag{1.14}$$

Hilbert spaces are useful in studying and generalizing the concepts of Fourier expansion, Fourier transforms, and are very important to the study of quantum mechanics. Hilbert spaces are studied under the functional analysis branch of mathematics.

1.7.1.1 Examples of Hilbert Spaces

Below we will list a few examples of Hilbert spaces[12]. You can verify that these are valid inner products at home.

- For C^n,

$$\mathbf{x} \cdot \mathbf{y} = \mathbf{y}^T \mathbf{x} = \begin{pmatrix} y_0^* & y_1^* & \cdots & y_{n-1}^* \end{pmatrix} \begin{pmatrix} x_0 \\ x_1 \\ \vdots \\ x_{n-1} \end{pmatrix} = \sum_{i=0}^{n-1} (x_i y_i^*)$$

- Space of finite energy complex functions: $L^2(\mathbb{R})$

$$\mathbf{f} \cdot \mathbf{g} = \int_{-\infty}^{\infty} f(t) g(t)^* dt$$

---

[10] http://cnx.org/content/m10755/latest/
[11] http://www-history.mcs.st-andrews.ac.uk/history/Mathematicians/Hilbert.html
[12] http://cnx.org/content/m10434/latest/

- Space of square-summable sequences: $\ell^2\left(\mathbb{Z}\right)$

$$\mathbf{x} \,\cdot\, \mathbf{y} = \sum_{i=-\infty}^{\infty} \left(x\left[i\right] y\left[i\right]^* \right)$$

1.8 Signal Expansions

1.8.1 Main Idea

When working with signals many times it is helpful to break up a signal into smaller, more manageable parts. Hopefully by now you have been exposed to the concept of eigenvectors[13] and there use in decomposing a signal into one of its possible basis. By doing this we are able to simplify our calculations of signals and systems through eigenfunctions of LTI systems[14].

Now we would like to look at an alternative way to represent signals, through the use of **orthonormal basis**. We can think of orthonormal basis as a set of building blocks we use to construct functions. We will build up the signal/vector as a weighted sum of basis elements.

Example 1.10:

The complex sinusoids $\frac{1}{\sqrt{T}}e^{j\omega_0 nt}$ for all $-\infty < n < \infty$ form an orthonormal basis for $L^2\left(\left[0,T\right]\right)$.

In our Fourier series[15] equation, $f\left(t\right) = \sum_{n=-\infty}^{\infty}\left(c_n e^{j\omega_0 nt}\right)$, the $\{c_n\}$ are just another representation of $f\left(t\right)$.

NOTE: For signals/vectors in a Hilbert Space[16], the expansion coefficients are easy to find.

1.8.2 Alternate Representation

Recall our definition of a **basis**: A set of vectors $\{b_i\}$ in a vector space S is a basis if

1. The b_i are linearly independent.

2. The b_i span[17] S. That is, we can find $\{\alpha_i\}$, where $\alpha_i \in C$ (scalars) such that

$$\mathbf{x} = \sum_i \left(\alpha_i b_i\right) \quad , \quad x \in S \tag{1.15}$$

where $\mathbf{x}$ is a vector in S, α is a scalar in C, and $\mathbf{b}$ is a vector in S.

Condition 2 in the above definition says we can **decompose** any vector in terms of the $\{b_i\}$. Condition 1 ensures that the decomposition is **unique** (think about this at home).

NOTE: The $\{\alpha_i\}$ provide an alternate representation of $\mathbf{x}$.

---

[13]http://cnx.org/content/m10736/latest/
[14]http://cnx.org/content/m10500/latest/
[15]http://cnx.org/content/m10496/latest/
[16]http://cnx.org/content/m10755/latest/#sec2
[17]http://cnx.org/content/m10734/latest/#span_sec

Example 1.11:

Let us look at simple example in $\mathbb{R}^2$, where we have the following vector:

$$\mathbf{x} = \begin{pmatrix} 1 \\ 2 \end{pmatrix}$$

Standard Basis: $\{e_0, e_1\} = \left\{ (1,0)^T, (0,1)^T \right\}$

$$\mathbf{x} = e_0 + 2e_1$$

Alternate Basis: $\{h_0, h_1\} = \left\{ (1,1)^T, (1,-1)^T \right\}$

$$\mathbf{x} = \frac{3}{2} h_0 + \frac{-1}{2} h_1$$

In general, given a basis $\{b_0, b_1\}$ and a vector $\mathbf{x} \in \mathbb{R}^2$, how do we find the α_0 and α_1 such that

$$\mathbf{x} = \alpha_0 b_0 + \alpha_1 b_1 \tag{1.16}$$

1.8.3 Finding the Alphas

Now let us address the question posed above about finding α_i's in general for $\mathbb{R}^2$. We start by rewriting Equation 1.16 so that we can stack our b_i's as columns in a $2{\times}2$ matrix.

$$\begin{pmatrix} \mathbf{x} \end{pmatrix} = \alpha_0 \begin{pmatrix} b_0 \end{pmatrix} + \alpha_1 \begin{pmatrix} b_1 \end{pmatrix} \tag{1.17}$$

$$\begin{pmatrix} \mathbf{x} \end{pmatrix} = \begin{pmatrix} \vdots & \vdots \\ b_0 & b_1 \\ \vdots & \vdots \end{pmatrix} \begin{pmatrix} \alpha_0 \\ \alpha_1 \end{pmatrix} \tag{1.18}$$

Example 1.12:

Here is a simple example, which shows a little more detail about the above equations.

$$\begin{aligned} \begin{pmatrix} x\,[0] \\ x\,[1] \end{pmatrix} &= \alpha_0 \begin{pmatrix} b_0\,[0] \\ b_0\,[1] \end{pmatrix} + \alpha_1 \begin{pmatrix} b_1\,[0] \\ b_1\,[1] \end{pmatrix} \\ &= \begin{pmatrix} \alpha_0 b_0\,[0] + \alpha_1 b_1\,[0] \\ \alpha_0 b_0\,[1] + \alpha_1 b_1\,[1] \end{pmatrix} \end{aligned} \tag{1.19}$$

$$\begin{pmatrix} x\,[0] \\ x\,[1] \end{pmatrix} = \begin{pmatrix} b_0\,[0] & b_1\,[0] \\ b_0\,[1] & b_1\,[1] \end{pmatrix} \begin{pmatrix} \alpha_0 \\ \alpha_1 \end{pmatrix} \tag{1.20}$$

1.8.3.1 Simplifying our Equation

To make notation simpler, we define the following two items from the above equations:

- **Basis Matrix**:

$$B = \begin{pmatrix} \vdots & \vdots \\ b_0 & b_1 \\ \vdots & \vdots \end{pmatrix}$$

- **Coefficient Vector**:

$$\alpha = \begin{pmatrix} \alpha_0 \\ \alpha_1 \end{pmatrix}$$

This gives us the following, concise equation:

$$\mathbf{x} = B\alpha \tag{1.21}$$

which is equivalent to $\mathbf{x} = \sum_{i=0}^{1} (\alpha_i b_i)$.

Example 1.13:

Given a standard basis, $\left\{ \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right\}$, then we have the following basis matrix:

$$B = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

To get the α_i's, we solve for the coefficient vector in Equation 1.21

$$\alpha = B^{-1}\mathbf{x} \tag{1.22}$$

Where B^{-1} is the inverse matrix[18] of B.

1.8.3.2 Examples

Example 1.14:

Let us look at the standard basis first and try to calculate α from it.

$$B = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = I$$

Where I is the **identity matrix**. In order to solve for α let us find the inverse of B first (which is obviously very trivial in this case):

$$B^{-1} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

Therefore we get,

$$\alpha = B^{-1}\mathbf{x} = \mathbf{x}$$

Example 1.15:

Let us look at a ever-so-slightly more complicated basis of $\left\{ \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ -1 \end{pmatrix} \right\} = \{h_0, h_1\}$ Then our basis matrix and inverse basis matrix becomes:

$$B = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$$

$$B^{-1} = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{-1}{2} \end{pmatrix}$$

---

[18]http://cnx.org/content/m2113/latest/

and for this example it is given that

$$\mathbf{x} = \left(\begin{array}{c} 3 \\ 2 \end{array} \right)$$

Now we solve for α

$$\alpha = B^{-1}\mathbf{x} = \left(\begin{array}{cc} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{-1}{2} \end{array} \right) \left(\begin{array}{c} 3 \\ 2 \end{array} \right) = \left(\begin{array}{c} 2.5 \\ 0.5 \end{array} \right)$$

and we get

$$\mathbf{x} = 2.5h_0 + 0.5h_1$$

Exercise 1.1:

Now we are given the following basis matrix and $\mathbf{x}$:

$$\{b_0, b_1\} = \left\{ \left(\begin{array}{c} 1 \\ 2 \end{array} \right), \left(\begin{array}{c} 3 \\ 0 \end{array} \right) \right\}$$

$$\mathbf{x} = \left(\begin{array}{c} 3 \\ 2 \end{array} \right)$$

For this problem, make a sketch of the bases and then represent $\mathbf{x}$ in terms of b_0 and b_1.

NOTE: A change of basis simply looks at $\mathbf{x}$ from a "different perspective." B^{-1} **transforms x** from the standard basis to our new basis, $\{b_0, b_1\}$. Notice that this is a totally mechanical procedure.

1.8.4 Extending the Dimension and Space

We can also extend all these ideas past just $\mathbb{R}^2$ and look at them in $\mathbb{R}^n$ and C^n. This procedure extends naturally to higher (>2) dimensions. Given a basis $\{b_0, b_1, \ldots, b_{n-1}\}$ for $\mathbb{R}^n$, we want to find $\{\alpha_0, \alpha_1, \ldots, \alpha_{n-1}\}$ such that

$$\mathbf{x} = \alpha_0 b_0 + \alpha_1 b_1 + \cdots + \alpha_{n-1} b_{n-1} \tag{1.23}$$

Again, we will set up a basis matrix

$$B = \left(\begin{array}{ccccc} b_0 & b_1 & b_2 & \ldots & b_{n-1} \end{array} \right)$$

where the columns equal the basis vectors and it will always be an n×n matrix (although the above matrix does not appear to be square since we left terms in vector notation). We can then proceed to rewrite Equation 1.21

$$\mathbf{x} = \left(\begin{array}{cccc} b_0 & b_1 & \ldots & b_{n-1} \end{array} \right) \left(\begin{array}{c} \alpha_0 \\ \vdots \\ \alpha_{n-1} \end{array} \right) = B\alpha$$

and

$$\alpha = B^{-1}\mathbf{x}$$

1.9 Fourier Analysis

Fourier analysis is fundamental to understanding the behavior of signals and systems. This is a result of the fact that sinusoids are Eigenfunctions[19] of linear, time-invariant (LTI)[20]

---

[19]http://cnx.org/content/m10500/latest/
[20]http://cnx.org/content/m10084/latest/

systems. This is to say that if we pass any particular sinusoid through a LTI system, we get a scaled version of that same sinusoid on the output. Then, since Fourier analysis allows us to redefine the signals in terms of sinusoids, all we need to do is determine how any given system effects all possible sinusoids (its transfer function[21]) and we have a complete understanding of the system. Furthermore, since we are able to define the passage of sinusoids through a system as multiplication of that sinusoid by the transfer function at the same frequency, we can convert the passage of any signal through a system from convolution[22] (in time) to multiplication (in frequency). These ideas are what give Fourier analysis its power.

Now, after hopefully having sold you on the value of this method of analysis, we must examine exactly what we mean by Fourier analysis. The four Fourier transforms that comprise this analysis are the Fourier Series[23], Continuous-Time Fourier Transform (Section 2.10), Discrete-Time Fourier Transform (Section 2.11) and Discrete Fourier Transform[24]. For this document, we will view the Laplace Transform[25] and Z-Transform (Section 4.3) as simply extensions of the CTFT and DTFT respectively. All of these transforms act essentially the same way, by converting a signal in time to an equivalent signal in frequency (sinusoids). However, depending on the nature of a specific signal *i.e.* whether it is finite- or infinite-length and whether it is discrete- or continuous-time) there is an appropriate transform to convert the signal into the frequency domain. Below is a table of the four Fourier transforms and when each is appropriate. It also includes the relevant convolution for the specified space.

<div align="center">Table of Fourier Representations</div>

Transform	Time Domain	Frequency Domain	Convolution
Continuous-Time Fourier Series	$L^2\left([0,T)\right)$	$l^2\left(\mathbb{Z}\right)$	Continuous-Time Circular
Continuous-Time Fourier Transform	$L^2\left(\mathbb{R}\right)$	$L^2\left(\mathbb{R}\right)$	Continuous-Time Linear
Discrete-Time Fourier Transform	$l^2\left(\mathbb{Z}\right)$	$L^2\left([0,2\pi)\right)$	Discrete-Time Linear
Discrete Fourier Transform	$l^2\left([0,N-1]\right)$	$l^2\left([0,N-1]\right)$	Discrete-Time Circular

1.10 Continuous-Time Fourier Transform (CTFT)

1.10.1 Introduction

Due to the large number of continuous-time signals that are present, the Fourier series[26] provided us the first glimpse of how me we may represent some of these signals in a general manner: as a superposition of a number of sinusoids. Now, we can look at a way to represent continuous-time nonperiodic signals using the same idea of superposition. Below we will present the **Continuous-Time Fourier Transform** (CTFT), also referred to as just the Fourier Transform (FT). Because the CTFT now deals with nonperiodic signals, we must now find a way to include *all* frequencies in the general equations.

1.10.1.1 Equations

Continuous-Time Fourier Transform

---

[21]http://cnx.org/content/m0028/latest/
[22]http://cnx.org/content/m10088/latest/
[23]http://cnx.org/content/m10097/latest/
[24]http://cnx.org/content/m0502/latest/
[25]http://cnx.org/content/m10110/latest/
[26]http://cnx.org/content/m0039/latest/

Figure 1.23: Mapping $L^2(\mathbb{R})$ in the time domain to $L^2(\mathbb{R})$ in the frequency domain.

$$\mathcal{F}(\Omega) = \int_{-\infty}^{\infty} f(t) e^{-(j\Omega t)} dt \tag{1.24}$$

Inverse CTFT

$$f(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \mathcal{F}(\Omega) e^{j\Omega t} d\Omega \tag{1.25}$$

WARNING: Do not be confused by notation - it is not uncommon to see the above formula written slightly different. One of the most common differences among many professors is the way that the exponential is written. Above we used the radial frequency variable Ω in the exponential, where $\Omega = 2\pi f$, but one will often see professors include the more explicit expression, $j2\pi ft$, in the exponential. Click here[27] for an overview of the notation used in Connexion's DSP modules.

The above equations for the CTFT and its inverse come directly from the Fourier series and our understanding of its coefficients. For the CTFT we simply utilize integration rather than summation to be able to express the aperiodic signals. This should make sense since for the CTFT we are simply extending the ideas of the Fourier series to include nonperiodic signals, and thus the entire frequency spectrum. Look at the Derivation of the Fourier Transform[28] for a more in depth look at this.

1.10.2 Relevant Spaces

The Continuous-Time Fourier Transform maps infinite-length, continuous-time signals in L^2 to infinite-length, continuous-frequency signals in L^2. Review the Fourier Analysis (Section 2.9) for an overview of all the spaces used in Fourier analysis.

For more information on the characteristics of the CTFT, please look at the module on Properties of the Fourier Transform[29].

1.10.3 Example Problems

Exercise 1.2:

Find the Fourier Transform (CTFT) of the function

$$f(t) = \begin{cases} e^{-(\alpha t)} & \text{if } t \geq 0 \\ 0 & \text{otherwise} \end{cases} \tag{1.26}$$

---

[27]http://cnx.org/content/m10161/latest/
[28]http://cnx.org/content/m0046/latest/
[29]http://cnx.org/content/m10100/latest/

Figure 1.24: Mapping $l^2\left(\mathbb{Z}\right)$ in the time domain to $L^2\left([0, 2\pi)\right)$ in the frequency domain.

Exercise 1.3:

Find the inverse Fourier transform of the square wave defined as

$$X\left(\Omega\right) = \begin{cases} 1 \text{ if } |\Omega| \le M \\ 0 \text{ otherwise} \end{cases} \tag{1.27}$$

1.11 Discrete-Time Fourier Transform (DTFT)

Discrete-Time Fourier Transform

$$X\left(\omega\right) = \sum_{n=-\infty}^{\infty} \left(x\left(n\right) e^{-\left(j\omega n\right)} \right) \tag{1.28}$$

Inverse Discrete-Time Fourier Transform

$$x\left(n\right) = \frac{1}{2\pi} \int_0^{2\pi} X\left(\omega\right) e^{j\omega n} d\omega \tag{1.29}$$

1.11.1 Relevant Spaces

The Discrete-Time Fourier Transform [30] maps infinite-length, discrete-time signals in l^2 to finite-length (or 2π-periodic), continuous-frequency signals in L^2.

1.12 DFT as a Matrix Operation

1.12.1 Matrix Review

Recall:

- Vectors in $\mathbb{R}^N$:

$$\mathbf{x} = \begin{pmatrix} x_0 \\ x_1 \\ \cdots \\ x_{N-1} \end{pmatrix} \ , \ \ x_i \in \mathbb{R}$$

---

[30]http://cnx.org/content/m10247/latest/

- Vectors in C^N:

$$\mathbf{x} = \begin{pmatrix} x_0 \\ x_1 \\ \dots \\ x_{N-1} \end{pmatrix} \quad , \quad x_i \in C$$

- Transposition:

    1. transpose:

$$\mathbf{x}^T = \begin{pmatrix} x_0 & x_1 & \dots & x_{N-1} \end{pmatrix}$$

    2. conjugate:

$$\mathbf{x}^H = \begin{pmatrix} x_0{}^* & x_1{}^* & \dots & x_{N-1}{}^* \end{pmatrix}$$

- Inner product[31]:

    1. real:

$$\mathbf{x}^T \mathbf{y} = \sum_{i=0}^{N-1} (x_i y_i)$$

    2. complex:

$$\mathbf{x}^H \mathbf{y} = \sum_{i=0}^{N-1} (x_n{}^* y_n)$$

- Matrix Multiplication:

$$A\mathbf{x} = \begin{pmatrix} a_{00} & a_{01} & \dots & a_{0,N-1} \\ a_{10} & a_{11} & \dots & a_{1,N-1} \\ \vdots & \vdots & \dots & \vdots \\ a_{N-1,0} & a_{N-1,1} & \dots & a_{N-1,N-1} \end{pmatrix} \begin{pmatrix} x_0 \\ x_1 \\ \dots \\ x_{N-1} \end{pmatrix} = \begin{pmatrix} y_0 \\ y_1 \\ \dots \\ y_{N-1} \end{pmatrix}$$

$$y_k = \sum_{n=0}^{N-1} (a_{kn} x_n)$$

- Matrix Transposition:

$$A^T = \begin{pmatrix} a_{00} & a_{10} & \dots & a_{N-1,0} \\ a_{01} & a_{11} & \dots & a_{N-1,1} \\ \vdots & \vdots & \dots & \vdots \\ a_{0,N-1} & a_{1,N-1} & \dots & a_{N-1,N-1} \end{pmatrix}$$

Matrix transposition involved simply swapping the rows with columns.

$$A^H = A^{T^*}$$

The above equation is Hermitian transpose.

$$\left[\mathbf{A}^T \right]_{kn} = \mathbf{A}_{nk}$$

$$\left[A^H \right]_{kn} = [\mathbf{A}^*]_{nk}$$

---

[31]http://cnx.org/content/m10755/latest/

1.12.2 Representing DFT as Matrix Operation

Now let's represent the DFT[32] in vector-matrix notation.

$$\mathbf{x} = \begin{pmatrix} x[0] \\ x[1] \\ \dots \\ x[N-1] \end{pmatrix}$$

$$\mathbf{X} = \begin{pmatrix} X[0] \\ X[1] \\ \dots \\ X[N-1] \end{pmatrix} \in C^N$$

Here $\mathbf{x}$ is the vector of time samples and $\mathbf{X}$ is the vector of DFT coefficients. How are $\mathbf{x}$ and $\mathbf{X}$ related:

$$X[k] = \sum_{n=0}^{N-1} \left(x[n] e^{-\left(j\frac{2\pi}{N}kn\right)} \right)$$

where

$$a_{kn} = \left(e^{-\left(j\frac{2\pi}{N}\right)} \right)^{kn} = W_N{}^{kn}$$

so

$$\mathbf{X} = W\mathbf{x}$$

where $\mathbf{X}$ is the DFT vector, W is the matrix and $\mathbf{x}$ the time domain vector.

$$W_{kn} = \left(e^{-\left(j\frac{2\pi}{N}\right)} \right)^{kn}$$

$$\mathbf{X} = W \begin{pmatrix} x[0] \\ x[1] \\ \dots \\ x[N-1] \end{pmatrix}$$

IDFT:

$$x[n] = \frac{1}{N} \sum_{k=0}^{N-1} \left(X[k] \left(e^{j\frac{2\pi}{N}} \right)^{nk} \right)$$

where

$$\left(e^{j\frac{2\pi}{N}} \right)^{nk} = W_N{}^{nk*}$$

$W_N{}^{nk*}$ is the matrix Hermitian transpose. So,

$$\mathbf{x} = \frac{1}{N} W^H \mathbf{X}$$

where $\mathbf{x}$ is the time vector, $\frac{1}{N}W^H$ is the inverse DFT matrix, and $\mathbf{X}$ is the DFT vector.

1.13 The FFT Algorithm

Definition 1: FFT
(Fast Fourier Transform) An efficient **computational algorithm** for computing
the DFT[33].

---

[32]http://cnx.org/content/m10249/latest/
[33]http://cnx.org/content/m10249/latest/

Figure 1.25

Figure 1.26

1.13.1 The Fast Fourier Transform FFT

DFT can be *expensive* to compute *directly*

$$X[k] = \sum_{n=0}^{N-1} \left(x[n] e^{-\left(j2\pi \frac{k}{N} n\right)} \right) \quad , \quad 0 \le k \le N-1$$

For each k, we must execute:

- N complex multiplies
- $N-1$ complex adds

The total cost of direct computation of an N-point DFT is

- N^2 complex multiplies
- $N(N-1)$ complex adds

How many adds and mults of *real* numbers are required?

This " $O\left(N^2\right)$ " computation rapidly gets out of hand, as N gets large:

N	1	10	100	1000	10^6
N^2	1	100	10,000	10^6	10^{12}

The FFT provides us with a much more *efficient* way of computing the DFT. The FFT requires only " $O\left(N\log N\right)$ " computations to compute the N-point DFT.

N	10	100	1000	10^6
N^2	100	10,000	10^6	10^{12}
$N\log_{10}N$	10	200	3000	6×10^6

How long is $10^{12}\mu$sec? More than 10 days! How long is $6 \times 10^6\mu$sec?
The FFT and digital computers revolutionized DSP (1960 - 1980).

1.13.2 How does the FFT work?

- The FFT exploits the **symmetries** of the complex exponentials $W_N{}^{kn} = e^{-\left(j\frac{2\pi}{N}kn\right)}$

- $W_N{}^{kn}$ are called "**twiddle factors**"

  Symmetry 1.1:   Complex Conjugate Symmetry

  $$W_N{}^{k(N-n)} = W_N{}^{-(kn)} = W_N{}^{kn*}$$

  $$e^{-\left(j2\pi\frac{k}{N}(N-n)\right)} = e^{j2\pi\frac{k}{N}n} = e^{-\left(j2\pi\frac{k}{N}n\right)^*}$$

  Symmetry 1.2:   Periodicity in n and k

  $$W_N{}^{kn} = W_N{}^{k(N+n)} = W_N{}^{(k+N)n}$$

  $$e^{-\left(j\frac{2\pi}{N}kn\right)} = e^{-\left(j\frac{2\pi}{N}k(N+n)\right)} = e^{-\left(j\frac{2\pi}{N}(k+N)n\right)}$$

  $$W_N = e^{-\left(j\frac{2\pi}{N}\right)}$$

1.13.3 Decimation in Time FFT

- Just one of *many* different FFT algorithms

- The *idea* is to build a DFT out of smaller and smaller DFTs by decomposing $x[n]$ into smaller and smaller subsequences.

- Assume $N = 2^m$ (a power of 2)

1.13.3.1 Derivation

N is *even*, so we can complete $X[k]$ by separating $x[n]$ into *two* subsequences each of length $\frac{N}{2}$.

$$x[n] \rightarrow \begin{cases} \frac{N}{2} & \text{if } n = \text{even} \\ \frac{N}{2} & \text{if } n = \text{odd} \end{cases}$$

$$X[k] = \sum_{n=0}^{N-1}\left(x[n]W_N{}^{kn}\right) \quad,\quad 0 \le k \le N-1$$

$$X[k] = \sum\left(x[n]W_N{}^{kn}\right) + \sum\left(x[n]W_N{}^{kn}\right)$$

where $0 \le r \le \frac{N}{2} - 1$. So

$$\begin{aligned} X[k] &= \sum_{r=0}^{\frac{N}{2}-1}\left(x[2r]W_N{}^{2kr}\right) + \sum_{r=0}^{\frac{N}{2}-1}\left(x[2r+1]W_N{}^{(2r+1)k}\right) \\ &= \sum_{r=0}^{\frac{N}{2}-1}\left(x[2r]\left(W_N{}^2\right)^{kr}\right) + W_N{}^k\sum_{r=0}^{\frac{N}{2}-1}\left(x[2r+1]\left(W_N{}^2\right)^{kr}\right) \end{aligned} \tag{1.30}$$

where $W_N{}^2 = e^{-\left(j\frac{2\pi}{N}2\right)} = e^{-\left(j\frac{2\pi}{\frac{N}{2}}\right)} = W_{\frac{N}{2}}$. So

$$X[k] = \sum_{r=0}^{\frac{N}{2}-1}\left(x[2r]W_{\frac{N}{2}}{}^{kr}\right) + W_N{}^k\sum_{r=0}^{\frac{N}{2}-1}\left(x[2r+1]W_{\frac{N}{2}}{}^{kr}\right)$$

where $\sum_{r=0}^{\frac{N}{2}-1}\left(x\left[2r\right]W_{\frac{N}{2}}{}^{kr}\right)$ is $\frac{N}{2}$-point DFT of even samples ($G\left[k\right]$) and $\sum_{r=0}^{\frac{N}{2}-1}\left(x\left[2r+1\right]W_{\frac{N}{2}}{}^{kr}\right)$ is $\frac{N}{2}$-point DFT of odd samples ($H\left[k\right]$).

$$X\left[k\right]=G\left[k\right]+W_N{}^kH\left[k\right] \quad,\quad 0\le k\le N-1$$

Decomposition of an N-point DFT as a sum of 2 $\frac{N}{2}$-point DFTs.

Why would we want to do this? *Because it is more efficient!*

RECALL: Cost to compute an N-point DFT is approximately N^2 complex mults and adds.

But decomposition into 2 $\frac{N}{2}$-point DFTs + combination requires only

$$\left(\frac{N}{2}\right)^2+\left(\frac{N}{2}\right)^2+N=\frac{N^2}{2}+N$$

where the first part is the number of complex mults and adds for $\frac{N}{2}$-point DFT, $G\left[k\right]$. The second part is the number of complex mults and adds for $\frac{N}{2}$-point DFT, $H\left[k\right]$. The third part is the number of complex mults and adds for combination. And the total is $\frac{N^2}{2}+N$ complex mults and adds.

Example 1.16: Savings

For $N=1000$,

$$N^2=10^6$$
$$\frac{N^2}{2}+N=\frac{10^6}{2}+1000$$

Because 1000 is small compared to 500,000,

$$\frac{N^2}{2}+N\approx\frac{10^6}{2}$$

So why stop here?! Keep decomposing. Break each of the $\frac{N}{2}$-point DFTs into two $\frac{N}{4}$-point DFTs, *etc.*, ...

We can keep decomposing:

$$\frac{N}{2^1}=\left\{\frac{N}{2},\frac{N}{4},\frac{N}{8},\ldots,\frac{N}{2^{m-1}},\frac{N}{2^m}\right\}=1$$

where

$$m=\log_2N= \text{ times}$$

Computational cost: N-pt DFT [U+F577] two $\frac{N}{2}$-pt DFTs. The cost is $N^2\to2\left(\frac{N}{2}\right)^2+N$. So replacing each $\frac{N}{2}$-pt DFT with two $\frac{N}{4}$-pt DFTs will reduce cost to

$$2\left(2\left(\frac{N}{4}\right)^2+\frac{N}{2}\right)+N=4\left(\frac{N}{4}\right)^2+2N=\frac{N^2}{2^2}+2N=\frac{N^2}{2^p}+pN$$

As we keep going $p=\{3,4,\ldots,m\}$, where $m=\log_2N$. We get the cost

$$\frac{N^2}{2^{\log_2N}}+N\log_2N=\frac{N^2}{N}+N\log_2N=N+N\log_2N$$

$N+N\log_2N$ is the total number of complex adds and mults.

For large N, cost $\approx N\log_2N$ or " $O\left(N\log_2N\right)$", since $(N\log_2N\gg N)$ for large N.

NOTE: Weird order of time samples

Figure 1.27: $N = 8$ point FFT. Summing nodes $W_n{}^k$ twiddle multiplication factors.

Figure 1.28: This is called "butterflies."

Solutions to Exercises in Chapter 1

Solution to Exercise 1.1:

In order to represent $\mathbf{x}$ in terms of b_0 and b_1 we will follow the same steps we used in the above example.

$$B = \begin{pmatrix} 1 & 2 \\ 3 & 0 \end{pmatrix}$$

$$B^{-1} = \begin{pmatrix} 0 & \frac{1}{2} \\ \frac{1}{3} & \frac{-1}{6} \end{pmatrix}$$

$$\alpha = B^{-1}\mathbf{x} = \begin{pmatrix} 1 \\ \frac{2}{3} \end{pmatrix}$$

And now we can write $\mathbf{x}$ in terms of b_0 and b_1.

$$\mathbf{x} = b_0 + \frac{2}{3}b_1$$

And we can easily substitute in our known values of b_0 and b_1 to verify our results.

Solution to Exercise 1.2:

In order to calculate the Fourier transform, all we need to use is Equation 1.24, complex exponentials[34], and basic calculus.

$$
\begin{aligned}
\mathcal{F}(\Omega) &= \int_{-\infty}^{\infty} f(t) e^{-(j\Omega t)} dt \\
&= \int_{0}^{\infty} e^{-(\alpha t)} e^{-(j\Omega t)} dt \\
&= \int_{0}^{\infty} e^{(-t)(\alpha + j\Omega)} dt \\
&= 0 - \frac{-1}{\alpha + j\Omega}
\end{aligned}
\tag{1.31}
$$

$$\mathcal{F}(\Omega) = \frac{1}{\alpha + j\Omega} \tag{1.32}$$

---
[34]http://cnx.org/content/m10060/latest/

Solution to Exercise 1.3:

Here we will use Equation 1.25 to find the inverse FT given that $t \neq 0$.

$$
\begin{aligned}
x(t) &= \frac{1}{2\pi} \int_{-M}^{M} e^{j\Omega t} d\Omega \\
&= \frac{1}{2\pi} e^{j\Omega t}\big|_{\Omega, \Omega = e^{jw}} \\
&= \frac{1}{\pi t} \sin(Mt)
\end{aligned}
\tag{1.33}
$$

$$
x(t) = \frac{M}{\pi} \left(\text{sinc} \frac{Mt}{\pi} \right)
\tag{1.34}
$$

Chapter 2

Sampling and Frequency Analysis

2.1 Introduction

Contents of Sampling chapter

- Introduction(Current module)
- Proof (Section 3.2)
- Illustrations (Section 3.3)
- Matlab Example (Section 3.4)
- Hold operation[1]
- System view (Section 3.5)
- Aliasing applet[2]
- Exercises[3]
- Table of formulas[4]

2.1.1 Why sample?

This section introduces sampling. Sampling is the necessary fundament for all digital signal processing and communication. Sampling can be defined as the process of measuring an analog signal at distinct points.

Digital representation of analog signals offers advantages in terms of

- robustness towards noise, meaning we can send more bits/s
- use of flexible processing equipment, in particular the computer
- more reliable processing equipment
- easier to adapt complex algorithms

---

[1] http://cnx.org/content/m11458/latest/
[2] http://cnx.org/content/m11448/latest/
[3] http://cnx.org/content/m11442/latest/
[4] http://cnx.org/content/m11450/latest/

Figure 2.1: Claude Elwood Shannon (1916-2001)

2.1.2 Claude E. Shannon

Claude Shannon[5] has been called the father of information theory, mainly due to his landmark papers on the "Mathematical theory of communication"[6]. Harry Nyquist[7] was the first to state the sampling theorem in 1928, but it was not proven until Shannon proved it 21 years later in the paper "Communications in the presence of noise"[8].

2.1.3 Notation

In this chapter we will be using the following notation

- Original analog signal $x(t)$

- Sampling frequency F_s

- Sampling interval T_s (Note that: $F_s = \frac{1}{T_s}$)

- Sampled signal $x_s(n)$. (Note that $x_s(n) = x(nT_s)$)

- Real angular frequency Ω

- Digital angular frequency ω. (Note that: $\omega = \Omega T_s$)

2.1.4 The Sampling Theorem

THE SAMPLING THEOREM: When sampling an analog signal the sampling frequency must be greater than twice the highest frequency component of the analog signal to be able to reconstruct the original signal from the sampled version.

Finished? Have at look at: Proof (Section 3.2); Illustrations (Section 3.3); Matlab Example (Section 3.4); Aliasing applet[9]; Hold operation[10]; System view (Section 3.5); Exercises[11]

2.2 Proof

SAMPLING THEOREM: In order to recover the signal $x(t)$ from it's samples exactly, it is necessary to sample $x(t)$ at a rate greater than twice it's highest frequency component.

---

[5]http://www.research.att.com/~njas/doc/ces5.html
[6]http://cm.bell-labs.com/cm/ms/what/shannonday/shannon1948.pdf
[7]http://www.wikipedia.org/wiki/Harry_Nyquist
[8]http://www.stanford.edu/class/ee104/shannonpaper.pdf
[9]http://cnx.org/content/m11448/latest/
[10]http://cnx.org/content/m11458/latest/
[11]http://cnx.org/content/m11442/latest/

2.2.1 Introduction

As mentioned earlier (pg 45), sampling is the necessary fundament when we want to apply digital signal processing on analog signals.

Here we present the proof of the sampling theorem. The proof is divided in two. First we find an expression for the spectrum of the signal resulting from sampling the original signal $x(t)$. Next we show that the signal $x(t)$ can be recovered from the samples. Often it is easier using the frequency domain when carrying out a proof, and this is also the case here.

Key points in the proof

- We find an equation (Equation 2.8) for the spectrum of the sampled signal
- We find a simple method to reconstruct (Equation 2.14) the original signal
- The sampled signal has a periodic spectrum...
- ...and the period is $2\pi F_s$

2.2.2 Proof part 1 - Spectral considerations

By sampling $x(t)$ every T_s second we obtain $x_s(n)$. The inverse fourier transform of this time discrete signal[12] is

$$x_s(n) = \frac{1}{2\pi} \int_{-\pi}^{\pi} X_s\left(e^{j\omega}\right) e^{j\omega n} d\omega \tag{2.1}$$

For convenience we express the equation in terms of the real angular frequency Ω using $\omega = \Omega T_s$. We then obtain

$$x_s(n) = \frac{T_s}{2\pi} \int_{\frac{-\pi}{T_s}}^{\frac{\pi}{T_s}} X_s\left(e^{j\Omega T_s}\right) e^{j\Omega T_s n} d\Omega \tag{2.2}$$

The inverse fourier transform of a continuous signal is

$$x(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} X(j\Omega) e^{j\Omega t} d\Omega \tag{2.3}$$

From this equation we find an expression for $x(nT_s)$

$$x(nT_s) = \frac{1}{2\pi} \int_{-\infty}^{\infty} X(j\Omega) e^{j\Omega nT_s} d\Omega \tag{2.4}$$

To account for the difference in region of integration we split the integration in Equation 2.4 into subintervals of length $\frac{2\pi}{T_s}$ and then take the sum over the resulting integrals to obtain the complete area.

$$x(nT_s) = \frac{1}{2\pi} \sum_{k=-\infty}^{\infty} \left(\int_{\frac{(2k-1)\pi}{T_s}}^{\frac{(2k+1)\pi}{T_s}} X(j\Omega) e^{j\Omega nT_s} d\Omega \right) \tag{2.5}$$

Then we change the integration variable, setting $\Omega = \eta + \frac{2\pi k}{T_s}$

$$x(nT_s) = \frac{1}{2\pi} \sum_{k=-\infty}^{\infty} \left(\int_{\frac{-\pi}{T_s}}^{\frac{\pi}{T_s}} X\left(j\left(\eta + \frac{2\pi k}{T_s}\right)\right) e^{j\left(\eta + \frac{2\pi k}{T_s}\right) nT_s} d\eta \right) \tag{2.6}$$

---
[12]http://cnx.org/content/m11476/latest/

We obtain the final form by observing that $e^{j2\pi kn} = 1$, reinserting $\eta = \Omega$ and multiplying by $\frac{T_s}{T_s}$

$$x\left(nT_s\right) = \frac{T_s}{2\pi} \int_{\frac{-\pi}{T_s}}^{\frac{\pi}{T_s}} \sum_{k=-\infty}^{\infty} \left(\frac{1}{T_s} X\left(j\left(\Omega + \frac{2\pi k}{T_s}\right)\right)\right) e^{j\Omega nT_s}\, d\Omega \qquad (2.7)$$

To make $x_s\left(n\right) = x\left(nT_s\right)$ for all values of n, the integrands in Equation 2.2 and Equation 2.7 have to agree, that is

$$X_s\left(e^{j\Omega T_s}\right) = \frac{1}{T_s} \sum_{k=-\infty}^{\infty} \left(X\left(j\left(\Omega + \frac{2\pi k}{T_s}\right)\right)\right) \qquad (2.8)$$

This is a central result. We see that the digital spectrum consists of a sum of shifted versions of the original, analog spectrum. Observe the periodicity!

We can also express this relation in terms of the digital angular frequency $\omega = \Omega T_s$

$$X_s\left(e^{j\omega}\right) = \frac{1}{T_s} \sum_{k=-\infty}^{\infty} \left(X\left(j\frac{\omega + 2\pi k}{T_s}\right)\right) \qquad (2.9)$$

This concludes the first part of the proof. Now we want to find a reconstruction formula, so that we can recover $x\left(t\right)$ from $x_s\left(n\right)$.

2.2.3 Proof part II - Signal reconstruction

For a bandlimited (Figure 2.3) signal the inverse fourier transform is

$$x\left(t\right) = \frac{1}{2\pi} \int_{\frac{-\pi}{T_s}}^{\frac{\pi}{T_s}} X\left(j\Omega\right) e^{j\Omega t}\, d\Omega \qquad (2.10)$$

In the interval we are integrating we have: $X_s\left(e^{j\Omega T_s}\right) = \frac{X(j\Omega)}{T_s}$. Substituting this relation into Equation 2.10 we get

$$x\left(t\right) = \frac{T_s}{2\pi} \int_{\frac{-\pi}{T_s}}^{\frac{\pi}{T_s}} X_s\left(e^{j\Omega T_s}\right) e^{j\Omega t}\, d\Omega \qquad (2.11)$$

Using the DTFT[13] relation for $X_s\left(e^{j\Omega T_s}\right)$ we have

$$x\left(t\right) = \frac{T_s}{2\pi} \int_{\frac{-\pi}{T_s}}^{\frac{\pi}{T_s}} \sum_{n=-\infty}^{\infty} \left(x_s\left(n\right) e^{-(j\Omega nT_s)} e^{j\Omega t}\right) d\Omega \qquad (2.12)$$

Interchanging integration and summation (under the assumption of convergence) leads to

$$x\left(t\right) = \frac{T_s}{2\pi} \sum_{n=-\infty}^{\infty} \left(x_s\left(n\right) \int_{\frac{-\pi}{T_s}}^{\frac{\pi}{T_s}} e^{j\Omega(t-nT_s)}\, d\Omega\right) \qquad (2.13)$$

Finally we perform the integration and arrive at the important reconstruction formula

$$x\left(t\right) = \sum_{n=-\infty}^{\infty} \left(x_s\left(n\right) \frac{\sin\left(\frac{\pi}{T_s}\left(t - nT_s\right)\right)}{\frac{\pi}{T_s}\left(t - nT_s\right)}\right) \qquad (2.14)$$

(Thanks to R.Loos for pointing out an error in the proof.)

---

[13]http://cnx.org/content/m11450/latest/

2.2.4 Summary

SPECTRUM SAMPLED SIGNAL: $X_s\left(e^{j\Omega T_s}\right) = \frac{1}{T_s}\sum_{k=-\infty}^{\infty}\left(X\left(j\left(\Omega + \frac{2\pi k}{T_s}\right)\right)\right)$

RECONSTRUCTION FORMULA: $x\left(t\right) = \sum_{n=-\infty}^{\infty}\left(x_s\left(n\right)\frac{\sin\left(\frac{\pi}{T_s}\left(t-nT_s\right)\right)}{\frac{\pi}{T_s}\left(t-nT_s\right)}\right)$

Go to Introduction (Section 3.1); Illustrations (Section 3.3); Matlab Example (Section 3.4); Hold operation[14]; Aliasing applet[15]; System view (Section 3.5); Exercises[16] ?

2.3 Illustrations

In this module we illustrate the processes involved in sampling and reconstruction. To see how all these processes work together as a whole, take a look at the system view (Section 3.5). In Sampling and reconstruction with Matlab (Section 3.4) we provide a Matlab script for download. The matlab script shows the process of sampling and reconstruction **live**.

2.3.1 Basic examples

Example 2.1:

To sample an analog signal with 3000 Hz as the highest frequency component requires sampling at 6000 Hz or above.

Example 2.2:

The sampling theorem can also be applied in two dimensions, i.e. for image analysis. A 2D sampling theorem has a simple physical interpretation in image analysis: Choose the sampling interval such that it is less than or equal to half of the smallest interesting detail in the image.

2.3.2 The process of sampling

We start off with an analog signal. This can for example be the sound coming from your stereo at home or your friend talking.

The signal is then sampled uniformly. Uniform sampling implies that we sample every T_s seconds. In Figure 2.2 we see an analog signal. The analog signal has been sampled at times $t = nT_s$.   In signal processing it is often more convenient and easier to work in the frequency domain. So let's look at at the signal in frequency domain, Figure 2.3. For illustration purposes we take the frequency content of the signal as a triangle. (If you Fourier transform the signal in Figure 2.2 you will not get such a nice triangle.)   Notice that the signal in Figure 2.3 is bandlimited. We can see that the signal is bandlimited because $X\left(j\Omega\right)$ is zero outside the interval $[-\Omega_g, \Omega_g]$. Equivalently we can state that the signal has no angular frequencies above Ω_g, corresponding to no frequencies above $F_g = \frac{\Omega_g}{2\pi}$.

Now let's take a look at the sampled signal in the frequency domain. While proving (Section 3.2) the sampling theorem we found the the spectrum of the sampled signal consists

---

[14]http://cnx.org/content/m11458/latest/
[15]http://cnx.org/content/m11448/latest/
[16]http://cnx.org/content/m11442/latest/

Figure 2.2: Analog signal, samples are marked with dots.

Figure 2.3:   The spectrum $X\left(j\Omega\right)$.

Figure 2.4: The spectrum X_s. Sampling frequency is OK.

of a sum of shifted versions of the analog spectrum. Mathematically this is described by the following equation:

$$X_s \left(e^{j\Omega T_s} \right) = \frac{1}{T_s} \sum_{k=-\infty}^{\infty} \left(X \left(j \left(\Omega + \frac{2\pi k}{T_s} \right) \right) \right) \tag{2.15}$$

2.3.2.1 Sampling fast enough

In Figure 2.4 we show the result of sampling $x(t)$ according to the sampling theorem (Section 2.1.4). This means that when sampling the signal in Figure 2.2/Figure 2.3 we use $F_s \geq 2F_g$. Observe in Figure 2.4 that we have the same spectrum as in Figure 2.3 for $\Omega \in [-\Omega_g, \Omega_g]$, except for the scaling factor $\frac{1}{T_s}$. This is a consequence of the sampling frequency. As mentioned in the proof (Key points in the proof) the spectrum of the sampled signal is periodic with period $2\pi F_s = \frac{2\pi}{T_s}$.

So now we are, according to the sample theorem (Section 2.1.4), able to reconstruct the original signal *exactly*. How we can do this will be explored further down under reconstruction (Section 2.3.3). But first we will take a look at what happens when we sample too slowly.

2.3.2.2 Sampling too slowly

If we sample $x(t)$ too slowly, that is $F_s < 2F_g$, we will get overlap between the repeated spectra, see Figure 2.5. According to Equation 2.15 the resulting spectra is the sum of these. This overlap gives rise to the concept of aliasing.

> ALIASING: If the sampling frequency is less than twice the highest frequency component, then frequencies in the original signal that are above half the sampling rate will be "aliased" and will appear in the resulting signal as lower frequencies.

The consequence of aliasing is that we cannot recover the original signal, so aliasing has to be avoided. Sampling too slowly will produce a sequence $x_s(n)$ that could have orginated from a number of signals. So there is *no* chance of recovering the original signal. To learn more about aliasing, take a look at this module[17]. (Includes an applet for demonstration!)

To avoid aliasing we have to sample fast enough. But if we can't sample fast enough (possibly due to costs) we can include an Anti-Aliasing filter. This will not able us to get an exact reconstruction but can still be a good solution.

---

[17]http://cnx.org/content/m11448/latest/

Figure 2.5: The spectrum X_s. Sampling frequency is too low.

ANTI-ALIASING FILTER: Typically a low-pass filter that is applied before sampling to ensure that no components with frequencies greater than half the sample frequency remain.

Example 2.3:

The stagecoach effect

In older western movies you can observe aliasing on a stagecoach when it starts to roll. At first the spokes appear to turn forward, but as the stagecoach increase its speed the spokes appear to turn backward. This comes from the fact that the sampling rate, here the number of frames per second, is too low. We can view each frame as a sample of an image that is changing continuously in time. (Applet illustrating the stagecoach effect[18])

2.3.3 Reconstruction

Given the signal in Figure 2.4 we want to recover the original signal, but the question is how?

When there is no overlapping in the spectrum, the spectral component given by $k = 0$ (see Equation 2.15),is equal to the spectrum of the analog signal. This offers an oppurtunity to use a simple reconstruction process. Remember what you have learned about filtering. What we want is to change signal in Figure 2.4 into that of Figure 2.3. To achieve this we have to remove all the extra components generated in the sampling process. To remove the extra components we apply an ideal analog low-pass filter as shown in Figure 2.6 As we see the ideal filter is rectangular in the frequency domain. A rectangle in the frequency domain corresponds to a sinc[19] function in time domain (and vice versa).

Then we have reconstructed the original spectrum, and as we know *if two signals are identical in the frequency domain, they are also identical in the time domain.* End of reconstruction.

2.3.4 Conclusions

The Shannon sampling theorem requires that the input signal prior to sampling is band-limited to at most half the sampling frequency. Under this condition the samples give an

---

[18]http://flowers.ofthenight.org/wagonWheel/wagonWheel.html
[19]http://ccrma-www.stanford.edu/~jos/Interpolation/sinc_function.html

Figure 2.6: $H(j\Omega)$ The ideal reconstruction filter.

exact signal representation. It is truly remarkable that such a broad and useful class signals can be represented that easily!

We also looked into the problem of reconstructing the signals form its samples. Again the simplicity of the *principle* is striking: linear filtering by an ideal low-pass filter will do the job. However, the ideal filter is impossible to create, but that is another story...

Go to? Introduction (Section 3.1); Proof (Section 3.2); Illustrations (Section 3.3); Matlab Example (Section 3.4); Aliasing applet[20]; Hold operation[21]; System view (Section 3.5); Exercises[22]

2.4 Sampling and Reconstruction with Matlab

2.4.1 Matlab files

Samprecon.m[23]

Introduction (Section 3.1); Proof (Section 3.2); Illustrations (Section 3.3); Aliasing applet[24]; Hold operation (Section 3.4); System view (Section 3.5); Exercises[25]

2.5 Systems View of Sampling and Reconstruction

2.5.1 Ideal reconstruction system

Figure 2.7 shows the ideal reconstruction system based on the results of the Sampling theorem proof (Section 3.2).

Figure 2.7 consists of a sampling device which produces a time-discrete sequence $x_s(n)$. The reconstruction filter, $h(t)$, is an ideal analog sinc[26] filter, with $h(t) = \text{sinc}\left(\frac{t}{T_s}\right)$. We can't apply the time-discrete sequence $x_s(n)$ directly to the analog filter $h(t)$. To solve this problem we turn the sequence into an analog signal using delta functions[27]. Thus we write $x_s(t) = \sum_{n=-\infty}^{\infty}(x_s(n)\delta(t-nT))$. But when will the system produce an output $\hat{x}(t) = x(t)$? According to the sampling theorem (Section 2.1.4) we have $\hat{x}(t) = x(t)$ when the sampling frequency, F_s, is at least twice the highest frequency component of $x(t)$.

[20]http://cnx.org/content/m11448/latest/
[21]http://cnx.org/content/m11458/latest/
[22]http://cnx.org/content/m11442/latest/
[23]http://cnx.rice.edu/content/m11549/latest/Samprecon.m
[24]http://cnx.org/content/m11448/latest/
[25]http://cnx.org/content/m11442/latest/
[26]http://ccrma-www.stanford.edu/~jos/Interpolation/sinc_function.html
[27]http://cnx.org/content/m11450/latest/

Figure 2.7: Ideal reconstruction system

Figure 2.8:   Ideal reconstruction system with anti-aliasing filter (pg 51)

2.5.2 Ideal system including anti-aliasing

To be sure that the reconstructed signal is free of aliasing it is customary to apply a lowpass filter, an anti-aliasing filter (pg 51), before sampling as shown in Figure 2.8.   Again we ask the question of when the system will produce an output $\hat{x}(t) = s(t)$? If the signal is entirely confined within the passband of the lowpass filter we will get perfect reconstruction if F_s is high enough.

But if the anti-aliasing filter removes the "higher" frequencies, (which in fact is the job of the anti-aliasing filter), we will *never* be able to *exactly* reconstruct the original signal, $s(t)$. If we sample fast enough we can reconstruct $x(t)$, which in most cases is satisfying.

The reconstructed signal, $\hat{x}(t)$, will not have aliased frequencies. This is essential for further use of the signal.

2.5.3 Reconstruction with hold operation

To make our reconstruction system realizable there are many things to look into. Among them are the fact that any practical reconstruction system must input finite length pulses into the reconstruction filter. This can be accomplished by the hold operation[28]. To alleviate the distortion caused by the hold opeator we apply the output from the hold device to a compensator. The compensation can be as accurate as we wish, this is cost and application consideration.   By the use of the hold component the reconstruction will not be exact, but as mentioned above we can get as close as we want. Introduction (Section 3.1); Proof (Section 3.2); Illustrations (Section 3.3); Matlab example (Section 3.4); Hold operation[30];

---

[28]http://cnx.org/content/m11458/latest/

[30]http://cnx.org/content/m11458/latest/

Figure 2.9:   More practical reconstruction system with a hold component[29]

Aliasing applet[31]; Exercises[32]

2.6 Sampling CT Signals: A Frequency Domain Perspective

2.6.1 Understanding Sampling in the Frequency Domain

We want to relate $x_c(t)$ directly to $x[n]$. Compute the CTFT of

$$x_s(t) = \sum_{n=-\infty}^{\infty} (x_c(nT)\delta(t - nT))$$

$$
\begin{aligned}
X_s(\Omega) &= \int_{-\infty}^{\infty} \left(\sum_{n=-\infty}^{\infty} (x_c(nT)\delta(t - nT))\right) e^{(-j)\Omega t} dt \\
&= \sum_{n=-\infty}^{\infty} \left(x_c(nT) \int_{-\infty}^{\infty} \delta(t - nT) e^{(-j)\Omega t} dt\right) \\
&= \sum_{n=-\infty}^{\infty} \left(x[n] e^{(-j)\Omega nT}\right) \\
&= \sum_{n=-\infty}^{\infty} \left(x[n] e^{(-j)\omega n}\right) \\
&= X(\omega)
\end{aligned}
\tag{2.16}
$$

where $\omega \equiv \Omega T$ and $X(\omega)$ is the DTFT of $x[n]$.

RECALL:

$$X_s(\Omega) = \frac{1}{T} \sum_{k=-\infty}^{\infty} (X_c(\Omega - k\Omega_s))$$

$$
\begin{aligned}
X(\omega) &= \frac{1}{T} \sum_{k=-\infty}^{\infty} (X_c(\Omega - k\Omega_s)) \\
&= \frac{1}{T} \sum_{k=-\infty}^{\infty} \left(X_c\left(\frac{\omega - 2\pi k}{T}\right)\right)
\end{aligned}
\tag{2.17}
$$

where this last part is 2π-periodic.

2.6.1.1 Sampling

Example 2.4:   Speech

Speech is intelligible if bandlimited by a CT lowpass filter to the band ± 4 kHz. We can sample speech as slowly as _ _ _ _ _?

2.6.2 Relating x[n] to sampled x(t)

Recall the following equality:

$$x_s(t) = \sum_n (x(nT)\delta(t - nT))$$

Recall the CTFT relation:

$$\left(x(\alpha t) \leftrightarrow \frac{1}{\alpha} X\left(\frac{\Omega}{\alpha}\right)\right) \tag{2.18}$$

where α is a scaling of time and $\frac{1}{\alpha}$ is a scaling in frequency.

$$X_s(\Omega) \equiv X(\Omega T) \tag{2.19}$$

---

[31]http://cnx.org/content/m11448/latest/
[32]http://cnx.org/content/m11442/latest/

Figure 2.10

Figure 2.11

Figure 2.12: Note that there is no mention of T or Ω_s!

2.7 The DFT: Frequency Domain with a Computer Analysis

2.7.1 Introduction

We just covered ideal (and non-ideal) (time) sampling of CT signals (Section 3.6). This enabled DT signal processing solutions for CT applications (Figure 2.14):

Much of the theoretical analysis of such systems relied on frequency domain representations. How do we carry out these frequency domain analysis on the computer? Recall the following relationships:

$$x[n] \overset{\text{DTFT}}{\leftrightarrow} X(\omega)$$

$$x(t) \overset{\text{CTFT}}{\leftrightarrow} X(\Omega)$$

where ω and Ω are continuous frequency variables.

2.7.1.1 Sampling DTFT

Consider the DTFT of a discrete-time (DT) signal $x[n]$. Assume $x[n]$ is of finite duration N (*i.e.*, an N-point signal).

$$X(\omega) = \sum_{n=0}^{N-1} \left(x[n] e^{(-j)\omega n} \right) \tag{2.20}$$

where $X(\omega)$ is the continuous function that is indexed by the real-valued parameter $-\pi \leq \omega \leq \pi$. The other function, $x[n]$, is a discrete function that is indexed by integers.

Figure 2.13

Figure 2.14

Finite Duration DT Signal

Figure 2.15

Sample X(ω)

Figure 2.16

We want to work with $X(\omega)$ on a computer. Why not just *sample* $X(\omega)$?

$$
\begin{aligned}
X[k] &= X\left(\tfrac{2\pi}{N}k\right) \\
&= \sum_{n=0}^{N-1}\left(x[n]\,e^{(-j)2\pi\frac{k}{N}n}\right)
\end{aligned}
\tag{2.21}
$$

In Equation 2.21 we sampled at $\omega = \frac{2\pi}{N}k$ where $k = \{0,1,\ldots,N-1\}$ and $X[k]$ for $k = \{0,\ldots,N-1\}$ is called the **Discrete Fourier Transform (DFT)** of $x[n]$.

Example 2.5:

The DTFT of the image in Figure 2.15 is written as follows:

$$
X(\omega) = \sum_{n=0}^{N-1}\left(x[n]\,e^{(-j)\omega n}\right)
\tag{2.22}
$$

where ω is any 2π-interval, for example $-\pi \leq \omega \leq \pi$.

where again we sampled at $\omega = \frac{2\pi}{N}k$ where $k = \{0,1,\ldots,M-1\}$. For example, we take

$$M = 10$$

. In the following section (Section 2.7.1.1.1) we will discuss in more detail how we should choose M, the number of samples in the 2π interval.

(This is precisely how we would plot $X(\omega)$ in Matlab.)

2.7.1.1.1 Choosing M

2.7.1.1.1.1 Case 1

Given N (length of $x[n]$), choose $(M \gg N)$ to obtain a dense sampling of the DTFT (Figure 2.17):

2.7.1.1.1.2 Case 2

Choose M as small as possible (to minimize the amount of computation).

In general, we require $M \geq N$ in order to represent all information in

$$x[n] \quad,\quad n = \{0,\ldots,N-1\}$$

Figure 2.17

Let's concentrate on $M = N$:

$$x\,[n] \overset{\text{DFT}}{\leftrightarrow} X\,[k]$$

for $n = \{0, \ldots, N-1\}$ and $k = \{0, \ldots, N-1\}$

$$\text{numbers} \leftrightarrow \text{N numbers}$$

2.7.2 Discrete Fourier Transform (DFT)

Define

$$X\,[k] \equiv X\left(\frac{2\pi k}{N}\right) \tag{2.23}$$

where $N = \text{length}\,(x\,[n])$ and $k = \{0, \ldots, N-1\}$. In this case, $M = N$.

DFT

$$X\,[k] = \sum_{n=0}^{N-1} \left(x\,[n]\, e^{(-j)2\pi \frac{k}{N} n}\right) \tag{2.24}$$

Inverse DFT (IDFT)

$$x\,[n] = \frac{1}{N} \sum_{k=0}^{N-1} \left(X\,[k]\, e^{j2\pi \frac{k}{N} n}\right) \tag{2.25}$$

2.7.2.1 Interpretation

Represent $x\,[n]$ in terms of a sum of N complex sinusoids[33] of amplitudes $X\,[k]$ and frequencies

$$\omega_k = \frac{2\pi k}{N} \quad , \quad k \in \{0, \ldots, N-1\}$$

Think: Fourier Series with fundamental frequency $\frac{2\pi}{N}$

---

[33]http://cnx.org/content/m10060/latest/

Figure 2.18

2.7.2.1.1 Remark 1

IDFT treats $x[n]$ as though it were N-periodic.

$$x[n] = \frac{1}{N}\sum_{k=0}^{N-1}\left(X[k]e^{j2\pi\frac{k}{N}n}\right) \tag{2.26}$$

where $n \in \{0,\ldots,N-1\}$

Exercise 2.1:

What about other values of n?

2.7.2.1.2 Remark 2

Proof that the IDFT inverts the DFT for $n \in \{0,\ldots,N-1\}$

$$\begin{array}{rcl}\frac{1}{N}\sum_{k=0}^{N-1}\left(X[k]e^{j2\pi\frac{k}{N}n}\right) & = & \frac{1}{N}\sum_{k=0}^{N-1}\left(\sum_{m=0}^{N-1}\left(x[m]e^{(-j)2\pi\frac{k}{N}m}e^{j2\pi\frac{k}{N}n}\right)\right) \\ & = & ??? \end{array} \tag{2.27}$$

Example 2.6: Computing DFT

Given the following discrete-time signal (Figure 2.18) with $N = 4$, we will compute the DFT using two different methods (the DFT Formula and Sample DTFT):

1. DFT Formula

$$\begin{array}{rcl}X[k] & = & \sum_{n=0}^{N-1}\left(x[n]e^{(-j)2\pi\frac{k}{N}n}\right) \\ & = & 1 + e^{(-j)2\pi\frac{k}{4}} + e^{(-j)2\pi\frac{k}{4}2} + e^{(-j)2\pi\frac{k}{4}3} \\ & = & 1 + e^{(-j)\frac{\pi}{2}k} + e^{(-j)\pi k} + e^{(-j)\frac{3}{2}\pi k}\end{array} \tag{2.28}$$

Using the above equation, we can solve and get the following results:

$$x[0] = 4$$
$$x[1] = 0$$
$$x[2] = 0$$
$$x[3] = 0$$

Figure 2.19

2. Sample DTFT. Using the same figure, Figure 2.18, we will take the DTFT of
   the signal and get the following equations:

$$
\begin{aligned}
X(\omega) &= \sum_{n=0}^{3} \left(e^{(-j)\omega n} \right) \\
&= \frac{1 - e^{(-j)4\omega}}{1 - e^{(-j)\omega}} \\
&= \text{???}
\end{aligned}
\tag{2.29}
$$

Our sample points will be:

$$
\omega_k = \frac{2\pi k}{4} = \frac{\pi}{2} k
$$

where $k = \{0, 1, 2, 3\}$ (Figure 2.19).

2.7.3 Periodicity of the DFT

DFT $X[k]$ consists of *samples* of DTFT, so $X(\omega)$, a 2π-periodic DTFT signal, can be
converted to $X[k]$, an N-periodic DFT.

$$
X[k] = \sum_{n=0}^{N-1} \left(x[n] e^{(-j)2\pi \frac{k}{N} n} \right)
\tag{2.30}
$$

where $e^{(-j)2\pi \frac{k}{N} n}$ is an N-periodic basis function (See Figure 2.20).

Also, recall,

$$
\begin{aligned}
x[n] &= \frac{1}{N} \sum_{n=0}^{N-1} \left(X[k] e^{j2\pi \frac{k}{N} n} \right) \\
&= \frac{1}{N} \sum_{n=0}^{N-1} \left(X[k] e^{j2\pi \frac{k}{N} (n+mN)} \right) \\
&= \text{???}
\end{aligned}
\tag{2.31}
$$

Example 2.7:   Illustration

NOTE: When we deal with the DFT, we need to remember that, in effect,
this treats the signal as an N-periodic sequence.

Figure 2.20

Figure 2.21

Figure 2.22

2.7.4 A Sampling Perspective

Think of sampling the continuous function $X(\omega)$, as depicted in Figure 2.22. $S(\omega)$ will represent the sampling function applied to $X(\omega)$ and is illustrated in Figure 2.22 as well. This will result in our discrete-time sequence, $X[k]$.

RECALL: Remember the multiplication in the frequency domain is equal to convolution in the time domain!

2.7.4.1 Inverse DTFT of S(ω)

$$\sum_{k=-\infty}^{\infty}\left(\delta\left(\omega-\frac{2\pi k}{N}\right)\right) \tag{2.32}$$

Figure 2.23

Given the above equation, we can take the DTFT and get the following equation:

$$N \sum_{m=-\infty}^{\infty} \left(\delta\left[n - mN\right] \right) \equiv S\left[n\right] \qquad (2.33)$$

Exercise 2.2:

Why does Equation 2.33 equal $S\left[n\right]$?

So, in the time-domain we have (Figure 2.23):

2.7.5 Connections

Combine signals in Figure 2.24 to get signals in Figure 2.25.

2.8 Discrete-Time Processing of CT Signals

2.8.1 DT Processing of CT Signals

2.8.1.1 Analysis

$$Y_c\left(\Omega\right) = H_{\mathrm{LP}}\left(\Omega\right) Y\left(\Omega T\right) \qquad (2.34)$$

where we know that $Y\left(\omega\right) = X\left(\omega\right) G\left(\omega\right)$ and $G\left(\omega\right)$ is the frequency response of the DT LTI system. Also, remember that

$$\omega \equiv \Omega T$$

So,

$$Y_c\left(\Omega\right) = H_{\mathrm{LP}}\left(\Omega\right) G\left(\Omega T\right) X\left(\Omega T\right) \qquad (2.35)$$

where $Y_c\left(\Omega\right)$ and $H_{\mathrm{LP}}\left(\Omega\right)$ are CTFTs and $G\left(\Omega T\right)$ and $X\left(\Omega T\right)$ are DTFTs.

Figure 2.24

Figure 2.25

DSP System

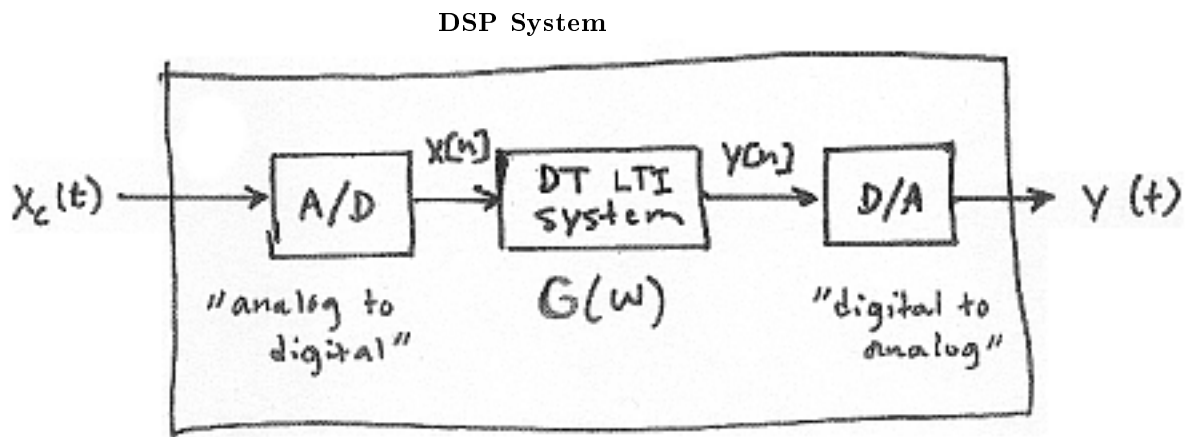**Figure 2.26**

Figure 2.27

RECALL:

$$X(\omega) = \frac{2\pi}{T} \sum_{k=-\infty}^{\infty} \left(X_c \left(\frac{\omega - 2\pi k}{T} \right) \right)$$

OR

$$X(\Omega T) = \frac{2\pi}{T} \sum_{k=-\infty}^{\infty} (X_c(\Omega - k\Omega_s))$$

Therefore our final output signal, $Y_c(\Omega)$, will be:

$$Y_c(\Omega) = H_{\mathrm{LP}}(\Omega) \, G(\Omega T) \left(\frac{2\pi}{T} \sum_{k=-\infty}^{\infty} (X_c(\Omega - k\Omega_s)) \right) \tag{2.36}$$

Now, if $X_c(\Omega)$ is bandlimited to $\left[-\left(\frac{\Omega_s}{2} \right), \frac{\Omega_s}{2} \right]$ and we use the usual lowpass reconstruction filter in the D/A, Figure 2.27:
Then,

$$Y_c(\Omega) = \begin{cases} G(\Omega T) \, X_c(\Omega) \text{ if } |\Omega| < \frac{\Omega_s}{2} \\ 0 \text{ otherwise} \end{cases} \tag{2.37}$$

Figure 2.28

2.8.1.2 Summary

For bandlimited signals sampled at or above the Nyquist rate, we can relate the input and output of the DSP system by:

$$Y_c(\Omega) = G_{\text{eff}}(\Omega) X_c(\Omega) \qquad (2.38)$$

where

$$G_{\text{eff}}(\Omega) = \left\{ \begin{array}{c} G(\Omega T) \text{ if } |\Omega| < \frac{\Omega_s}{2} \\ 0 \text{ otherwise} \end{array} \right.$$

2.8.1.2.1 Note

$G_{\text{eff}}(\Omega)$ is LTI if and only if the following two conditions are satisfied:

1. $G(\omega)$ is LTI (in DT).

2. $X_c(T)$ is bandlimited and sampling rate equal to or greater than Nyquist. For example, if we had a simple pulse described by

$$X_c(t) = u(t - T_0) - u(t - T_1)$$

where $T_1 > T_0$. If the sampling period $T > T_1 - T_0$, then some samples might "miss" the pulse while others might not be "missed." This is what we term **time-varying behavior**.

Figure 2.29

Example 2.8:

If $\frac{2\pi}{T} > 2B$ and $\omega_1 < BT$, determine and sketch $Y_c(\Omega)$ using Figure 2.29.

2.8.2 Application: 60Hz Noise Removal

Unfortunately, in real-world situations electrodes also pick up ambient 60 Hz signals from lights, computers, *etc.*. In fact, usually this "60 Hz noise" is much greater in amplitude than the EKG signal shown in Figure 2.30. Figure 2.31 shows the EKG signal; it is barely noticeable as it has become overwhelmed by noise.

2.8.2.1 DSP Solution

2.8.2.2 Sampling Period/Rate

First we must note that $|Y(\Omega)|$ is **bandlimited** to ± 60 Hz. Therefore, the minimum rate should be 120 Hz. In order to get the best results we should set

$$f_s = 240\text{Hz}$$

.

$$\Omega_s = 2\pi \left(240\frac{\text{rad}}{s}\right)$$

2.8.2.3 Digital Filter

Therefore, we want to design a digital filter that will remove the 60Hz component and preserve the rest.

Figure 2.30

**Figure 2.31:**   Our EKG signal, $y(t)$, is overwhelmed by noise.

Figure 2.32

Figure 2.33

Figure 2.34

Figure 2.35

2.9 Short Time Fourier Transform

2.9.1 Short Time Fourier Transform

The Fourier transforms (FT, DTFT, DFT, *etc.*) do not clearly indicate how the frequency content of a signal changes over time.

That information is hidden in the phase - it is not revealed by the plot of the magnitude of the spectrum.

> NOTE: To see how the frequency content of a signal changes over time, we can cut the signal into blocks and compute the spectrum of each block.

To improve the result,

1. blocks are overlapping

2. each block is multiplied by a window that is tapered at its endpoints.

Several parameters must be chosen:

- Block length, R.

- The type of window.

- Amount of overlap between blocks. (Figure 2.36)

- Amount of zero padding, if any.

The short-time Fourier transform is defined as

$$
\begin{aligned}
X\left(\omega,m\right) &= \left(\text{STFT}\left(x\left(n\right)\right) \mathbin{\text{[U+2254]}} \text{DTFT}\left(x\left(n-m\right)w\left(n\right)\right)\right) \\
&= \sum_{n=-\infty}^{\infty}\left(x\left(n-m\right)w\left(n\right)e^{-(j\omega n)}\right) \\
&= \sum_{n=0}^{R-1}\left(x\left(n-m\right)w\left(n\right)e^{-(j\omega n)}\right)
\end{aligned}
\tag{2.39}
$$

where $w\left(n\right)$ is the window function of length R.

1. The STFT of a signal $x\left(n\right)$ is a function of two variables: time and frequency.

2. The block length is determined by the support of the window function $w\left(n\right)$.

3. A graphical display of the magnitude of the STFT, $|X\left(\omega,m\right)|$, is called the **spectrogram** of the signal. It is often used in speech processing.

4. The STFT of a signal is invertible.

5. One can choose the block length. A long block length will provide higher frequency resolution (because the main-lobe of the window function will be narrow). A short block length will provide higher time resolution because less averaging across samples is performed for each STFT value.

6. A **narrow-band spectrogram** is one computed using a relatively long block length R, (long window function).

7. A **wide-band spectrogram** is one computed using a relatively short block length R, (short window function).

STFT: Overlap Parameter

NO OVERLAP

|←— DFT 1 —→|
|←— DFT 2 —→|
|←— DFT 3 —→|
|←— DFT 4 —→|

R/4 OVERLAP

|←— DFT 1 —→|
|←— DFT 2 —→|
|←— DFT 3 —→|
|←— DFT 4 —→|

R/2 OVERLAP

|←— DFT 1 —→|
|←— DFT 2 —→|
|←— DFT 3 —→|
|←— DFT 4 —→|

The parameter L

|←— DFT 2 —→|
|←— DFT 1 —→|
|←— DFT 3 —→|
|←— DFT 4 —→|

|←— L —→|←— L —→|←— L —→|

L is the number of samples between adjacent blocks.

Figure 2.36

2.9.1.1 Sampled STFT

To numerically evaluate the STFT, we sample the frequency axis ω in N equally spaced samples from $\omega = 0$ to $\omega = 2\pi$.

$$\omega_k = \frac{2\pi}{N}k \quad, \quad 0 \le k \le N - 1 \tag{2.40}$$

We then have the discrete STFT,

$$\begin{aligned}
\left(X^d\left(k, m\right) \text{ [U+2254] } X\left(\tfrac{2\pi}{N}k, m\right) \right) &= \sum_{n=0}^{R-1} \left(x\left(n-m\right) w\left(n\right) e^{-(j\omega n)} \right) \\
&= \sum_{n=0}^{R-1} \left(x\left(n-m\right) w\left(n\right) W_N^{-(kn)} \right) \\
&= \mathrm{DFT}_N \left(x\left(n-m\right) w\left(n\right) |_{n=0}^{R-1}, 0, \ldots 0 \right)
\end{aligned} \tag{2.41}$$

where $0, \ldots 0$ is $N - R$.

In this definition, the overlap between adjacent blocks is $R - 1$. The signal is shifted along the window one sample at a time. That generates more points than is usually needed, so we also sample the STFT along the time direction. That means we usually evaluate

$$X^d\left(k, Lm\right)$$

where L is the time-skip. The relation between the time-skip, the number of overlapping samples, and the block length is

$$\text{Overlap} = R - L$$

Exercise 2.3:

Match each signal to its spectrogram in Figure 2.37.

2.9.1.2 Spectrogram Example

The matlab program for producing the figures above (Figure 2.38 and Figure 2.39).

```
% LOAD DATA
load mtlb;
x = mtlb;

figure(1), clf
plot(0:4000,x)
xlabel('n')
ylabel('x(n)')

% SET PARAMETERS
R = 256;                % R: block length
window = hamming(R);    % window function of length R
N = 512;                % N: frequency discretization
L = 35;                 % L: time lapse between blocks
fs = 7418;              % fs: sampling frequency
overlap = R - L;
```

(a)

Figure 2.38

Figure 2.39

```
% COMPUTE SPECTROGRAM
[B,f,t] = specgram(x,N,fs,window,overlap);

% MAKE PLOT
figure(2), clf
imagesc(t,f,log10(abs(B)));
colormap('jet')
axis xy
xlabel('time')
ylabel('frequency')
title('SPECTROGRAM, R = 256')
```

2.9.1.3 Effect of window length R

Here is another example to illustrate the frequency/time resolution trade-off (See figures - Figure 2.40, Figure 2.41, and Figure 2.42).

2.9.1.4 Effect of L and N

A spectrogram is computed with different parameters:

$$L \in \{1, 10\}$$

$$N \in \{32, 256\}$$

- L = time lapse between blocks.
- N = FFT length (Each block is zero-padded to length N.)

In each case, the block length is 30 samples.

Exercise 2.4:

For each of the four spectrograms in Figure 2.43 can you tell what L and N are?

L and N do not effect the time resolution or the frequency resolution. They only affect the 'pixelation'.

2.9.1.5 Effect of R and L

Shown below are four spectrograms of the same signal. Each spectrogram is computed using a different set of parameters.

$$R \in \{120, 256, 1024\}$$

$$L \in \{35, 250\}$$

where

- R = block length
- L = time lapse between blocks.

**Narrow-band spectrogram: better frequency resolution**

Figure 2.40

Wide-band spectrogram: better time resolution


Figure 2.41

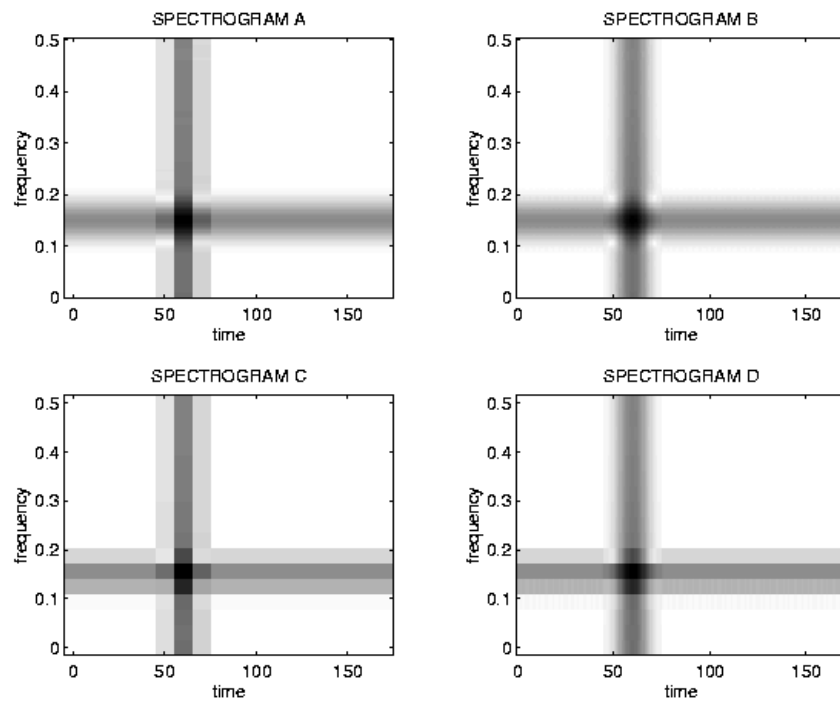**Effect of Window Length R**

(a)

(b)

Figure 2.42

(a)

(b)
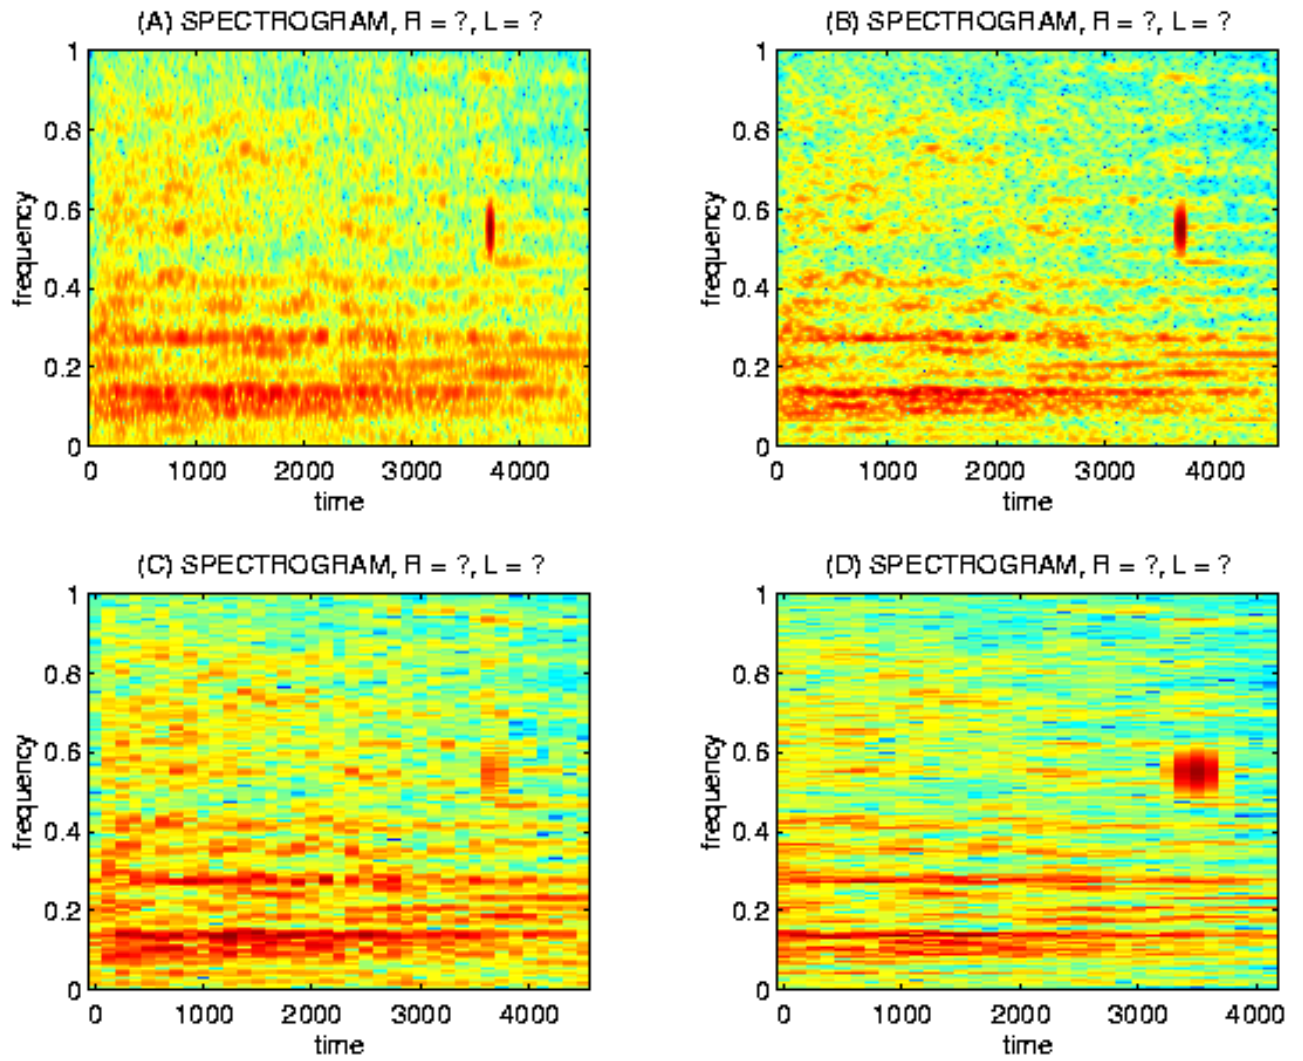
Figure 2.43

Figure 2.44

Figure 2.45

Exercise 2.5:

For each of the four spectrograms in Figure 2.44, match the above values of L and R.

If you like, you may listen to this signal with the `soundsc` command; the data is in the file: `stft_data.m`. Here (Figure 2.45) is a figure of the signal.

2.10 Spectrograms

We know how to acquire analog signals for digital processing (pre-filtering[34], sampling[35], and A/D conversion[36]) and to compute spectra of discrete-time signals (using the FFT algorithm[37]), let's put these various components together to learn how the spectrogram shown in Figure 2.46, which is used to analyze speech[38], is calculated. The speech was sampled at a rate of 11.025 kHz and passed through a 16-bit A/D converter.

> POINT OF INTEREST: Music compact discs (CDs) encode their signals at a sampling rate of 44.1 kHz. We'll learn the rationale for this number later. The 11.025 kHz sampling rate for the speech is 1/4 of the CD sampling rate, and was the lowest available sampling rate commensurate with speech signal bandwidths available on my computer.

Exercise 2.6:

Looking at Figure 2.46 the signal lasted a little over 1.2 seconds. How long was the sampled signal (in terms of samples)? What was the datarate during the sampling process in bps (bits per second)? Assuming the computer storage is organized in terms of bytes (8-bit quantities), how many bytes of computer memory does the speech consume?

---

[34] http://cnx.org/content/m0050/latest/#five
[35] http://cnx.org/content/m0050/latest/
[36] http://cnx.org/content/m0051/latest/
[37] http://cnx.org/content/m10250/latest/
[38] http://cnx.org/content/m0049/latest/

Figure 2.46

Figure 2.47: The top waveform is a segment 1024 samples long taken from the beginning of the "Rice University" phrase. Computing Figure 2.46 involved creating frames, here demarked by the vertical lines, that were 256 samples long and finding the spectrum of each. If a rectangular window is applied (corresponding to extracting a frame from the signal), oscillations appear in the spectrum (middle of bottom row). Applying a Hanning window gracefully tapers the signal toward frame edges, thereby yielding a more accurate computation of the signal's spectrum at that moment of time.

The resulting discrete-time signal, shown in the bottom of Figure 2.46, clearly changes its character with time. To display these spectral changes, the long signal was sectioned into **frames**: comparatively short, contiguous groups of samples. Conceptually, a Fourier transform of each frame is calculated using the FFT. Each frame is not so long that significant signal variations are retained within a frame, but not so short that we lose the signal's spectral character.

An important detail emerges when we examine each framed signal (Figure 2.47). At the frame's edges, the signal may change very abruptly, a feature not present in the original signal. A transform of such a segment reveals a curious oscillation in the spectrum, an artifact directly related to this sharp amplitude change. A better way to frame signals for spectrograms is to apply a **window**: Shape the signal values within a frame so that the signal decays gracefully as it nears the edges. This shaping is accomplished by multiplying the framed signal by the sequence $w(n)$. In sectioning the signal, we essentially applied a rectangular window: $w(n) = 1$, $0 \leq n \leq N - 1$. A much more graceful window is the **Hanning window**; it has the cosine shape $w(n) = \frac{1}{2}\left(1 - \cos\left(\frac{2\pi n}{N}\right)\right)$. As shown in Figure 2.47, this shaping greatly reduces spurious oscillations in each frame's spectrum. Considering the spectrum of the Hanning windowed frame, we find that the oscillations

Hanning speech

Figure 2.48: In comparison with the original speech segment shown in the upper plot, the non-overlapped Hanning windowed version shown below it is very ragged. Clearly, spectral information extracted from the bottom plot could well miss important features present in the original.

resulting from applying the rectangular window obscured a formant (the one located at a little more than half the Nyquist frequency).

Exercise 2.7:

What might be the source of these oscillations? To gain some insight, what is the length- $2N$ discrete Fourier transform of a length-N pulse? The pulse emulates the rectangular window, and certainly has edges. Compare your answer with the length- $2N$ transform of a length- N Hanning window.

If you examine the windowed signal sections in sequence to examine windowing's affect on signal amplitude, we see that we have managed to amplitude-modulate the signal with the periodically repeated window (Figure 2.48). To alleviate this problem, frames are overlapped (typically by half a frame duration). This solution requires more Fourier transform calculations than needed by rectangular windowing, but the spectra are much better behaved and spectral changes are much better captured.

Spectrograms, such as shown in the speech spectrogram (Figure 2.46), are sectioned into overlapping, equal-length frames, with a Hanning window applied to each frame. The spectra of each of these is calculated, and displayed in spectrograms with frequency extending vertically, window time location running horizontally, and spectral magnitude color-coded. Figure 2.49 illustrates these computations.

Exercise 2.8:

Why the specific values of 256 for N and 512 for K? Another issue is how was the length-512 transform of each length-256 windowed frame computed?

2.11 Filtering with the DFT

2.11.1 Introduction

$$
\begin{aligned}
y\left[n\right] &= x\left[n\right] * h\left[n\right] \\
&= \sum_{k=-\infty}^{\infty} \left(x\left[k\right] h\left[n-k\right] \right)
\end{aligned}
\tag{2.42}
$$

Hanning windows

Figure 2.49: The original speech segment and the sequence of overlapping Hanning windows applied to it are shown in the upper portion. Frames were 256 samples long and a Hanning window was applied with a half-frame overlap. A length-512 FFT of each frame was computed, with the magnitude of the first 257 FFT values displayed vertically, with spectral amplitude values color-coded.

Figure 2.50

Figure 2.51

$$Y\left(\omega\right) = X\left(\omega\right) H\left(\omega\right) \tag{2.43}$$

Assume that $H\left(\omega\right)$ is specified.

Exercise 2.9:

How can we implement $X\left(\omega\right) H\left(\omega\right)$ in a computer?

Recall that the DFT treats N-point sequences as if they are periodically extended (Figure 2.51):

Figure 2.52

$$\tilde{\tilde{y}}[n] = x[n] \,\, \textcircled{N} \,\, h[n]$$

Figure 2.53: The above symbol for the circular convolution is for an N-periodic extension.

2.11.2 Compute IDFT of Y[k]

$$
\begin{aligned}
\tilde{y}[n] &= \frac{1}{N} \sum_{k=0}^{N-1} \left(Y[k] \, e^{j2\pi \frac{k}{N} n} \right) \\
&= \frac{1}{N} \sum_{k=0}^{N-1} \left(X[k] \, H[k] \, e^{j2\pi \frac{k}{N} n} \right) \\
&= \frac{1}{N} \sum_{k=0}^{N-1} \left(\left(\sum_{m=0}^{N-1} \left(x[m] \, e^{-\left(j2\pi \frac{k}{N} m \right)} \right) \right) H[k] \, e^{j2\pi \frac{k}{N} n} \right) \\
&= \sum_{m=0}^{N-1} \left(x[m] \left(\frac{1}{N} \sum_{k=0}^{N-1} \left(H[k] \, e^{j2\pi \frac{k}{N}(n-m)} \right) \right) \right) \\
&= \sum_{m=0}^{N-1} \left(x[m] \, h\left[((n-m))_N\right] \right)
\end{aligned}
\tag{2.44}
$$

And the IDFT periodically extends $h[n]$:

$$\tilde{h}[n-m] = h\left[((n-m))_N\right]$$

This computes as shown in Figure 2.52:

$$\tilde{y}[n] = \sum_{m=0}^{N-1} \left(x[m] \, h\left[((n-m))_N\right] \right) \tag{2.45}$$

is called **circular convolution** and is denoted by Figure 2.53.

Figure 2.54

Figure 2.55

2.11.2.1 DFT Pair

Note that in general:

Example 2.9:   Regular vs. Circular Convolution

To begin with, we are given the following two length-3 signals:

$$x[n] = \{1, 2, 3\}$$

$$h[n] = \{1, 0, 2\}$$

We can zero-pad these signals so that we have the following discrete sequences:

$$x[n] = \{\ldots, 0, 1, 2, 3, 0, \ldots\}$$

$$h[n] = \{\ldots, 0, 1, 0, 2, 0, \ldots\}$$

where $x[0] = 1$ and $h[0] = 1$.

- Regular Convolution:

$$y[n] = \sum_{m=0}^{2} (x[m] \, h[n-m]) \tag{2.46}$$

Using the above convolution formula (refer to the link if you need a review of convolution (Section 2.5)), we can calculate the resulting value for $y[0]$ to $y[4]$. Recall that because we have two length-3 signals, our convolved signal will be length-5.

    · $n = 0$

$$\{\ldots, 0, 0, 0, 1, 2, 3, 0, \ldots\}$$

$$\{\ldots, 0, 2, 0, 1, 0, 0, 0, \ldots\}$$

$$\begin{aligned} y[0] &= 1 \times 1 + 2 \times 0 + 3 \times 0 \\ &= 1 \end{aligned} \tag{2.47}$$

Regular Convolution Result

Figure 2.56: Result is finite duration, not periodic!

· $n = 1$

$$\{\dots, 0, 0, 1, 2, 3, 0, \dots\}$$
$$\{\dots, 0, 2, 0, 1, 0, 0, \dots\}$$

$$
\begin{aligned}
y[1] &= 1 \times 0 + 2 \times 1 + 3 \times 0 \\
&= 2
\end{aligned}
\tag{2.48}
$$

· $n = 2$

$$\{\dots, 0, 1, 2, 3, 0, \dots\}$$
$$\{\dots, 0, 2, 0, 1, 0, \dots\}$$

$$
\begin{aligned}
y[2] &= 1 \times 2 + 2 \times 0 + 3 \times 1 \\
&= 5
\end{aligned}
\tag{2.49}
$$

· $n = 3$

$$y[3] = 4 \tag{2.50}$$

· $n = 4$

$$y[4] = 6 \tag{2.51}$$

- Circular Convolution:

$$\tilde{y}[n] = \sum_{m=0}^{2} \left(x[m] \, h\left[((n-m))_N\right] \right) \tag{2.52}$$

And now with circular convolution our $h[n]$ changes and becomes a periodically extended signal:

$$h\left[((n))_N\right] = \{\dots, 1, 0, 2, 1, 0, 2, 1, 0, 2, \dots\} \tag{2.53}$$

· $n = 0$

$$\{\dots, 0, 0, 0, 1, 2, 3, 0, \dots\}$$
$$\{\dots, 1, 2, 0, 1, 2, 0, 1, \dots\}$$

$$
\begin{aligned}
\tilde{y}[0] &= 1 \times 1 + 2 \times 2 + 3 \times 0 \\
&= 5
\end{aligned}
\tag{2.54}
$$

Circular Convolution Result

Figure 2.57:   Result is 3-periodic.

Circular Convolution from Regular

Figure 2.58:   The left plot (the circular convolution results) has a "wrap-around" effect due to periodic extension.

$\cdot\, n = 1$

$$\{\ldots, 0, 0, 0, 1, 2, 3, 0, \ldots\}$$

$$\{\ldots, 0, 1, 2, 0, 1, 2, 0, \ldots\}$$

$$\tilde{y}\,[1] \quad = \quad 1 \times 1 + 2 \times 1 + 3 \times 2 \qquad\qquad (2.55)$$
$$= \quad 8$$

$\cdot\, n = 2$

$$\tilde{y}\,[2] = 5 \qquad\qquad (2.56)$$

$\cdot\, n = 3$

$$\tilde{y}\,[3] = 5 \qquad\qquad (2.57)$$

$\cdot\, n = 4$

$$\tilde{y}\,[4] = 8 \qquad\qquad (2.58)$$

Figure 2.58 illustrates the relationship between circular convolution and regular convolution using the previous two figures:

No Image.

Figure 2.59:   The sequence from 0 to 4 (the underlined part of the sequence) is the regular convolution result. From this illustration we can see that it is 5-periodic!

2.11.2.2 Regular Convolution from Periodic Convolution

1. "Zero-pad" $x[n]$ and $h[n]$ to avoid the overlap (wrap-around) effect. We will zero-pad the two signals to a length-5 signal (5 being the duration of the regular convolution result):

$$x[n] = \{1, 2, 3, 0, 0\}$$

$$h[n] = \{1, 0, 2, 0, 0\}$$

2. Now take the DFTs of the zero-padded signals:

$$
\begin{aligned}
\tilde{y}[n] &= \tfrac{1}{N} \sum_{k=0}^{4} \left(X[k] H[k] e^{j2\pi \frac{k}{5} n} \right) \\
&= \sum_{m=0}^{4} \left(x[m] h[((n-m))_5] \right)
\end{aligned}
\tag{2.59}
$$

Now we can plot this result (Figure 2.59):

GENERAL RESULT: We can compute the regular convolution result of a convolution of an M-point signal $x[n]$ with an N-point signal $h[n]$ by padding each signal with zeros to obtain two $M + N - 1$ length sequences and computing the circular convolution (or equivalently computing the IDFT of $H[k] X[k]$, the product of the DFTs of the zero-padded signals) (Figure 2.60).

2.11.3 DSP System

1. Sample finite duration continuous-time input $x(t)$ to get $x[n]$ where $n = \{0, \ldots, M-1\}$.

2. Zero-pad $x[n]$ and $h[n]$ to length $M + N - 1$.

3. Compute DFTs $X[k]$ and $H[k]$

4. Compute IDFTs of $X[k] H[k]$

$$y[n] = \tilde{y}[n]$$

where $n = \{0, \ldots, M + N - 1\}$.

5. Reconstruct $y(t)$

Figure 2.60: Note that the lower two images are simply the top images that have been zero-padded.

No Image.

Figure 2.61: The system has a length N impulse response, $h[n]$

Figure 2.62: Fourier Transform (FT) relationship between the two functions.

(a)   (b)

Figure 2.63

2.12 Image Restoration Basics

2.12.1 Image Restoration

In many applications (*e.g.*, satellite imaging, medical imaging, astronomical imaging, poor-quality family portraits) the imaging system introduces a slight distortion. Often images are slightly blurred and image restoration aims at **deblurring** the image.

The blurring can usually be modeled as an LSI system with a given PSF $h\,[m,n]$.

The observed image is

$$g\,[m,n] = h\,[m,n] * f\,[m,n] \tag{2.60}$$

$$G\,(u,v) = H\,(u,v)\,F\,(u,v) \tag{2.61}$$

$$F\,(u,v) = \frac{G\,(u,v)}{H\,(u,v)} \tag{2.62}$$

Example 2.10:   Image Blurring

Above we showed the equations for representing the common model for blurring an image. In Figure 2.63 we have an original image and a PSF function that we wish to apply to the image in order to model a basic blurred image.

Once we apply the PSF to the original image, we receive our blurred image that is shown in Figure 2.64:

Figure 2.64

2.12.1.1 Frequency Domain Analysis

Example 2.10 looks at the original images in its typical form; however, it is often useful to look at our images and PSF in the frequency domain. In Figure 2.65, we take another look at the image blurring example above and look at how the images and results would appear in the frequency domain if we applied the fourier transforms.
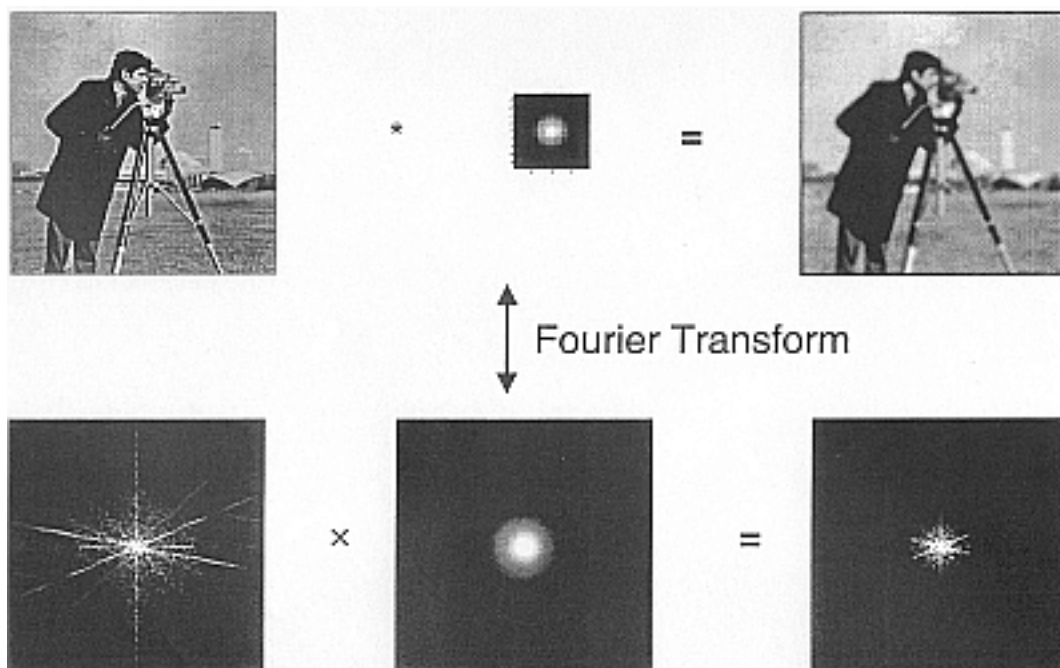
Figure 2.65

Solutions to Exercises in Chapter 2

Solution to Exercise 2.1:

$$x[n+N] = ???$$

Solution to Exercise 2.2:

$S[n]$ is N-periodic, so it has the following Fourier Series[39]:

$$\begin{aligned} c_k &= \frac{1}{N} \int_{-\left(\frac{N}{2}\right)}^{\frac{N}{2}} \delta[n] e^{(-j)2\pi \frac{k}{N} n} dn \\ &= \frac{1}{N} \end{aligned} \tag{2.63}$$

$$S[n] = \sum_{k=-\infty}^{\infty} \left(e^{(-j)2\pi \frac{k}{N} n} \right) \tag{2.64}$$

where the DTFT of the exponential in the above equation is equal to $\delta\left(\omega - \frac{2\pi k}{N}\right)$.

Solution to Exercise 2.3:

Solution to Exercise 2.4:

Solution to Exercise 2.5:

Solution to Exercise 2.6:

Number of samples equals $1.2 \times 11025 = 13230$. The datarate is $11025 \times 16 = 176.4$ kbps. The storage required would be 26460 bytes.

Solution to Exercise 2.7:

The oscillations are due to the boxcar window's Fourier transform, which equals the sinc function.

Solution to Exercise 2.8:

These numbers are powers-of-two, and the FFT algorithm can be exploited with these lengths. To compute a longer transform than the input signal's duration, we simply zero-pad the signal.

Solution to Exercise 2.9:

Discretize (sample) $X(\omega)$ and $H(\omega)$. In order to do this, we should take the DFTs of $x[n]$ and $h[n]$ to get $X[k]$ and $X[k]$. Then we will compute

$$\tilde{y}[n] = \text{IDFT}(X[k] H[k])$$

Does $\tilde{y}[n] = y[n]$?

---

[39]http://cnx.org/content/m10496/latest/

Chapter 3

Digital Filtering

3.1 Difference Equation

3.1.1 Introduction

One of the most important concepts of DSP is to be able to properly represent the input/output relationship to a given LTI system. A linear constant-coefficient **difference equation** (LCCDE) serves as a way to express just this relationship in a discrete-time system. Writing the sequence of inputs and outputs, which represent the characteristics of the LTI system, as a difference equation help in understanding and manipulating a system.

> **Definition 2: difference equation**
> An equation that shows the relationship between consecutive values of a sequence and the differences among them. They are often rearranged as a recursive formula so that a systems output can be computed from the input signal and past outputs.

Example:

$$y[n] + 7y[n-1] + 2y[n-2] = x[n] - 4x[n-1] \qquad (3.1)$$

3.1.2 General Formulas from the Difference Equation

As stated briefly in the definition above, a difference equation is a very useful tool in describing and calculating the output of the system described by the formula for a given sample n. The key property of the difference equation is its ability to help easily find the transform, $H(z)$, of a system. In the following two subsections, we will look at the general form of the difference equation and the general conversion to a z-transform directly from the difference equation.

3.1.2.1 Difference Equation

The general form of a linear, constant-coefficient difference equation (LCCDE), is shown below:

$$\sum_{k=0}^{N} (a_k y[n-k]) = \sum_{k=0}^{M} (b_k x[n-k]) \qquad (3.2)$$

We can also write the general form to easily express a recursive output, which looks like this:

$$y[n] = -\left(\sum_{k=1}^{N}\left(a_k y[n-k]\right)\right) + \sum_{k=0}^{M}\left(b_k x[n-k]\right) \tag{3.3}$$

From this equation, note that $y[n-k]$ represents the outputs and $x[n-k]$ represents the inputs. The value of N represents the **order** of the difference equation and corresponds to the memory of the system being represented. Because this equation relies on past values of the output, in order to compute a numerical solution, certain past outputs, referred to as the **initial conditions**, must be known.

3.1.2.2 Conversion to Z-Transform

Using the above formula, Equation 3.2, we can easily generalize the **transfer function**, $H(z)$, for any difference equation. Below are the steps taken to convert any difference equation into its transfer function, *i.e.* z-transform. The first step involves taking the Fourier Transform[1] of all the terms in Equation 3.2. Then we use the linearity property to pull the transform inside the summation and the time-shifting property of the z-transform to change the time-shifting terms to exponentials. Once this is done, we arrive at the following equation: $a_0 = 1$.

$$Y(z) = -\left(\sum_{k=1}^{N}\left(a_k Y(z) z^{-k}\right)\right) + \sum_{k=0}^{M}\left(b_k X(z) z^{-k}\right) \tag{3.4}$$

$$\begin{aligned} H(z) &= \frac{Y(z)}{X(z)} \\ &= \frac{\sum_{k=0}^{M}\left(b_k z^{-k}\right)}{1+\sum_{k=1}^{N}\left(a_k z^{-k}\right)} \end{aligned} \tag{3.5}$$

3.1.2.3 Conversion to Frequency Response

Once the z-transform has been calculated from the difference equation, we can go one step further to define the frequency response of the system, or filter, that is being represented by the difference equation.

> NOTE: Remember that the reason we are dealing with these formulas is to be able to aid us in filter design. A LCCDE is one of the easiest ways to represent FIR filters. By being able to find the frequency response, we will be able to look at the basic properties of any filter represented by a simple LCCDE.

Below is the general formula for the frequency response of a z-transform. The conversion is simple a matter of taking the z-transform formula, $H(z)$, and replacing every instance of z with e^{jw}.

$$\begin{aligned} H(w) &= H(z)\big|_{z,z=e^{jw}} \\ &= \frac{\sum_{k=0}^{M}\left(b_k e^{-(jwk)}\right)}{\sum_{k=0}^{N}\left(a_k e^{-(jwk)}\right)} \end{aligned} \tag{3.6}$$

Once you understand the derivation of this formula, look at the module concerning Filter Design from the Z-Transform[2] for a look into how all of these ideas of the Z-transform (Section 4.2), Difference Equation, and Pole/Zero Plots (Section 4.4) play a role in filter design.

---

[1] http://cnx.org/content/m0046/latest/
[2] http://cnx.org/content/m10548/latest/

3.1.3 Example

Example 3.1:   Finding Difference Equation

Below is a basic example showing the opposite of the steps above: given a transfer function one can easily calculate the systems difference equation.

$$H\left(z\right) = \frac{\left(z+1\right)^2}{\left(z-\frac{1}{2}\right)\left(z+\frac{3}{4}\right)} \tag{3.7}$$

Given this transfer function of a time-domain filter, we want to find the difference equation. To begin with, expand both polynomials and divide them by the highest order z.

$$\begin{aligned} H\left(z\right) &= \frac{\left(z+1\right)\left(z+1\right)}{\left(z-\frac{1}{2}\right)\left(z+\frac{3}{4}\right)} \\ &= \frac{z^2+2z+1}{z^2+2z+1-\frac{3}{8}} \\ &= \frac{1+2z^{-1}+z^{-2}}{1+\frac{1}{4}z^{-1}-\frac{3}{8}z^{-2}} \end{aligned} \tag{3.8}$$

From this transfer function, the coefficients of the two polynomials will be our a_k and b_k values found in the general difference equation formula, Equation 3.2. Using these coefficients and the above form of the transfer function, we can easily write the difference equation:

$$x\left[n\right] + 2x\left[n-1\right] + x\left[n-2\right] = y\left[n\right] + \frac{1}{4}y\left[n-1\right] - \frac{3}{8}y\left[n-2\right] \tag{3.9}$$

In our final step, we can rewrite the difference equation in its more common form showing the recursive nature of the system.

$$y\left[n\right] = x\left[n\right] + 2x\left[n-1\right] + x\left[n-2\right] + \frac{-1}{4}y\left[n-1\right] + \frac{3}{8}y\left[n-2\right] \tag{3.10}$$

3.1.4 Solving a LCCDE

In order for a linear constant-coefficient difference equation to be useful in analyzing a LTI system, we must be able to find the systems output based upon a known input, $x\left(n\right)$, and a set of initial conditions. Two common methods exist for solving a LCCDE: the **direct method** and the **indirect method**, the later being based on the z-transform. Below we will briefly discuss the formulas for solving a LCCDE using each of these methods.

3.1.4.1 Direct Method

The final solution to the output based on the direct method is the sum of two parts, expressed in the following equation:

$$y\left(n\right) = y_h\left(n\right) + y_p\left(n\right) \tag{3.11}$$

The first part, $y_h\left(n\right)$, is referred to as the **homogeneous solution** and the second part, $y_h\left(n\right)$, is referred to as **particular solution**. The following method is very similar to that used to solve many differential equations, so if you have taken a differential calculus course or used differential equations before then this should seem very familiar.

3.1.4.1.1 Homogeneous Solution

We begin by assuming that the input is zero, $x(n) = 0$. Now we simply need to solve the homogeneous difference equation:

$$\sum_{k=0}^{N} (a_k y[n-k]) = 0 \tag{3.12}$$

In order to solve this, we will make the assumption that the solution is in the form of an exponential. We will use lambda, λ, to represent our exponential terms. We now have to solve the following equation:

$$\sum_{k=0}^{N} \left(a_k \lambda^{n-k}\right) = 0 \tag{3.13}$$

We can expand this equation out and factor out all of the lambda terms. This will give us a large polynomial in parenthesis, which is referred to as the **characteristic polynomial**. The roots of this polynomial will be the key to solving the homogeneous equation. If there are all distinct roots, then the general solution to the equation will be as follows:

$$y_h(n) = C_1(\lambda_1)^n + C_2(\lambda_2)^n + \cdots + C_N(\lambda_N)^n \tag{3.14}$$

However, if the characteristic equation contains multiple roots then the above general solution will be slightly different. Below we have the modified version for an equation where λ_1 has K multiple roots:

$$y_h(n) = C_1(\lambda_1)^n + C_1 n(\lambda_1)^n + C_1 n^2(\lambda_1)^n + \cdots + C_1 n^{K-1}(\lambda_1)^n + C_2(\lambda_2)^n + \cdots + C_N(\lambda_N)^n \tag{3.15}$$

3.1.4.1.2 Particular Solution

The particular solution, $y_p(n)$, will be any solution that will solve the general difference equation:

$$\sum_{k=0}^{N} (a_k y_p(n-k)) = \sum_{k=0}^{M} (b_k x(n-k)) \tag{3.16}$$

In order to solve, our guess for the solution to $y_p(n)$ will take on the form of the input, $x(n)$. After guessing at a solution to the above equation involving the particular solution, one only needs to plug the solution into the difference equation and solve it out.

3.1.4.2 Indirect Method

The indirect method utilizes the relationship between the difference equation and z-transform, discussed earlier (Section 3.1.2), to find a solution. The basic idea is to convert the difference equation into a z-transform, as described above (Section 3.1.2.2), to get the resulting output, $Y(z)$. Then by inverse transforming this and using partial-fraction expansion, we can arrive at the solution.

3.2 The Z Transform: Definition

3.2.1 Basic Definition of the Z-Transform

The **z-transform** of a sequence is defined as

$$X(z) = \sum_{n=-\infty}^{\infty} \left(x[n] z^{-n} \right) \tag{3.17}$$

Sometimes this equation is referred to as the **bilateral z-transform**. At times the z-transform is defined as

$$X(z) = \sum_{n=0}^{\infty} \left(x[n] z^{-n} \right) \tag{3.18}$$

which is known as the **unilateral z-transform**.

There is a close relationship between the z-transform and the **Fourier transform** of a discrete time signal, which is defined as

$$X\left(e^{j\omega}\right) = \sum_{n=-\infty}^{\infty} \left(x[n] e^{-(j\omega n)} \right) \tag{3.19}$$

Notice that that when the z^{-n} is replaced with $e^{-(j\omega n)}$ the z-transform reduces to the Fourier Transform. When the Fourier Transform exists, $z = e^{j\omega}$, which is to have the magnitude of z equal to unity.

3.2.2 The Complex Plane

In order to get further insight into the relationship between the Fourier Transform and the Z-Transform it is useful to look at the complex plane or **z-plane**. Take a look at the complex plane:

The Z-plane is a complex plane with an imaginary and real axis referring to the complex-valued variable z. The position on the complex plane is given by $re^{j\omega}$, and the angle from the positive, real axis around the plane is denoted by ω. $X(z)$ is defined everywhere on this plane. $X\left(e^{j\omega}\right)$ on the other hand is defined only where $|z| = 1$, which is referred to as the unit circle. So for example, $\omega = 1$ at $z = 1$ and $\omega = \pi$ at $z = -1$. This is useful because, by representing the Fourier transform as the z-transform on the unit circle, the periodicity of Fourier transform is easily seen.

3.2.3 Region of Convergence

The region of convergence, known as the **ROC**, is important to understand because it defines the region where the z-transform exists. The ROC for a given $x[n]$, is defined as the range of z for which the z-transform converges. Since the z-transform is a **power series**, it converges when $x[n] z^{-n}$ is absolutely summable. Stated differently,

$$\sum_{n=-\infty}^{\infty} \left(|x[n] z^{-n}| \right) < \infty \tag{3.20}$$

must be satisfied for convergence. This is best illustrated by looking at the different ROC's of the z-transforms of $\alpha^n u[n]$ and $\alpha^n u[n-1]$.

Figure 3.1

Figure 3.2: $x[n] = \alpha^n u[n]$ where $\alpha = 0.5$.

Figure 3.3: ROC for $x[n] = \alpha^n u[n]$ where $\alpha = 0.5$

Example 3.2:

For

$$x[n] = \alpha^n u[n] \tag{3.21}$$

$$
\begin{aligned}
X(z) &= \sum_{n=-\infty}^{\infty} \left(x[n] z^{-n} \right) \\
&= \sum_{n=-\infty}^{\infty} \left(\alpha^n u[n] z^{-n} \right) \\
&= \sum_{n=0}^{\infty} \left(\alpha^n z^{-n} \right) \\
&= \sum_{n=0}^{\infty} \left(\left(\alpha z^{-1} \right)^n \right)
\end{aligned} \tag{3.22}
$$

This sequence is an example of a right-sided exponential sequence because it is nonzero for $n \geq 0$. It only converges when $\left| \alpha z^{-1} \right| < 1$. When it converges,

$$
\begin{aligned}
X(z) &= \frac{1}{1 - \alpha z^{-1}} \\
&= \frac{z}{z - \alpha}
\end{aligned} \tag{3.23}
$$

If $\left| \alpha z^{-1} \right| \geq 1$, then the series, $\sum_{n=0}^{\infty} \left(\left(\alpha z^{-1} \right)^n \right)$ does not converge. Thus the ROC is the range of values where

$$\left| \alpha z^{-1} \right| < 1 \tag{3.24}$$

or, equivalently,

$$|z| > |\alpha| \tag{3.25}$$

Example 3.3:

For

$$x[n] = \left(-\left(\alpha^n \right) \right) u[-n-1] \tag{3.26}$$

Figure 3.4: $x[n] = (-(\alpha^n))u[-n-1]$ where $\alpha = 0.5$.

$$
\begin{array}{rcl}
X(z) & = & \sum_{n=-\infty}^{\infty}(x[n]z^{-n}) \\
     & = & \sum_{n=-\infty}^{\infty}((-(\alpha^n))u[-\text{n}-1]z^{-n}) \\
     & = & -\left(\sum_{n=-\infty}^{-1}(\alpha^n z^{-n})\right) \\
     & = & -\left(\sum_{n=-\infty}^{-1}\left((\alpha^{-1}z)^{-n}\right)\right) \\
     & = & -\left(\sum_{n=1}^{\infty}\left((\alpha^{-1}z)^{n}\right)\right) \\
     & = & 1 - \sum_{n=0}^{\infty}\left((\alpha^{-1}z)^{n}\right)
\end{array}
\tag{3.27}
$$

The ROC in this case is the range of values where

$$
|\alpha^{-1}z| < 1 \tag{3.28}
$$

or, equivalently,

$$
|z| < |\alpha| \tag{3.29}
$$

If the ROC is satisfied, then

$$
\begin{array}{rcl}
X(z) & = & 1 - \frac{1}{1-\alpha^{-1}z} \\
     & = & \frac{z}{z-\alpha}
\end{array}
\tag{3.30}
$$

3.3 Table of Common z-Transforms

The table below provides a number of unilateral and bilateral **z-transforms (Section 4.2)**. The table also specifies the region of convergence[3].

---

[3]http://cnx.org/content/m10622/latest/

Figure 3.5: ROC for $x[n] = (-(\alpha^n))\,u[-n-1]$

NOTE: The notation for z found in the table below may differ from that found in other tables. For example, the basic z-transform of $u[n]$ can be written as either of the following two expressions, which are equivalent:

$$\frac{z}{z-1} = \frac{1}{1-z^{-1}} \tag{3.31}$$

Signal	Z-Transform	ROC
$\delta[n-k]$	z^{-k}	All z
$u[n]$	$\frac{z}{z-1}$	$|z| > 1$
$-(u[-n-1])$	$\frac{z}{z-1}$	$|z| < 1$
$nu[n]$	$\frac{z}{(z-1)^2}$	$|z| > 1$
$n^2 u[n]$	$\frac{z(z+1)}{(z-1)^3}$	$|z| > 1$
$n^3 u[n]$	$\frac{z(z^2+4z+1)}{(z-1)^4}$	$|z| > 1$
$(-(\alpha^n))\,u[-n-1]$	$\frac{z}{z-\alpha}$	$|z| < |\alpha|$
$\alpha^n u[n]$	$\frac{z}{z-\alpha}$	$|z| > |\alpha|$
$n\alpha^n u[n]$	$\frac{\alpha z}{(z-\alpha)^2}$	$|z| > |\alpha|$
$n^2 \alpha^n u[n]$	$\frac{\alpha z(z+\alpha)}{(z-\alpha)^3}$	$|z| > |\alpha|$
$\frac{\prod_{k=1}^{m}(n-k+1)}{\alpha^m m!}\alpha^n u[n]$	$\frac{z}{(z-\alpha)^{m+1}}$	
$\gamma^n \cos(\alpha n)\,u[n]$	$\frac{z(z-\gamma\cos(\alpha))}{z^2-(2\gamma\cos(\alpha))z+\gamma^2}$	$|z| > |\gamma|$
$\gamma^n \sin(\alpha n)\,u[n]$	$\frac{z\gamma\sin(\alpha)}{z^2-(2\gamma\cos(\alpha))z+\gamma^2}$	$|z| > |\gamma|$

3.4 Understanding Pole/Zero Plots on the Z-Plane

3.4.1 Introduction to Poles and Zeros of the Z-Transform

Once the Z-transform of a system has been determined, one can use the information contained in function's polynomials to graphically represent the function and easily observe many defining characteristics. The Z-transform will have the below structure, based on Rational Functions[4]:

$$X(z) = \frac{P(z)}{Q(z)} \tag{3.32}$$

The two polynomials, $P(z)$ and $Q(z)$, allow us to find the poles and zeros[5] of the Z-Transform.

Definition 3: zeros

1. The value(s) for z where $P(z) = 0$.
2. The complex frequencies that make the overall gain of the filter transfer function zero.

Definition 4: poles

1. The value(s) for z where $Q(z) = 0$.
2. The complex frequencies that make the overall gain of the filter transfer function infinite.

Example 3.4:

Below is a simple transfer function with the poles and zeros shown below it.

$$H(z) = \frac{z+1}{\left(z - \frac{1}{2}\right)\left(z + \frac{3}{4}\right)}$$

The zeros are: $\{-1\}$

The poles are: $\left\{\frac{1}{2}, -\left(\frac{3}{4}\right)\right\}$

3.4.2 The Z-Plane

Once the poles and zeros have been found for a given Z-Transform, they can be plotted onto the Z-Plane. The Z-plane is a complex plane with an imaginary and real axis referring to the complex-valued variable z. The position on the complex plane is given by $re^{j\theta}$ and the angle from the positive, real axis around the plane is denoted by θ. When mapping poles and zeros onto the plane, poles are denoted by an "x" and zeros by an "o". The below figure shows the Z-Plane, and examples of plotting zeros and poles onto the plane can be found in the following section.

---

[4]http://cnx.org/content/m10593/latest/
[5]http://cnx.org/content/m10112/latest/

Z-Plane

Figure 3.6

3.4.3 Examples of Pole/Zero Plots

This section lists several examples of finding the poles and zeros of a transfer function and then plotting them onto the Z-Plane.

Example 3.5:   Simple Pole/Zero Plot

$$H(z) = \frac{z}{\left(z - \frac{1}{2}\right)\left(z + \frac{3}{4}\right)}$$

The zeros are: $\{0\}$

The poles are: $\left\{\frac{1}{2}, -\left(\frac{3}{4}\right)\right\}$

Example 3.6:   Complex Pole/Zero Plot

$$H(z) = \frac{(z - j)(z + j)}{\left(z - \left(\frac{1}{2} - \frac{1}{2}j\right)\right)\left(z - \left(\frac{1}{2} + \frac{1}{2}j\right)\right)}$$

The zeros are: $\{j, -j\}$

The poles are: $\left\{-1, \frac{1}{2} + \frac{1}{2}j, \frac{1}{2} - \frac{1}{2}j\right\}$

MATLAB - If access to MATLAB is readily available, then you can use its functions to easily create pole/zero plots. Below is a short program that plots the poles and zeros from the above example onto the Z-Plane.

```
% Set up vector for zeros
```

Pole/Zero Plot

Figure 3.7:   Using the zeros and poles found from the transfer function, the one zero is mapped to zero and the two poles are placed at $\frac{1}{2}$ and $-\left(\frac{3}{4}\right)$

Pole/Zero Plot

Figure 3.8:   Using the zeros and poles found from the transfer function, the zeros are mapped to $\pm j$, and the poles are placed at -1, $\frac{1}{2} + \frac{1}{2}j$ and $\frac{1}{2} - \frac{1}{2}j$

```
z = [j ; -j];

% Set up vector for poles
p = [-1 ; .5+.5j ; .5-.5j];

figure(1);
zplane(z,p);
title('Pole/Zero Plot for Complex Pole/Zero Plot Example');
```

3.4.4 Pole/Zero Plot and Region of Convergence

The region of convergence (ROC) for $X(z)$ in the complex Z-plane can be determined from the pole/zero plot. Although several regions of convergence may be possible, where each one corresponds to a different impulse response, there are some choices that are more practical. A ROC can be chosen to make the transfer function causal and/or stable depending on the pole/zero plot.

Filter Properties from ROC

- If the ROC extends outward from the outermost pole, then the system is **causal**.

- If the ROC includes the unit circle, then the system is **stable**.

Below is a pole/zero plot with a possible ROC of the Z-transform in the Simple Pole/Zero Plot (Example 3.5) discussed earlier. The shaded region indicates the ROC chosen for the filter. From this figure, we can see that the filter will be both causal and stable since the above listed conditions are both met.

    Example 3.7:

$$H(z) = \frac{z}{\left(z - \frac{1}{2}\right)\left(z + \frac{3}{4}\right)}$$

3.4.5 Frequency Response and the Z-Plane

The reason it is helpful to understand and create these pole/zero plots is due to their ability to help us easily design a filter. Based on the location of the poles and zeros, the magnitude response of the filter can be quickly understood. Also, by starting with the pole/zero plot, one can design a filter and obtain its transfer function very easily. Refer to this module[6] for information on the relationship between the pole/zero plot and the frequency response.

3.5 Filtering in the Frequency Domain

Because we are interested in actual computations rather than analytic calculations, we must consider the details of the discrete Fourier transform. To compute the length-N DFT, we assume that the signal has a duration less than or equal to N. Because frequency responses

---

[6]http://cnx.org/content/m10548/latest/

Region of Convergence for the Pole/Zero Plot

Figure 3.9: The shaded area represents the chosen ROC for the transfer function.

have an explicit frequency-domain specification[7] in terms of filter coefficients, we don't have a direct handle on which signal has a Fourier transform equaling a given frequency response. Finding this signal is quite easy. First of all, note that the discrete-time Fourier transform of a unit sample equals one for all frequencies. Because of the input and output of linear, shift-invariant systems are related to each other by $Y\left(e^{j2\pi f}\right) = H\left(e^{j2\pi f}\right) X\left(e^{j2\pi f}\right)$, a *unit-sample input, which has $X\left(e^{j2\pi f}\right) = 1$, results in the output's Fourier transform equaling the system's transfer function.*

Exercise 3.1:

This statement is a very important result. Derive it yourself.

In the time-domain, the output for a unit-sample input is known as the system's **unit-sample response**, and is denoted by $h(n)$. Combining the frequency-domain and time-domain interpretations of a linear, shift-invariant system's unit-sample response, we have that $h(n)$ and the transfer function are Fourier transform pairs *in terms of the discrete-time Fourier transform.*

$$\left(h\left(n\right) \leftrightarrow H\left(e^{j2\pi f}\right)\right) \tag{3.33}$$

Returning to the issue of how to use the DFT to perform filtering, we can analytically specify the frequency response, and derive the corresponding length-N DFT by sampling the frequency response.

$$H\left(k\right) = H\left(e^{\frac{j2\pi k}{N}}\right) \quad , \quad k = \{0, \dots, N-1\} \tag{3.34}$$

Computing the inverse DFT yields a length-N signal *no matter what the actual duration of the unit-sample response might be.* If the unit-sample response has a duration less than or equal to N (it's a FIR filter), computing the inverse DFT of the sampled frequency response indeed yields the unit-sample response. If, however, the duration exceeds N, errors are encountered. The nature of these errors is easily explained by appealing to the Sampling

---

[7]http://cnx.org/content/m0510/latest/#dtsinf

Theorem. By sampling in the frequency domain, we have the potential for aliasing in the time domain (sampling in one domain, be it time or frequency, can result in aliasing in the other) unless we sample fast enough. Here, the duration of the unit-sample response determines the minimal sampling rate that prevents aliasing. For FIR systems — they by definition have finite-duration unit sample responses — the number of required DFT samples equals the unit-sample response's duration: $N \geq q$.

Exercise 3.2:

Derive the minimal DFT length for a length-q unit-sample response using the Sampling Theorem. Because sampling in the frequency domain causes repetitions of the unit-sample response in the time domain, sketch the time-domain result for various choices of the DFT length N.

Exercise 3.3:

Express the unit-sample response of a FIR filter in terms of difference equation coefficients. Note that the corresponding question for IIR filters is far more difficult to answer: Consider the example[8].

For IIR systems, we cannot use the DFT to find the system's unit-sample response: aliasing of the unit-sample response will *always* occur. Consequently, we can only implement an IIR filter accurately in the time domain with the system's difference equation. *Frequency-domain implementations are restricted to FIR filters.*

Another issue arises in frequency-domain filtering that is related to time-domain aliasing, this time when we consider the output. Assume we have an input signal having duration N_x that we pass through a FIR filter having a length-$q + 1$ unit-sample response. What is the duration of the output signal? The difference equation for this filter is

$$y(n) = b_0 x(n) + \cdots + b_q x(n - q) \tag{3.35}$$

This equation says that the output depends on current and past input values, with the input value q samples previous defining the extent of the filter's *memory* of past input values. For example, the output at index N_x depends on $x(N_x)$ (which equals zero), $x(N_x - 1)$, through $x(N_x - q)$. Thus, the output returns to zero only after the last input value passes through the filter's memory. As the input signal's last value occurs at index $N_x - 1$, the last nonzero output value occurs when $n - q = N_x - 1$ or $n = q + N_x - 1$. Thus, the output signal's duration equals $q + N_x$.

Exercise 3.4:

In words, we express this result as "The output's duration equals the input's duration plus the filter's duration minus one.". Demonstrate the accuracy of this statement.

The main theme of this result is that a filter's output extends longer than either its input or its unit-sample response. Thus, to avoid aliasing when we use DFTs, the dominant factor is not the duration of input or of the unit-sample response, but of the output. Thus, the number of values at which we must evaluate the frequency response's DFT must be at least $q + N_x$ and we must compute the same length DFT of the input. To accommodate a shorter signal than DFT length, we simply **zero-pad** the input: Ensure that for indices extending beyond the signal's duration that the signal is zero. Frequency-domain filtering, diagrammed in Figure 3.10, is accomplished by storing the filter's frequency response as

Figure 3.10: To filter a signal in the frequency domain, first compute the DFT of the input, multiply the result by the sampled frequency response, and finally compute the inverse DFT of the product. The DFT's length *must* be at least the sum of the input's and unit-sample response's duration minus one. We calculate these discrete Fourier transforms using the fast Fourier transform algorithm, of course.

the DFT $H(k)$, computing the input's DFT $X(k)$, multiplying them to create the output's DFT $Y(k) = H(k)X(k)$, and computing the inverse DFT of the result to yield $y(n)$.

Before detailing this procedure, let's clarify why so many new issues arose in trying to develop a frequency-domain implementation of linear filtering. The frequency-domain relationship[9] between a filter's input and output is *always* true: $Y\left(e^{j2\pi f}\right) = H\left(e^{j2\pi f}\right)X\left(e^{j2\pi f}\right)$. This Fourier transforms in this result are discrete-time Fourier transforms; for example, $X\left(e^{j2\pi f}\right) = \sum_n \left(x(n)e^{-(j2\pi fn)}\right)$. Unfortunately, using this relationship to perform filtering is restricted to the situation when we have analytic formulas for the frequency response and the input signal. The reason why we had to "invent" the discrete Fourier transform (DFT) has the same origin: The spectrum resulting from the discrete-time Fourier transform depends on the *continuous* frequency variable f. That's fine for analytic calculation, but computationally we would have to make an uncountably infinite number of computations.

> NOTE: Did you know that two kinds of infinities can be meaningfully defined? A **countably infinite** quantity means that it can be associated with a limiting process associated with integers. An **uncountably infinite** quantity cannot be so associated. The number of rational numbers is countably infinite (the numerator and denominator correspond to locating the rational by row and column; the total number so-located can be counted, voila!); the number of irrational numbers is uncountably infinite. Guess which is "bigger?"

The DFT computes the Fourier transform at a finite set of frequencies — samples the true spectrum — which can lead to aliasing in the time-domain unless we sample sufficiently fast. The sampling interval here is $\frac{1}{K}$ for a length-K DFT: faster sampling to avoid aliasing thus requires a longer transform calculation. Since the longest signal among the input, unit-sample response and output is the output, it is that signal's duration that determines the transform length. We simply extend the other two signals with zeros (zero-pad) to compute their DFTs.

Example 3.8:

Suppose we want to average daily stock prices taken over last year to yield a running weekly average (average over five trading sessions). The filter we want is a length-5 averager (as shown in the unit-sample response[10]), and the input's duration is 253 (365 calendar days minus weekend days and holidays). The output

---

[8]http://cnx.org/content/m10251/latest/#p0
[9]http://cnx.org/content/m0510/latest/#equation4
[10]http://cnx.org/content/m10251/latest/#fig1002

Figure 3.11: The blue line shows the Dow Jones Industrial Average from 1997, and the red one the length-5 boxcar-filtered result that provides a running weekly of this market index. Note the "edge" effects in the filtered output.

duration will be $253 + 5 - 1 = 257$, and this determines the transform length we need to use. Because we want to use the FFT, we are restricted to power-of-two transform lengths. We need to choose any FFT length that exceeds the required DFT length. As it turns out, 256 is a power of two ($2^8 = 256$), and this length just undershoots our required length. To use frequency domain techniques, we must use length-512 fast Fourier transforms.

Figure 3.11 shows the input and the filtered output. The MATLAB programs that compute the filtered output in the time and frequency domains are

```
Time Domain
h = [1 1 1 1 1]/5;
y = filter(h,1,[djia zeros(1,4)]);

Frequency Domain
h = [1 1 1 1 1]/5;
DJIA = fft(djia, 512);
H = fft(h, 512);
Y = H.*X;
y = ifft(Y);
```

> NOTE: The `filter` program has the feature that the length of its output equals the length of its input. To force it to produce a signal having the proper length, the program zero-pads the input appropriately.

MATLAB's `fft` function automatically zero-pads its input if the specified transform length (its second argument) exceeds the signal's length. The frequency domain result will have a small imaginary component — largest value is 2.2×10^{-11} —

because of the inherent finite precision nature of computer arithmetic. Because of the unfortunate misfit between signal lengths and favored FFT lengths, the number of arithmetic operations in the time-domain implementation is far less than those required by the frequency domain version: 514 versus 62,271. If the input signal had been one sample shorter, the frequency-domain computations would have been more than a factor of two less (28,696), but far more than in the time-domain implementation.

An interesting signal processing aspect of this example is demonstrated at the beginning and end of the output. The ramping up and down that occurs can be traced to assuming the input is zero before it begins and after it ends. The filter "sees" these initial and final values as the difference equation passes over the input. These artifacts can be handled in two ways: we can just ignore the edge effects or the data from previous and succeeding years' last and first week, respectively, can be placed at the ends.

3.6 Linear-Phase FIR Filters

3.6.1 THE AMPLITUDE RESPONSE

If the real and imaginary parts of $H^f(\omega)$ are given by

$$H^f(\omega) = \text{Re}(\omega) + j\text{Im}(\omega) \tag{3.36}$$

the magnitude and phase are defined as

$$|H^f(\omega)| = \sqrt{(\text{Re}(\omega))^2 + (\text{Im}(\omega))^2}$$

$$p(\omega) = \arctan\left(\frac{\text{Im}(\omega)}{\text{Re}(\omega)}\right)$$

so that

$$H^f(\omega) = |H^f(\omega)|e^{jp(\omega)} \tag{3.37}$$

With this definition, $|H^f(\omega)|$ is never negative and $p(\omega)$ is usually discontinuous, but it can be very helpful to write $H^f(\omega)$ as

$$H^f(\omega) = A(\omega)e^{j\theta(\omega)} \tag{3.38}$$

where $A(\omega)$ can be positive and negative, and $\theta(\omega)$ continuous. $A(\omega)$ is called the **amplitude response**. Figure 3.12 illustrates the difference between $|H^f(\omega)|$ and $A(\omega)$.

A linear-phase phase filter is one for which the continuous phase $\theta(\omega)$ is linear.

$$H^f(\omega) = A(\omega)e^{j\theta(\omega)}$$

with

$$\theta(\omega) = (-M)\omega + B$$

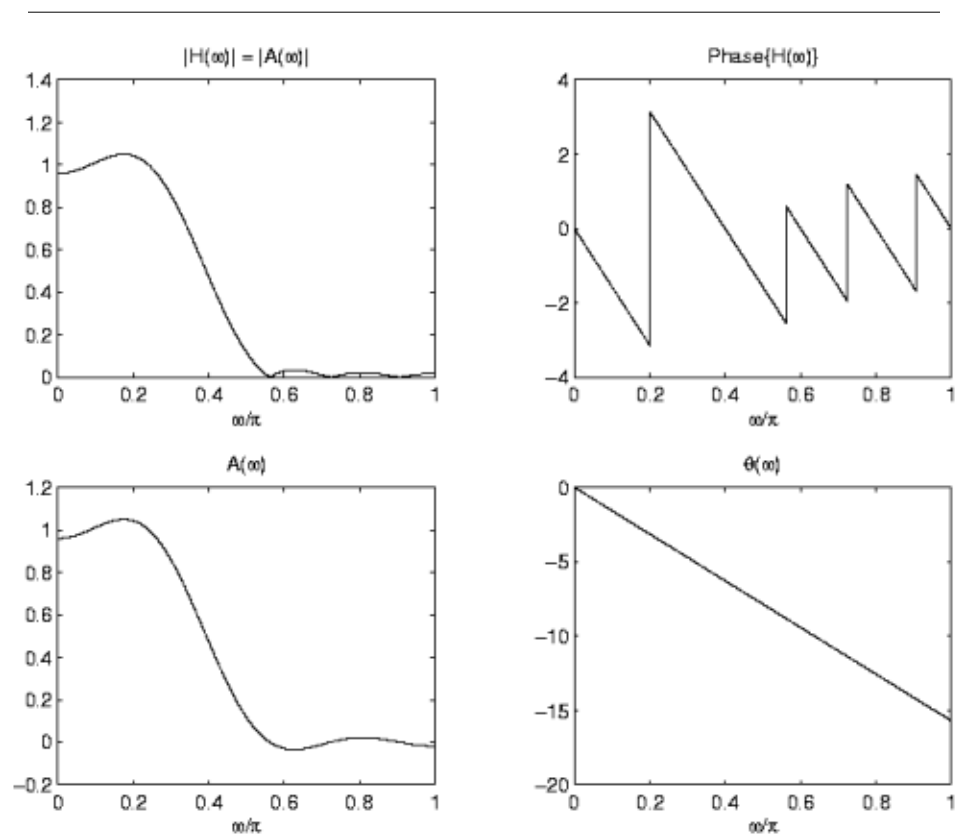We assume in the following that the impulse response $h(n)$ is real-valued.

Figure 3.12

3.6.2 WHY LINEAR-PHASE?

If a discrete-time cosine signal

$$x_1(n) = \cos(\omega_1 n + \phi_1)$$

is processed through a discrete-time filter with frequency response

$$H^f(\omega) = A(\omega)e^{j\theta(\omega)}$$

then the output signal is given by

$$y_1(n) = A(\omega_1)\cos(\omega_1 n + \phi_1 + \theta(\omega_1))$$

or

$$y_1(n) = A(\omega_1)\cos\left(\omega_1\left(n + \frac{\theta(\omega_1)}{\omega_1}\right) + \phi_1\right)$$

The LTI system has the effect of scaling the cosine signal and delaying it by $\frac{\theta(\omega_1)}{\omega_1}$.

> **Exercise 3.5:**
>
> When does the system delay cosine signals with different frequencies by the same amount?

The function $\frac{\theta(\omega)}{\omega}$ is called the **phase delay**. A linear phase filter therefore has constant phase delay.

3.6.3 WHY LINEAR-PHASE: EXAMPLE

Consider a discrete-time filter described by the difference equation

$$y(n) = -0.1821x(n) + 0.7865x(n-1) - 0.6804x(n-2) + x(n-3) + 0.6804y(n-1) - 0.7865y(n-2) + 0.1821y(n-3) \tag{3.39}$$

When $\omega_1 = 0.31\pi$, then the delay is $\frac{-\theta(\omega_1)}{\omega_1} = 2.45$. The delay is illustrated in Figure 3.13:

Notice that the delay is fractional — the discrete-time samples are not exactly reproduced in the output. The fractional delay can be interpreted in this case as a delay of the underlying continuous-time cosine signal.

3.6.4 WHY LINEAR-PHASE: EXAMPLE (2)

Consider the same system given on the previous slide, but let us change the frequency of the cosine signal.

When $\omega_2 = 0.47\pi$, then the delay is $\frac{-\theta(\omega_2)}{\omega_2} = 0.14$.

NOTE: For this example, the delay depends on the frequency, because this system does not have linear phase.

3.6.5 WHY LINEAR-PHASE: MORE

From the previous slides, we see that a filter will delay different frequency components of a signal by the same amount if the filter has linear phase (constant phase delay).

In addition, when a narrow band signal (as in AM modulation) goes through a filter, the envelop will be delayed by the **group delay** or **envelop delay** of the filter. The amount by which the envelop is delayed is independent of the carrier frequency only if the filter has linear phase.

Also, in applications like image processing, filters with non-linear phase can introduce artifacts that are visually annoying.

Figure 3.13

3.7 Filter Structures

A realizable filter must require only a finite number of computations per output sample. For linear, causal, time-Invariant filters, this restricts one to rational transfer functions of the form

$$H\left(z\right) = \frac{b_0 + b_1 z^{-1} + \cdots + b_m z^{-m}}{1 + a_1 z^{-1} + a_2 z^{-2} + \cdots + a_n z^{-n}}$$

Assuming no pole-zero cancellations, $H\left(z\right)$ is FIR if $a_i = 0$, $i > 0$, and IIR otherwise. Filter structures usually implement rational transfer functions as difference equations.

Whether FIR or IIR, a given transfer function can be implemented with many different filter structures. With infinite-precision data, coefficients, and arithmetic, all filter structures implementing the same transfer function produce the same output. However, different filter strucures may produce very different errors with quantized data and finite-precision or fixed-point arithmetic. The computational expense and memory usage may also differ greatly. Knowledge of different filter structures allows DSP engineers to trade off these factors to create the best implementation.

3.8 Overview of Digital Filter Design

Advantages of FIR filters

1. Straight forward conceptually and simple to implement

Figure 3.14

2. Can be implemented with fast convolution

3. Always stable

4. Relatively insensitive to quantization

5. Can have linear phase (same time delay of all frequencies)

Advantages of IIR filters

1. Better for approximating analog systems

2. For a given magnitude response specification, IIR filters often require much less computation than an equivalent FIR, particularly for narrow transition bands

Both FIR and IIR filters are very important in applications.

Generic Filter Design Procedure

1. Choose a desired response, based on application requirements

2. Choose a filter class

3. Choose a quality measure

4. Solve for the filter in class 2 optimizing criterion in 3

3.8.1 Perspective on FIR filtering

Most of the time, people do L^∞ optimal design, using the Parks-McClellan algorithm (Section 4.11). This is probably the second most important technique in "classical" signal processing (after the Cooley-Tukey (radix-2^{11}) FFT).

Most of the time, FIR filters are designed to have linear phase. The most important advantage of FIR filters over IIR filters is that they can have exactly linear phase. There are advanced design techniques for minimum-phase filters, constrained L^2 optimal designs, *etc.* (see chapter 8 of text). However, if only the *magnitude* of the response is important, IIR filers usually require much fewer operations and are typically used, so the bulk of FIR filter design work has concentrated on linear phase designs.

3.9 Window Design Method

The truncate-and-delay design procedure is the simplest and most obvious FIR design procedure.

Exercise 3.6:

Is it any Good?

3.9.1 L2 optimization criterion

find $h[n]$, $0 \le n \le M-1$, maximizing the energy difference between the desired response and the actual response: i.e., find

$$\min_{h[n]} \left\{ \int_{-\pi}^{\pi} (|H_d(\omega) - H(\omega)|)^2 d\omega \right\}$$

by Parseval's relationship[12]

$$
\begin{aligned}
\min_{h[n]} \left\{ \int_{-\pi}^{\pi} (|H_d(\omega) - H(\omega)|)^2 d\omega \right\} &= 2\pi \sum_{n=-\infty}^{\infty} \left((|h_d[n] - h[n]|)^2 \right) \\
&= 2\pi \left(\sum_{n=-\infty}^{-1} \left((|h_d[n] - h[n]|)^2 \right) + \sum_{n=0}^{M-1} \left((|h_d[n] - h[n]|)^2 \right) + \sum_{n=M}^{\infty} \right.
\end{aligned}
$$
(3.40)

Since $= h[n]$, $n < 0 n \ge M$ this becomes

$$\min_{h[n]} \left\{ \int_{-\pi}^{\pi} (|H_d(\omega) - H(\omega)|)^2 d\omega \right\} = \sum_{h=-\infty}^{-1} \left((|h_d[n]|)^2 \right) + \sum_{n=0}^{M-1} \left((|h[n] - h_d[n]|)^2 \right) + \sum_{n=M}^{\infty} \left((|h_d[n]|)^2 \right)$$

NOTE: $h[n]$ has no influence on the first and last sums.

The best we can do is let

$$h[n] = \begin{cases} h_d[n] & \text{if } 0 \le n \le M-1 \\ 0 & \text{if else} \end{cases}$$

Thus $h[n] = h_d[n] w[n]$,

$$w[n] = \begin{cases} 1 & \text{if } 0 \le n(M-1) \\ 0 & \text{if else} \end{cases}$$

is *optimal* in a least-total-sqaured-error (L_2, or energy) sense!

---

[11] http://cnx.org/content/m12016/latest/
[12] http://cnx.org/content/m0047/latest/

Exercise 3.7:

Why, then, is this design often considered undersirable?

For desired spectra with discontinuities, the least-square designs are poor in a minimax (worst-case, or L_∞) error sense.

3.9.2 Window Design Method

Apply a more gradual truncation to reduce "ringing" (Gibb's Phenomenon[13])

,

NOTE: $H(\omega) = H_d(\omega) * W(\omega)$

The window design procedure (except for the boxcar window) is ad-hoc and not optimal in any usual sense. However, it is very simple, so it is sometimes used for "quick-and-dirty" designs of if the error criterion is itself heurisitic.

3.10 Frequency Sampling Design Method for FIR filters

Given a desired frequency response, the frequency sampling design method designs a filter with a frequency response *exactly* equal to the desired response at a particular set of frequencies ω_k.

Procedure

$$H_d(\omega_k) = \sum_{n=0}^{M-1} \left(h(n) e^{-(j\omega_k n)} \right) \quad , \quad k = [o, 1, \ldots, N-1] \tag{3.41}$$

NOTE: Desired Response must incluce linear phase shift (if linear phase is desired)

Exercise 3.8:

What is $H_d(\omega)$ for an ideal lowpass filter, cotoff at ω_c?

NOTE: This set of linear equations can be written in matrix form

$$H_d(\omega_k) = \sum_{n=0}^{M-1} \left(h(n) e^{-(j\omega_k n)} \right) \tag{3.42}$$

$$\begin{pmatrix} H_d(\omega_0) \\ H_d(\omega_1) \\ \vdots \\ H_d(\omega_{N-1}) \end{pmatrix} = \begin{pmatrix} e^{-(j\omega_0 0)} & e^{-(j\omega_0 1)} & \ldots & e^{-(j\omega_0(M-1))} \\ e^{-(j\omega_1 0)} & e^{-(j\omega_1 1)} & \ldots & e^{-(j\omega_1(M-1))} \\ \vdots & \vdots & \vdots & \vdots \\ e^{-(j\omega_{M-1} 0)} & e^{-(j\omega_{M-1} 1)} & \ldots & e^{-(j\omega_{M-1}(M-1))} \end{pmatrix} \begin{pmatrix} h(0) \\ h(1) \\ \vdots \\ h(M-1) \end{pmatrix}$$
$$\tag{3.43}$$

or

$$H_d = W\mathbf{h}$$

So

$$\mathbf{h} = W^{-1} H_d \tag{3.44}$$

NOTE: W is a square matrix for $N = M$, and invertible as long as $\omega_i \neq \omega_j + 2\pi l$, $i \neq j$

---

[13]http://cnx.org/content/m10092/latest/

3.10.1 Important Special Case

What if the frequencies are equally spaced between 0 and 2π, i.e. $\omega_k = \frac{2\pi k}{M} + \alpha$

Then

$$H_d\left(\omega_k\right) = \sum_{n=0}^{M-1}\left(h\left(n\right)e^{-\left(j\frac{2\pi kn}{M}\right)}e^{-\left(j\alpha n\right)}\right) = \sum_{n=0}^{M-1}\left(\left(h\left(n\right)e^{-\left(j\alpha n\right)}\right)e^{-\left(j\frac{2\pi kn}{M}\right)}\right) = \text{DFT!}$$

so

$$h\left(n\right)e^{-\left(j\alpha n\right)} = \frac{1}{M}\sum_{k=0}^{M-1}\left(H_d\left(\omega_k\right)e^{+j\frac{2\pi nk}{M}}\right)$$

or

$$h\left[n\right] = \frac{e^{j\alpha n}}{M}\sum_{k=0}^{M-1}\left(H_d\left[\omega_k\right]e^{j\frac{2\pi nk}{M}}\right) = e^{j\alpha n}\text{IDFT}\left[H_d\left[\omega_k\right]\right]$$

3.10.2 Important Special Case 2

$h\left[n\right]$ symmetric, linear phase, and has real coefficients. Since $h\left[n\right] = h\left[-1\right]$, there are only $\frac{M}{2}$ degrees of freedom, and only $\frac{M}{2}$ linear equations are required.

$$
\begin{aligned}
H\left[\omega_k\right] &= \sum_{n=0}^{M-1}\left(h\left[n\right]e^{-\left(j\omega_k n\right)}\right) \\
&= \begin{cases} \sum_{n=0}^{\frac{M}{2}-1}\left(h\left[n\right]\left(e^{-\left(j\omega_k n\right)}+e^{-\left(j\omega_k\left(M-n-1\right)\right)}\right)\right) & \text{if M even} \\ \sum_{n=0}^{M-\frac{3}{2}}\left(+h\left[n\right]\left(e^{-\left(j\omega_k n\right)}+e^{-\left(j\omega_k\left(M-n-1\right)\right)}\right)\left(h\left[\frac{M-1}{2}\right]e^{-\left(j\omega_k\frac{M-1}{2}\right)}\right)\right) & \text{if M odd} \end{cases} \\
&= \begin{cases} e^{-\left(j\omega_k\frac{M-1}{2}\right)}2\sum_{n=0}^{\frac{M}{2}-1}\left(h\left[n\right]\cos\left(\omega_k\left(\frac{M-1}{2}-n\right)\right)\right) & \text{if M even} \\ e^{-\left(j\omega_k\frac{M-1}{2}\right)}2\sum_{n=0}^{M-\frac{3}{2}}\left(h\left[n\right]\cos\left(\omega_k\left(\frac{M-1}{2}-n\right)\right)+h\left[\frac{M-1}{2}\right]\right) & \text{if M odd} \end{cases}
\end{aligned}
$$
(3.45)

Removing linear phase from both sides yields

$$A\left(\omega_k\right) = \begin{cases} 2\sum_{n=0}^{\frac{M}{2}-1}\left(h\left[n\right]\cos\left(\omega_k\left(\frac{M-1}{2}-n\right)\right)\right) & \text{if M even} \\ 2\sum_{n=0}^{M-\frac{3}{2}}\left(h\left[n\right]\cos\left(\omega_k\left(\frac{M-1}{2}-n\right)\right)+h\left[\frac{M-1}{2}\right]\right) & \text{if M odd} \end{cases}$$

Due to symmetry of response for real coefficients, only $\frac{M}{2}$ ω_k on $\omega \in [0, \pi)$ need be specified, with the frequencies $-\omega_k$ thereby being implicitly defined also. Thus we have $\frac{M}{2}$ *real-valued* simultaneous linear equations to solve for $h\left[n\right]$.

3.10.2.1 Special Case 2a

$h\left[n\right]$ symmetric, odd length, linear phase, real coefficients, and ω_k equally spaced: $\omega_k = \frac{n\pi k}{M}$, $0 \le k \le M-1$

$$
\begin{aligned}
h\left[n\right] &= \text{IDFT}\left[H_d\left(\omega_k\right)\right] \\
&= \frac{1}{M}\sum_{k=0}^{M-1}\left(A\left(\omega_k\right)e^{-\left(j\frac{2\pi k}{M}\right)\frac{M-1}{2}}e^{j\frac{2\pi nk}{M}}\right) \\
&= \frac{1}{M}\sum_{k=0}^{M-1}\left(A\left(k\right)e^{j\left(\frac{2\pi k}{M}\left(n-\frac{M-1}{2}\right)\right)}\right)
\end{aligned}
$$
(3.46)

To yield real coefficients, $A\left(\omega\right)$ mus be symmetric

$$A\left(\omega\right) = A\left(-\omega\right) \Rightarrow A\left[k\right] = A\left[M-k\right]$$

$$
\begin{aligned}
h\left[n\right] &= \frac{1}{M}\left(A\left(0\right)+\sum_{k=1}^{\frac{M-1}{2}}\left(A\left[k\right]\left(e^{j\frac{2\pi k}{M}\left(n-\frac{M-1}{2}\right)}+e^{-\left(j2\pi k\left(n-\frac{M-1}{2}\right)\right)}\right)\right)\right)\\
&= \frac{1}{M}\left(A\left(0\right)+2\sum_{k=1}^{\frac{M-1}{2}}\left(A\left[k\right]\cos\left(\frac{2\pi k}{M}\left(n-\frac{M-1}{2}\right)\right)\right)\right)\\
&= \frac{1}{M}\left(A\left(0\right)+2\sum_{k=1}^{\frac{M-1}{2}}\left(A\left[k\right]\left(-1\right)^{k}\cos\left(\frac{2\pi k}{M}\left(n+\frac{1}{2}\right)\right)\right)\right)
\end{aligned}
\tag{3.47}
$$

Simlar equations exist for even lengths, anti-symmetric, and $\alpha=\frac{1}{2}$ filter forms.

3.10.3 Comments on frequency-sampled design

This method is simple conceptually and very efficient for equally spaced samples, since $h\left[n\right]$ can be computed using the IDFT.

$H\left(\omega\right)$ for a frequency sampled design goes *exactly* through the sample points, but it may be very far off from the desired response for $\omega\neq\omega_k$. This is the main problem with frequency sampled design.

Possible solution to this problem: specify more frequency samples than degrees of freedom, and minimize the total error in the frequency response at all of these samples.

3.10.4 Extended frequency sample design

For the samples $H\left(\omega_k\right)$ where $0\leq k\leq M-1$ and $N>M$, find $h\left[n\right]$, where $0\leq n\leq M-1$ minimizing $\parallel H_d\left(\omega_k\right)-H\left(\omega_k\right)\parallel$

For $\parallel l\parallel_{\infty}$ norm, this becomes a linear programming problem (standard packages availble!)

Here we will consider the $\parallel l\parallel_{2}$ norm.

To minimize the $\parallel l\parallel_{2}$ norm; that is, $\sum_{n=0}^{N-1}\left(|H_d\left(\omega_k\right)-H\left(\omega_k\right)|\right)$, we have an overdetermined set of linear equations:

$$
\begin{pmatrix}
e^{-\left(j\omega_0 0\right)} & \cdots & e^{-\left(j\omega_0\left(M-1\right)\right)}\\
\vdots & \vdots & \vdots\\
e^{-\left(j\omega_{N-1}0\right)} & \cdots & e^{-\left(j\omega_{N-1}\left(M-1\right)\right)}
\end{pmatrix}\mathbf{h}=
\begin{pmatrix}
H_d\left(\omega_0\right)\\
H_d\left(\omega_1\right)\\
\vdots\\
H_d\left(\omega_{N-1}\right)
\end{pmatrix}
$$

or

$$
W\mathbf{h}=H_d
$$

The minimum error norm solution is well known to be $\mathbf{h}=\left(W^*W\right)^{-1}W^*H_d$; $\left(W^*W\right)^{-1}W^*$ is well known as the pseudo-inverse matrix.

> NOTE: Extended frequency sampled design discourages radical behavior of the frequency response between samples for sufficiently closely spaced samples. However, the actual frequency response may no longer pass exactly through *any* of the $H_d\left(\omega_k\right)$.

3.11 Parks-McClellan FIR Filter Design

Very often, the approximation tolerances for a filter are given in terms of the maximum deviation within bands. For example, we might wish a lowpass filter in a (16-bit) CD player to have no more than $\frac{1}{2}$-bit deviation in the pass and stop bands.

$$
H\left(\omega\right)=\begin{cases}
1-\frac{1}{2^{17}}\leq|H\left(\omega\right)|\leq1+\frac{1}{2^{17}}\text{ if }|\omega|\leq\omega_p\\
\frac{1}{2^{17}}\geq|H\left(\omega\right)|\text{ if }\omega_p\leq|\omega|\leq\pi
\end{cases}
$$

Figure 3.15: The black boxes on the left and right are the passbands, the black boxes in the middle represent the stop band, and the space between the boxes are the transition bands. Note that overshoots may be allowed in the transition bands.

Typically, we would like to have the shortest-length filter achieving these specifications. See Figure 3.15.

Exercise 3.9:

Must there be a transition band?

3.11.1 Formal Statement of the L-∞ (Minimax) Design Problem

For a given filter length (M) and type (odd length, symmetric, linear phase, for example), and a relative error weighting function $W(\omega)$, find the filter coefficients minimizing the maximum error

$$\underset{\mathbf{h}}{argmin}\,\underset{\omega \in F}{argmax}|E(\omega)| = \underset{\mathbf{h}}{argmin}\| E(\omega) \|_{\infty}$$

where

$$E(\omega) = W(\omega)(H_d(\omega) - H(\omega))$$

and F is a compact subset of $\omega \in [0, \pi]$ (*i.e.*, all ω in the passbands and stop bands).

NOTE: Typically, we would often rather specify $\parallel E(\omega) \parallel_\infty \leq \delta$ and minimize over M and $\mathbf{h}$; however, the design techniques minimize δ for a given M. One then repeats the design procedure for different M until the minimum M satisfying the requirements is found.

We will discuss in detail the design only of odd-length symmetric linear-phase FIR filters. Even-length and anti-symmetric linear phase FIR filters are essentially the same except for a slightly different implicit weighting function. For arbitrary phase, exactly optimal design procedures have only recently been developed (1990).

3.11.2 Outline of L-∞ Filter Design

1. Using results from Approximation Theory, simple conditions for determining whether a given filter is L^∞ optimal are found.

2. Some ad-hoc iterative methods for finding a filter which satisfies these condition (and which is thus optimal) are developed.

   That is, the L^∞ filter design problem is actually solved *indirectly*.

3.11.3 Conditions for L-∞ Optimality of a Linear-phase FIR Filter

All conditions are based on Chebyshev's "Alternation Theorem," a mathematical fact from polynomial approximation theory.

3.11.3.1 Alternation Theorem

Let F be a compact subset on the real axis x, and let $P(x)$ be and Lth-order polynomial

$$P(x) = \sum_{k=0}^{L} \left(a_k x^k\right)$$

Also, let $D(x)$ be a desired function of x that is continuous on F, and $W(x)$ a positive, continuous weighting function on F. Define the error $E(x)$ on F as

$$E(x) = W(x)(D(x) - P(x))$$

and

$$\parallel E(x) \parallel_\infty = \underset{x \in F}{argmax} |E(x)|$$

A necessary and sufficient condition that $P(x)$ is the unique Lth-order polynomial minimizing $\parallel E(x) \parallel_\infty$ is that $E(x)$ exhibits *at least* $L+2$ "alternations;" that is, there must exist at least $L+2$ values of x, $x_k \in F$, $k = [0, 1, \ldots, L+1]$, such that $x_0 < x_1 < \cdots < x_{L+2}$ and such that $E(x_k) = -(E(x_{k+1})) = \pm(\parallel E \parallel_\infty)$

Exercise 3.10:

What does this have to do with linear-phase filter design?

3.11.4 Optimality Conditions for Even-length Symmetric Linear-phase Filters

For M even,

$$A(\omega) = \sum_{n=0}^{L} \left(h(L-n) \cos\left(\omega\left(n + \frac{1}{2} \right) \right) \right)$$

where $L = \frac{M}{2} - 1$ Using the trigonometric identity $\cos(\alpha + \beta) = \cos(\alpha - \beta) + 2\cos(\alpha)\cos(\beta)$ to pull out the $\frac{\omega}{2}$ term and then using the other trig identities (pg 134), it can be shown that $A(\omega)$ can be written as

$$A(\omega) = \cos\left(\frac{\omega}{2} \right) \sum_{k=0}^{L} \left(\alpha_k \cos^k(\omega) \right)$$

Again, this is a polynomial in $x = \cos(\omega)$, except for a weighting function out in front.

$$
\begin{aligned}
E(\omega) &= W(\omega)(A_d(\omega) - A(\omega)) \\
&= W(\omega)\left(A_d(\omega) - \cos\left(\frac{\omega}{2} \right) P(\omega) \right) \\
&= W(\omega)\cos\left(\frac{\omega}{2} \right)\left(\frac{A_d(\omega)}{\cos\left(\frac{\omega}{2} \right)} - P(\omega) \right)
\end{aligned}
\tag{3.48}
$$

which implies

$$E(x) = W'(x)\left(A_d'(x) - P(x) \right) \tag{3.49}$$

where

$$W'(x) = W\left((\cos(x))^{-1} \right) \cos\left(\frac{1}{2}(\cos(x))^{-1} \right)$$

and

$$A_d'(x) = \frac{A_d\left((\cos(x))^{-1} \right)}{\cos\left(\frac{1}{2}(\cos(x))^{-1} \right)}$$

Again, this is a polynomial approximation problem, so the alternation theorem holds. If $E(\omega)$ has at least $L + 2 = \frac{M}{2} + 1$ alternations, the even-length symmetric filter is optimal in an L^∞ sense.

The prototypical filter design problem:

$$W = \left\{ \begin{array}{l} 1 \text{ if } |\omega| \le \omega_p \\ \frac{\delta_s}{\delta_p} \text{ if } |\omega_s| \le |\omega| \end{array} \right.$$

See Figure 3.16.

3.11.5 L-∞ Optimal Lowpass Filter Design Lemma

1. The maximum possible number of alternations for a lowpass filter is $L + 3$: The proof is that the extrema of a polynomial occur only where the derivative is zero: $\frac{\partial}{\partial x} P(x) = 0$. Since $P'(x)$ is an $(L-1)$th-order polynomial, it can have at most $L - 1$ zeros. However, the mapping $x = \cos(\omega)$ implies that $\frac{\partial}{\partial \omega} A(\omega) = 0$ at $\omega = 0$ and $\omega = \pi$, for two more possible alternation points. Finally, the band edges can also be alternations, for a total of $L - 1 + 2 + 2 = L + 3$ possible alternations.

2. There must be an alternation at either $\omega = 0$ or $\omega = \pi$.

3. Alternations must occur at ω_p and ω_s. See Figure 3.16.

Figure 3.16

4. The filter must be equiripple except at possibly $\omega = 0$ or $\omega = \pi$. Again see Figure 3.16.

> NOTE: The alternation theorem doesn't directly suggest a method for computing the optimal filter. It simply tells us how to recognize that a filter *is* optimal, or *isn't* optimal. What we need is an intelligent way of guessing the optimal filter coefficients.

In matrix form, these $L + 2$ simultaneous equations become

$$\begin{pmatrix} 1 & \cos(\omega_0) & \cos(2\omega_0) & ... & \cos(L\omega_0) & \frac{1}{W(\omega_0)} \\ 1 & \cos(\omega_1) & \cos(2\omega_1) & ... & \cos(L\omega_1) & \frac{-1}{W(\omega_1)} \\ \vdots & \vdots & \ddots & ... & \vdots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \vdots & \vdots & \vdots & ... & \ddots & \vdots \\ 1 & \cos(\omega_{L+1}) & \cos(2\omega_{L+1}) & ... & \cos(L\omega_{L+1}) & \frac{\pm 1}{W(\omega_{L+1})} \end{pmatrix} \begin{pmatrix} h(L) \\ h(L-1) \\ \vdots \\ h(1) \\ h(0) \\ \delta \end{pmatrix} = \begin{pmatrix} A_d(\omega_0) \\ A_d(\omega_1) \\ \vdots \\ \vdots \\ \vdots \\ A_d(\omega_{L+1}) \end{pmatrix}$$

or

$$W \begin{pmatrix} \mathbf{h} \\ \delta \end{pmatrix} = A_d$$

So, for the given set of $L + 2$ extremal frequencies, we can solve for $\mathbf{h}$ and δ via $(\mathbf{h}, \delta)^T = W^{-1} A_d$. Using the FFT, we can compute $A(\omega)$ of $h(n)$, on a dense set of frequencies. If the old ω_k are, in fact the extremal locations of $A(\omega)$, then the alternation theorem is satisfied and $h(n)$ is **optimal**. If not, repeat the process with the new extremal locations.

3.11.6 Computational Cost

$O(L^3)$ for the matrix inverse and $N\log_2 N$ for the FFT ($N \geq 32L$, typically), *per iteration*!
  This method is expensive computationally due to the matrix inverse.
  A more efficient variation of this method was developed by Parks and McClellan (1972), and is based on the Remez exchange algorithm. To understand the Remez exchange algorithm, we first need to understand Lagrange Interpoloation[14].
  Now $A(\omega)$ is an Lth-order polynomial in $x = \cos(\omega)$, so Lagrange interpolation can be used to *exactly* compute $A(\omega)$ from $L + 1$ samples of $A(\omega_k)$, $k = [0, 1, 2, ..., L]$.
  Thus, given a set of extremal frequencies and knowing δ, samples of the amplitude response $A(\omega)$ can be computed *directly* from the

$$A(\omega_k) = \frac{(-1)^{k(+1)}}{W(\omega_k)}\delta + A_d(\omega_k) \tag{3.50}$$

without solving for the filter coefficients!
  This leads to computational savings!
  Note that Equation 3.50 is a set of $L + 2$ simultaneous equations, which can be solved for δ to obtain (Rabiner, 1975)

$$\delta = \frac{\sum_{k=0}^{L+1} (\gamma_k A_d(\omega_k))}{\sum_{k=0}^{L+1} \left(\frac{(-1)^{k(+1)}\gamma_k}{W(\omega_k)} \right)} \tag{3.51}$$

---

[14]http://cnx.org/content/m1/latest/

where
$$\gamma_k = \prod_{\substack{i=0 \\ i \neq k}}^{L+1} \left(\frac{1}{\cos\left(\omega_k\right) - \cos\left(\omega_i\right)} \right)$$

The result is the Parks-McClellan FIR filter design method, which is simply an application of the Remez exchange algorithm to the filter design problem. See Figure 3.17.

The cost per iteration is $O\left(16L^2\right)$, as opposed to $O\left(L^3\right)$; much more efficient for large L. Can also interpolate to DFT sample frequencies, take inverse FFT to get corresponding filter coefficients, and zeropad and take longer FFT to efficiently interpolate.

3.12 FIR Filter Design using MATLAB

3.12.1 FIR Filter Design Using MATLAB

3.12.1.1 Design by windowing

The MATLAB function `fir1()` designs conventional lowpass, highpass, bandpass, and bandstop linear-phase FIR filters based on the windowing method. The command

```
b = fir1(N,Wn)
```

returns in vector `b` the impulse response of a lowpass filter of order `N`. The cut-off frequency `Wn` must be between 0 and 1 with 1 corresponding to the half sampling rate.

The command

```
b = fir1(N,Wn,'high')
```

returns the impulse response of a highpass filter of order `N` with normalized cutoff frequency `Wn`.

Similarly, `b = fir1(N,Wn,'stop')` with `Wn` a two-element vector designating the stopband designs a bandstop filter.

Without explicit specification, the *Hamming window* is employed in the design. Other windowing functions can be used by specifying the windowing function as an extra argument of the function. For example, Blackman window can be used instead by the command `b = fir1(N, Wn, blackman(N))`.

3.12.1.2 Parks-McClellan FIR filter design

The MATLAB command

```
b = remez(N,F,A)
```
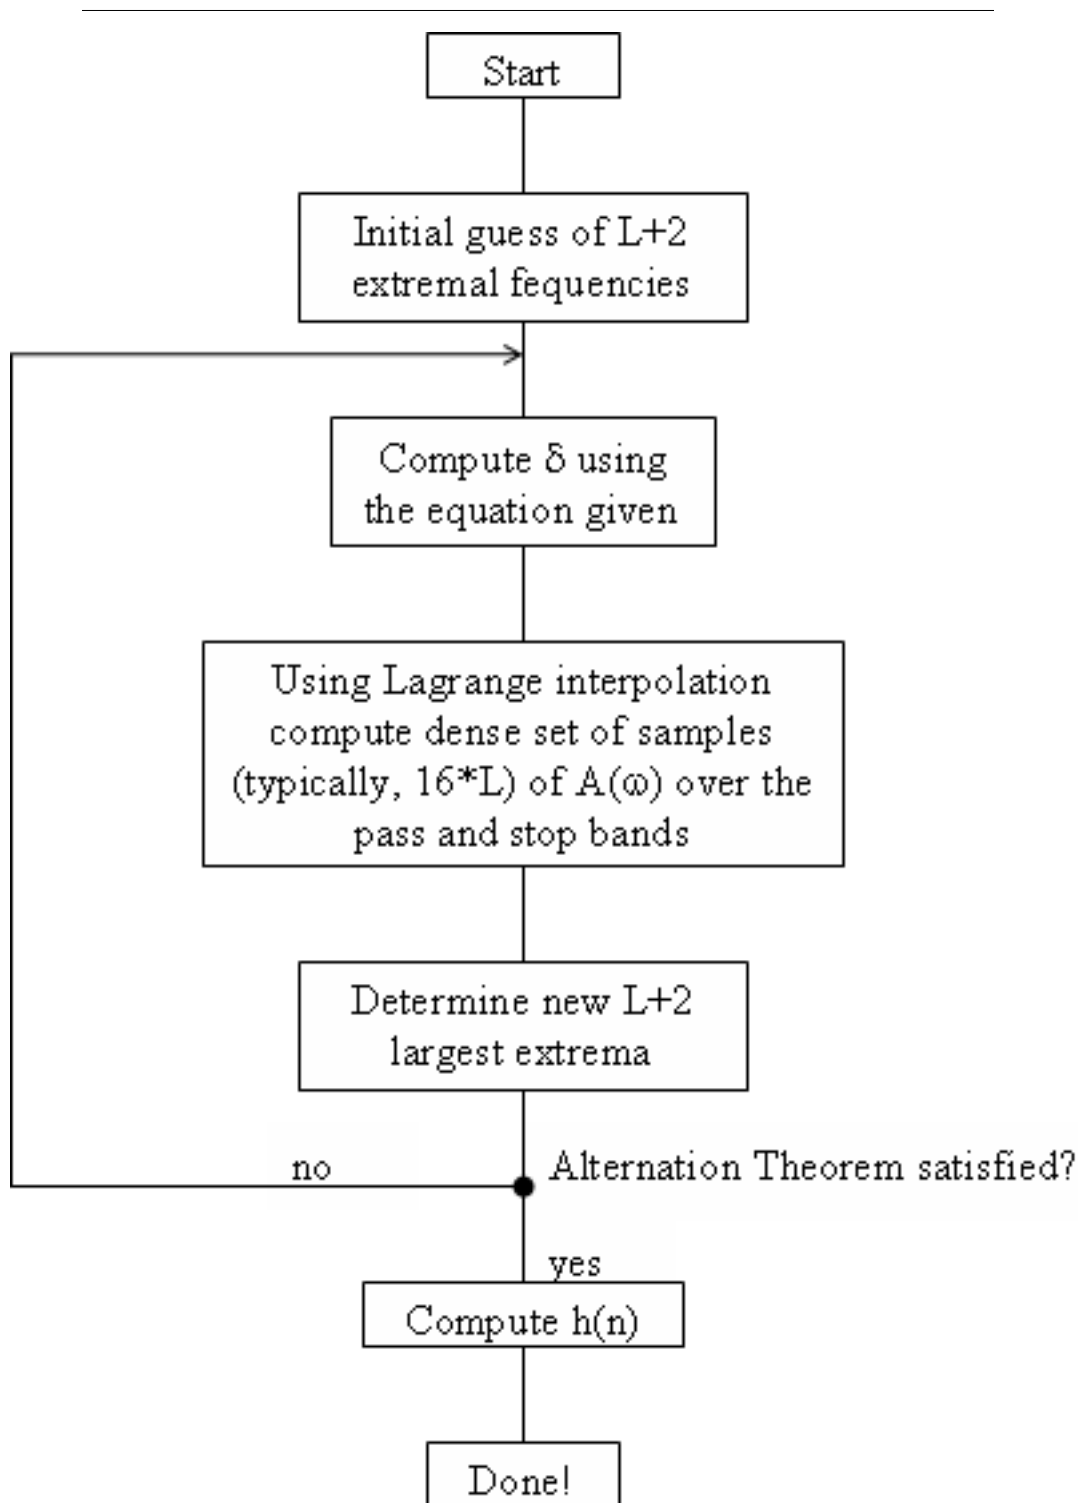
Figure 3.17: The initial guess of extremal frequencies is usually equally spaced in the band. Computing δ costs $O\left(L^2\right)$. Using Lagrange interpolation costs $O\left(16LL\right) \approx O\left(16L^2\right)$. Computing $h\left(n\right)$ costs $O\left(L^3\right)$, but it is only done once!

returns the impulse response of the length `N+1` linear phase FIR filter of order `N` designed by Parks-McClellan algorithm. `F` is a vector of frequency band edges in ascending order between 0 and 1 with 1 corresponding to the half sampling rate. `A` is a real vector of the same size as `F` which specifies the desired amplitude of the frequency response of the points `(F(k),A(k))` and `(F(k+1),A(k+1))` for odd `k`. For odd `k`, the bands between `F(k+1)` and `F(k+2)` is considered as transition bands.

3.13 MATLAB FIR Filter Design Exercise

3.13.1 FIR Filter Design MATLAB Exercise

3.13.1.1 Design by windowing

Exercise 3.11:

Assuming sampling rate at 48kHz, design an order-40 low-pass filter having cut-off frequency 10kHz by windowing method. In your design, use Hamming window as the windowing function.

3.13.1.2 Parks-McClellan Optimal Design

Exercise 3.12:

Assuming sampling rate at 48kHz, design an order-40 lowpass filter having transition band 10kHz-11kHz using the Parks-McClellan optimal FIR filter design algorithm.

Solutions to Exercises in Chapter 3

Solution to Exercise 3.1:

The DTFT of the unit sample equals a constant (equaling 1). Thus, the Fourier transform of the output equals the transfer function.

Solution to Exercise 3.2:

In sampling a discrete-time signal's Fourier transform L times equally over $[0, 2\pi)$ to form the DFT, the corresponding signal equals the periodic repetition of the original signal.

$$\left(S\left(k\right) \leftrightarrow \sum_{i=-\infty}^{\infty} \left(s\left(n - iL\right)\right) \right) \tag{3.52}$$

To avoid aliasing (in the time domain), the transform length must equal or exceed the signal's duration.

Solution to Exercise 3.3:

The difference equation for an FIR filter has the form

$$y\left(n\right) = \sum_{m=0}^{q} \left(b_m x\left(n - m\right)\right) \tag{3.53}$$

The unit-sample response equals

$$h\left(n\right) = \sum_{m=0}^{q} \left(b_m \delta\left(n - m\right)\right) \tag{3.54}$$

which corresponds to the representation described in a problem[15] of a length-q boxcar filter.

Solution to Exercise 3.4:

The unit-sample response's duration is $q + 1$ and the signal's N_x. Thus the statement is correct.

Solution to Exercise 3.5:

- $\frac{\theta(\omega)}{\omega} = $ constant
- $\theta\left(\omega\right) = K\omega$
- The phase is linear.

Solution to Exercise 3.6:

Yes; in fact it's optimal! (in a certain sense)

Solution to Exercise 3.7:   Gibbs Phenomenon

Solution to Exercise 3.8:

$$\begin{cases} e^{-\left(j\omega \frac{M-1}{2}\right)} \text{ if } -\omega_c \leq \omega \leq \omega_c \\ 0 \text{ if } -\pi \leq \omega < -\omega_c | \omega_c < \omega \leq \pi \end{cases}$$

---

[15]http://cnx.org/content/m10251/latest/#ex2001

Solution to Exercise 3.10:

It's the same problem! To show that, consider an odd-length, symmetric linear phase filter.

$$
\begin{aligned}
H(\omega) &= \sum_{n=0}^{M-1} \left(h(n) e^{-(j\omega n)} \right) \\
&= e^{-\left(j\omega\frac{M-1}{2}\right)} \left(h\left(\frac{M-1}{2}\right) + 2\sum_{n=1}^{L} \left(h\left(\frac{M-1}{2} - n\right) \cos(\omega n) \right) \right)
\end{aligned}
\tag{3.55}
$$

$$
A(\omega) = h(L) + 2\sum_{n=1}^{L} \left(h(L-n) \cos(\omega n) \right)
\tag{3.56}
$$

Where $\left(L \doteq \frac{M-1}{2} \right)$.

Using trigonometric identities (such as $\cos(n\alpha) = 2\cos((n-1)\alpha)\cos(\alpha) - \cos((n-2)\alpha)$), we can rewrite $A(\omega)$ as

$$
A(\omega) = h(L) + 2\sum_{n=1}^{L} \left(h(L-n) \cos(\omega n) \right) = \sum_{k=0}^{L} \left(\alpha_k \cos^k(\omega) \right)
$$

where the α_k are related to the $h(n)$ by a linear transformation. Now, let $x = \cos(\omega)$. This is a one-to-one mapping from $x \in [-1, 1]$ onto $\omega \in [0, \pi]$. Thus $A(\omega)$ is an Lth-order polynomial in $x = \cos(\omega)$!

> IMPLICATION: The alternation theorem holds for the L^∞ filter design problem, too!

Therefore, to determine whether or not a length-M, odd-length, symmetric linear-phase filter is optimal in an L^∞ sense, simply count the alternations in $E(\omega) = W(\omega)(A_d(\omega) - A(\omega))$ in the pass and stop bands. If there are $L + 2 = \frac{M+3}{2}$ or more alternations, $h(n)$, $0 \le n \le M - 1$ is the optimal filter!

Solution to Exercise 3.11:

```
b = fir1(40,10.0/48.0)
```

Solution to Exercise 3.12:

```
b = remez(40,[1 1 0 0],[0 10/48 11/48 1])
```

(a)

~0.11 + 1.0

-0.11

(b)

Figure 3.18: (a) $A(\omega)$, small M (b) $A(\omega)$, large M

Chapter 4

Statistical and Adaptive Signal Processing

4.1 Introduction to Random Signals and Processes

Before now, you have probably dealt strictly with the theory behind signals and systems, as well as look at some the basic characteristics of signals[1] and systems[2]. In doing so you have developed an important foundation; however, most electrical engineers do not get to work in this type of fantasy world. In many cases the signals of interest are very complex due to the randomness of the world around them, which leaves them noisy and often corrupted. This often causes the information contained in the signal to be hidden and distorted. For this reason, it is important to understand these random signals and how to recover the necessary information.

4.1.1 Signals: Deterministic vs. Stochastic

For this study of signals and systems, we will divide signals into two groups: those that have a fixed behavior and those that change randomly. As most of you have probably already dealt with the first type, we will focus on introducing you to random signals. Also, note that we will be dealing strictly with discrete-time signals since they are the signals we deal with in DSP and most real-world computations, but these same ideas apply to continuous-time signals.

4.1.1.1 Deterministic Signals

Most introductions to signals and systems deal strictly with **deterministic signals**. Each value of these signals are fixed and can be determined by a mathematical expression, rule, or table. Because of this, future values of any deterministic signal can be calculated from past values. For this reason, these signals are relatively easy to analyze as they do not change, and we can make accurate assumptions about their past and future behavior.

4.1.1.2 Stochastic Signals

Unlike deterministic signals, **stochastic signals**, or **random signals**, are not so nice. Random signals cannot be characterized by a simple, well-defined mathematical equation

---

[1] http://cnx.org/content/m10057/latest/
[2] http://cnx.org/content/m10084/latest/

Figure 4.1:   An example of a deterministic signal, the sine wave.

Figure 4.2:   We have taken the above sine wave and added random noise to it to come up with a noisy, or random, signal. These are the types of signals that we wish to learn how to deal with so that we can recover the original sine wave.

and their future values cannot be predicted. Rather, we must use probability and statistics to analyze their behavior. Also, because of their randomness, average values (Section 5.3) from a collection of signals are usually studied rather than analyzing one individual signal.

4.1.2 Random Process

As mentioned above, in order to study random signals, we want to look at a collection of these signals rather than just one instance of that signal. This collection of signals is called a **random process**.

> **Definition 5: random process**
> A family or ensemble of signals that correspond to every possible outcome of a certain signal measurement. Each signal in this collection is referred to as a **realization** or **sample function** of the process.

> **Example:**
> As an example of a random process, let us look at the Random Sinusoidal Process below. We use $f[n] = A\sin(\omega n + \phi)$ to represent the sinusoid with a given amplitude and phase. Note that the phase and amplitude of each sinusoid is based on a random number, thus making this a random process.
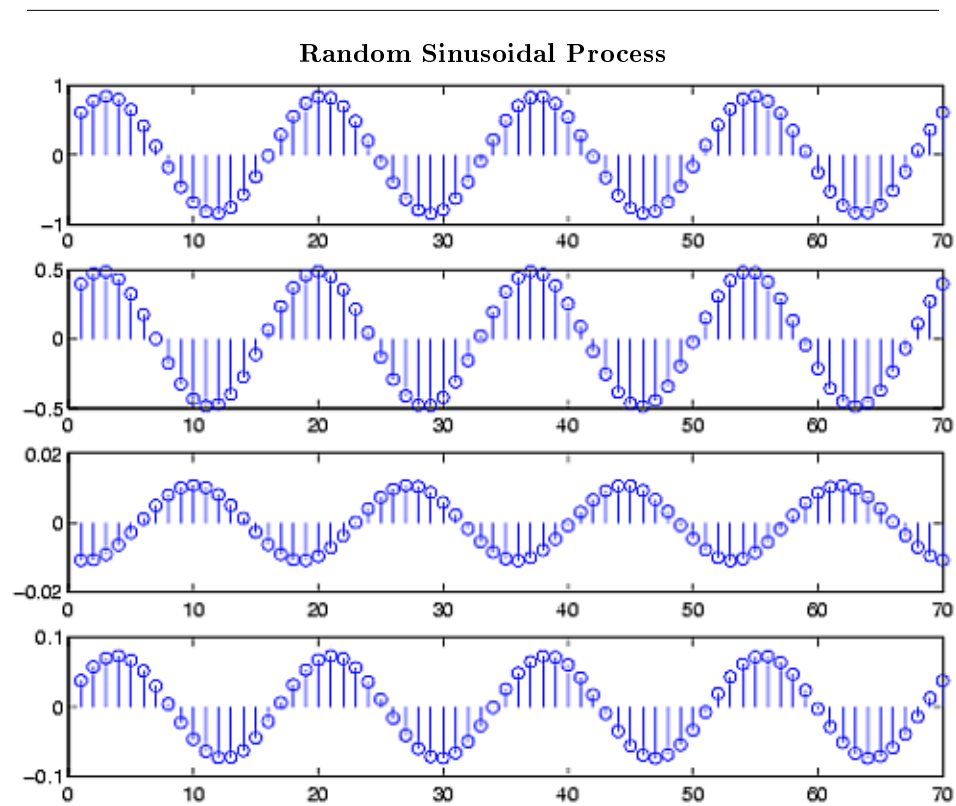
Figure 4.3: A random sinusoidal process, with the amplitude and phase being random numbers.

A random process is usually denoted by $X(t)$ or $X[n]$, with $x(t)$ or $x[n]$ used to represent an individual signal or waveform from this process.

In many notes and books, you might see the following notation and terms used to describe different types of random processes. For a **discrete random process**, sometimes just called a **random sequence**, t represents time that has a finite number of values. If t can take on any value of time, we have a **continuous random process**. Often times discrete and continuous refer to the amplitude of the process, and process or sequence refer to the nature of the time variable. For this study, we often just use **random process** to refer to a general collection of discrete-time signals, as seen above in Figure 4.3.

4.2 Stationary and Nonstationary Random Processes

4.2.1 Introduction

From the definition of a random process (Section 5.1), we know that all random processes are composed of random variables, each at its own unique point in time. Because of this, random processes have all the properties of random variables, such as mean, correlation, variances, etc.. When dealing with groups of signals or sequences it will be important for us to be able to show whether of not these statistical properties hold true for the entire random process. To do this, the concept of **stationary processes** has been developed. The general definition of a stationary process is:

> **Definition 6: stationary process**
> a random process where all of its statistical properties do not vary with time

Processes whose statistical properties do change are referred to as **nonstationary**.

Understanding the basic idea of stationarity will help you to be able to follow the more concrete and mathematical definition to follow. Also, we will look at various levels of stationarity used to describe the various types of stationarity characteristics a random process can have.

4.2.2 Distribution and Density Functions

In order to properly define what it means to be stationary from a mathematical standpoint, one needs to be somewhat familiar with the concepts of distribution and density functions. If you can remember your statistics then feel free to skip this section!

Recall that when dealing with a single random variable, the **probability distribution function** is a simply tool used to identify the probability that our observed random variable will be less than or equal to a given number. More precisely, let X be our random variable, and let x be our given value; from this we can define the distribution function as

$$F_x(x) = \Pr[X \le x] \tag{4.1}$$

This same idea can be applied to instances where we have multiple random variables as well. There may be situations where we want to look at the probability of event *X and Y* both occurring. For example, below is an example of a second-order **joint distribution function**.

$$F_x(x, y) = \Pr[X \le x, Y \le y] \tag{4.2}$$

While the distribution function provides us with a full view of our variable or processes probability, it is not always the most useful for calculations. Often times we will want to look at its derivative, the **probability density function (pdf)**. We define the the pdf as

$$f_x(x) = \frac{d}{dx} F_x(x) \tag{4.3}$$

$$f_x(x)\,dx = \Pr\left[x < X \le x + dx\right] \tag{4.4}$$

Equation 4.4 reveals some of the physical significance of the density function. This equations tells us the probability that our random variable falls within a given interval can be approximated by $f_x(x)\,dx$. From the pdf, we can now use our knowledge of integrals to evaluate probabilities from the above approximation. Again we can also define a **joint density function** which will include multiple random variables just as was done for the distribution function. The density function is used for a variety of calculations, such as finding the expected value or proving a random variable is stationary, to name a few.

> NOTE: The above examples explain the distribution and density functions in terms of a single random variable, X. When we are dealing with signals and random processes, remember that we will have a set of random variables where a different random variable will occur at each time instance of the random process, $X(t_k)$. In other words, the distribution and density function will also need to take into account the choice of time.

4.2.3 Stationarity

Below we will now look at a more in depth and mathematical definition of a stationary process. As was mentioned previously, various levels of stationarity exist and we will look at the most common types.

4.2.3.1 First-Order Stationary Process

A random process is classified as **first-order stationary** if its first-order probability density function remains equal regardless of any shift in time to its time origin. If we let x_{t_1} represent a given value at time t_1, then we define a first-order stationary as one that satisfies the following equation:

$$f_x(x_{t_1}) = f_x(x_{t_1+\tau}) \tag{4.5}$$

The physical significance of this equation is that our density function, $f_x(x_{t_1})$, is completely independent of t_1 and thus any time shift, τ.

The most important result of this statement, and the identifying characteristic of any first-order stationary process, is the fact that the mean is a constant, independent of any time shift. Below we show the results for a random process, X, that is a discrete-time signal, $x[n]$.

$$\begin{aligned} \overline{X} &= m_x[n] \\ &= E[x[n]] \\ &= \text{constant (independentofn)} \end{aligned} \tag{4.6}$$

4.2.3.2 Second-Order and Strict-Sense Stationary Process

A random process is classified as **second-order stationary** if its second-order probability density function does not vary over any time shift applied to both values. In other words, for values x_{t_1} and x_{t_2} then we will have the following be equal for an arbitrary time shift τ.

$$f_x(x_{t_1}, x_{t_2}) = f_x(x_{t_1+\tau}, x_{t_2+\tau}) \tag{4.7}$$

From this equation we see that the absolute time does not affect our functions, rather it only really depends on the time difference between the two variables. Looked at another way, this equation can be described as

$$\Pr\left[X\left(t_1\right) \leq x_1, X\left(t_2\right) \leq x_2\right] = \Pr\left[X\left(t_1 + \tau\right) \leq x_1, X\left(t_2 + \tau\right) \leq x_2\right] \tag{4.8}$$

These random processes are often referred to as **strict sense stationary (SSS)** when *all* of the distribution functions of the process are unchanged regardless of the time shift applied to them.

For a second-order stationary process, we need to look at the autocorrelation function (Section 5.5) to see its most important property. Since we have already stated that a second-order stationary process depends only on the time difference, then all of these types of processes have the following property:

$$\begin{aligned} R_{xx}\left(t, t+\tau\right) &= E\left[X\left(t+\tau\right)\right] \\ &= R_{xx}\left(\tau\right) \end{aligned} \tag{4.9}$$

4.2.3.3 Wide-Sense Stationary Process

As you begin to work with random processes, it will become evident that the strict requirements of a SSS process is more than is often necessary in order to adequately approximate our calculations on random processes. We define a final type of stationarity, referred to as **wide-sense stationary (WSS)**, to have slightly more relaxed requirements but ones that are still enough to provide us with adequate results. In order to be WSS a random process only needs to meet the following two requirements.

1. $\overline{X} = E\left[x\left[n\right]\right] = \text{constant}$

2. $E\left[X\left(t+\tau\right)\right] = R_{xx}\left(\tau\right)$

Note that a second-order (or SSS) stationary process will always be WSS; however, the reverse will not always hold true.

4.3 Random Processes: Mean and Variance

In order to study the characteristics of a random process (Section 5.1), let us look at some of the basic properties and operations of a random process. Below we will focus on the operations of the random signals that compose our random processes. We will denote our random process with X and a random variable from a random process or signal by x.

4.3.1 Mean Value

Finding the average value of a set of random signals or random variables is probably the most fundamental concepts we use in evaluating random processes through any sort of statistical method. *The mean of a random process is the average of all realizations of that process.* In order to find this average, we must look at a random signal over a range of time (possible values) and determine our average from this set of values. The **mean**, or average, of a random process, $x\left(t\right)$, is given by the following equation:

$$\begin{aligned} m_x\left(t\right) &= \mu_x\left(t\right) \\ &= \overline{X} \\ &= E\left[X\right] \\ &= \int_{-\infty}^{\infty} x f\left(x\right) dx \end{aligned} \tag{4.10}$$

This equation may seem quite cluttered at first glance, but we want to introduce you to the various notations used to represent the mean of a random signal or process. Throughout texts and other readings, remember that these will all equal the same thing. The symbol, $\mu_x(t)$, and the X with a bar over it are often used as a short-hand to represent an average, so you might see it in certain textbooks. The other important notation used is, $E[X]$, which represents the "expected value of X" or the mathematical expectation. This notation is very common and will appear again.

If the random variables, which make up our random process, are discrete or quantized values, such as in a binary process, then the integrals become summations over all the possible values of the random variable. In this case, our expected value becomes

$$E[x[n]] = \sum_x (\alpha \Pr[x[n] = \alpha]) \tag{4.11}$$

If we have two random signals or variables, their averages can reveal how the two signals interact. If the product of the two individual averages of both signals do *not* equal the average of the product of the two signals, then the two signals are said to be **linearly independent**, also referred to as **uncorrelated**.

In the case where we have a random process in which only one sample can be viewed at a time, then we will often not have all the information available to calculate the mean using the density function as shown above. In this case we must estimate the mean through the time-average mean (Section 4.3.4), discussed later. For fields such as signal processing that deal mainly with discrete signals and values, then these are the averages most commonly used.

4.3.1.1 Properties of the Mean

- The expected value of a constant, α, is the constant:

$$E[\alpha] = \alpha \tag{4.12}$$

- Adding a constant, α, to each term increases the expected value by that constant:

$$E[X + \alpha] = E[X] + \alpha \tag{4.13}$$

- Multiplying the random variable by a constant, α, multiplies the expected value by that constant.

$$E[\alpha X] = \alpha E[X] \tag{4.14}$$

- The expected value of the sum of two or more random variables, is the sum of each individual expected value.

$$E[X + Y] = E[X] + E[Y] \tag{4.15}$$

4.3.2 Mean-Square Value

If we look at the second **moment** of the term (we now look at x^2 in the integral), then we will have the **mean-square value** of our random process. As you would expect, this is written as

$$\begin{aligned} \overline{X^2} &= E[X^2] \\ &= \int_{-\infty}^{\infty} x^2 f(x)\, dx \end{aligned} \tag{4.16}$$

This equation is also often referred to as the **average power** of a process or signal.

4.3.3 Variance

Now that we have an idea about the average value or values that a random process takes, we are often interested in seeing just how spread out the different random values might be. To do this, we look at the **variance** which is a measure of this spread. The variance, often denoted by σ^2, is written as follows:

$$
\begin{aligned}
\sigma^2 &= \text{Var}(X) \\
&= E\left[(X - E[X])^2\right] \\
&= \int_{-\infty}^{\infty} (x - \overline{X})^2 f(x)\, dx
\end{aligned}
\tag{4.17}
$$

Using the rules for the expected value, we can rewrite this formula as the following form, which is commonly seen:

$$
\begin{aligned}
\sigma^2 &= \overline{X^2} - \left(\overline{X}\right)^2 \\
&= E\left[X^2\right] - (E[X])^2
\end{aligned}
\tag{4.18}
$$

4.3.3.1 Standard Deviation

Another common statistical tool is the standard deviation. Once you know how to calculate the variance, the standard deviation is simply *the square root of the variance*, or σ.

4.3.3.2 Properties of Variance

- The variance of a constant, α, equals zero:

$$
\begin{aligned}
\text{Var}(\alpha) &= \sigma(\alpha)^2 \\
&= 0
\end{aligned}
\tag{4.19}
$$

- Adding a constant, α, to a random variable does not affect the variance because the mean increases by the same value:

$$
\begin{aligned}
\text{Var}(X + \alpha) &= \sigma(X + \alpha)^2 \\
&= \sigma(X)^2
\end{aligned}
\tag{4.20}
$$

- Multiplying the random variable by a constant, α, increases the variance by the square of the constant:

$$
\begin{aligned}
\text{Var}(\alpha X) &= \sigma(\alpha X)^2 \\
&= \alpha^2 \sigma(X)^2
\end{aligned}
\tag{4.21}
$$

- The variance of the sum of two random variables only equals the sum of the variances if the variable are **independent**.

$$
\begin{aligned}
\text{Var}(X + Y) &= \sigma(X + Y)^2 \\
&= \sigma(X)^2 + \sigma(Y)^2
\end{aligned}
\tag{4.22}
$$

Otherwise, if the random variable are *not* independent, then we must also include the covariance of the product of the variables as follows:

$$
\text{Var}(X + Y) = \sigma(X)^2 + 2\text{Cov}(X, Y) + \sigma(Y)^2
\tag{4.23}
$$

4.3.4 Time Averages

In the case where we can not view the entire ensemble of the random process, we must use time averages to estimate the values of the mean and variance for the process. Generally, this will only give us acceptable results for independent and **ergodic** processes, meaning those processes in which each signal or member of the process seems to have the same statistical behavior as the entire process. The time averages will also only be taken over a finite interval since we will only be able to see a finite part of the sample.

4.3.4.1 Estimating the Mean

For the ergodic random process, $x(t)$, we will estimate the mean using the time averaging function defined as

$$\begin{aligned} \overline{X} &= E[X] \\ &= \frac{1}{T} \int_0^T X(t)\, dt \end{aligned} \qquad (4.24)$$

However, for most real-world situations we will be dealing with discrete values in our computations and signals. We will represent this mean as

$$\begin{aligned} \overline{X} &= E[X] \\ &= \frac{1}{N} \sum_{n=1}^N (X[n]) \end{aligned} \qquad (4.25)$$

4.3.4.2 Estimating the Variance

Once the mean of our random process has been estimated then we can simply use those values in the following variance equation (introduced in one of the above sections)

$$\sigma_x{}^2 = \overline{X^2} - \left(\overline{X}\right)^2 \qquad (4.26)$$

4.3.5 Example

Let us now look at how some of the formulas and concepts above apply to a simple example. We will just look at a single, continuous random variable for this example, but the calculations and methods are the same for a random process. For this example, we will consider a random variable having the probability density function described below and shown in Figure 4.4.

$$f(x) = \begin{cases} \frac{1}{10} & \text{if } 10 \le x \le 20 \\ 0 & \text{otherwise} \end{cases} \qquad (4.27)$$

First, we will use Equation 4.10 to solve for the mean value.

$$\begin{aligned} \overline{X} &= \int_{10}^{20} x \frac{1}{10}\, dx \\ &= \frac{1}{10} \left(\frac{x^2}{2}\right) \big|_{x=10}^{20} \\ &= \frac{1}{10}(200 - 50) \\ &= 15 \end{aligned} \qquad (4.28)$$

Using Equation 4.16 we can obtain the mean-square value for the above density function.

$$\begin{aligned} \overline{X^2} &= \int_{10}^{20} x^2 \frac{1}{10}\, dx \\ &= \frac{1}{10} \left(\frac{x^3}{3}\right) \big|_{x=10}^{20} \\ &= \frac{1}{10} \left(\frac{8000}{3} - \frac{1000}{3}\right) \\ &= 233.33 \end{aligned} \qquad (4.29)$$

Probability Density Function

Figure 4.4:   A uniform probability density function.

And finally, let us solve for the variance of this function.

$$\begin{aligned}
\sigma^2 &= \overline{X^2} - \left(\overline{X}\right)^2 \\
&= 233.33 - 15^2 \\
&= 8.33
\end{aligned} \tag{4.30}$$

4.4 Correlation and Covariance of a Random Signal

When we take the expected value (Section 5.3), or average, of a random process (Section 4.1.2), we measure several important characteristics about how the process behaves in general. This proves to be a very important observation. However, suppose we have several random processes measuring different aspects of a system. The relationship between these different processes will also be an important observation. The covariance and correlation are two important tools in finding these relationships. Below we will go into more details as to what these words mean and how these tools are helpful. Note that much of the following discussions refer to just random variables, but keep in mind that these variables can represent random signals or random processes.

4.4.1 Covariance

To begin with, when dealing with more than one random process, it should be obvious that it would be nice to be able to have a number that could quickly give us an idea of how similar the processes are. To do this, we use the **covariance**, which is analogous to the variance of a single variable.

> **Definition 7: Covariance**
> A measure of how much the deviations of two or more variables or processes match.

For two processes, X and Y, if they are *not* closely related then the covariance will be small, and if they are similar then the covariance will be large. Let us clarify this statement by describing what we mean by "related" and "similar." Two processes are "closely related" if their distribution spreads are almost equal and they are around the same, or a very slightly different, mean.

Mathematically, covariance is often written as σ_{xy} and is defined as

$$
\begin{aligned}
\operatorname{cov}(X,Y) &= \sigma_{xy} \\
&= E\left[\left(X - \overline{X}\right)\left(Y - \overline{Y}\right)\right]
\end{aligned}
\tag{4.31}
$$

This can also be reduced and rewritten in the following two forms:

$$
\sigma_{xy} = \overline{(xy)} - (\overline{x})(\overline{y})
\tag{4.32}
$$

$$
\sigma_{xy} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \left(X - \overline{X}\right)\left(Y - \overline{Y}\right) f(x,y)\, dx dy
\tag{4.33}
$$

4.4.1.1 Useful Properties

- If X and Y are independent and uncorrelated or one of them has zero mean value, then
$$
\sigma_{xy} = 0
$$

- If X and Y are orthogonal, then
$$
\sigma_{xy} = -\left(E\left[X\right]E\left[Y\right]\right)
$$

- The covariance is symmetric
$$
\operatorname{cov}(X,Y) = \operatorname{cov}(Y,X)
$$

- Basic covariance identity
$$
\operatorname{cov}(X + Y, Z) = \operatorname{cov}(X, Z) + \operatorname{cov}(Y, Z)
$$

- Covariance of equal variables
$$
\operatorname{cov}(X, X) = \operatorname{Var}(X)
$$

4.4.2 Correlation

For anyone who has any kind of statistical background, you should be able to see that the idea of dependence/independence among variables and signals plays an important role when dealing with random processes. Because of this, the **correlation** of two variables provides us with a measure of how the two variables affect one another.

Definition 8: Correlation
A measure of how much one random variable depends upon the other.

This measure of association between the variables will provide us with a clue as to how well the value of one variable can be predicted from the value of the other. The correlation is equal to the average of the product of two random variables and is defined as

$$
\begin{aligned}
\operatorname{cor}(X,Y) &= E\left[XY\right] \\
&= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} xy f(x,y)\, dx dy
\end{aligned}
\tag{4.34}
$$
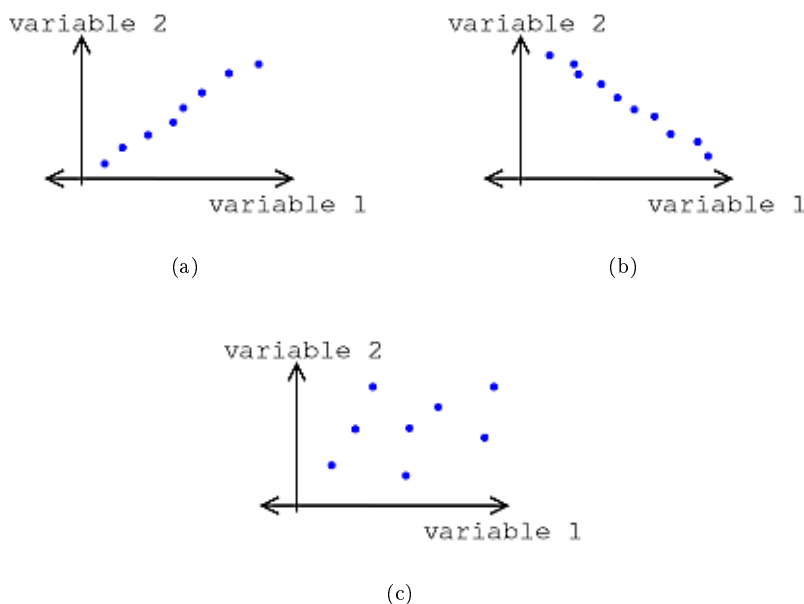
Figure 4.5: Types of Correlation (a) Positive Correlation (b) Negative Correlation (c) Uncorrelated (No Correlation)

4.4.2.1 Correlation Coefficient

It is often useful to express the correlation of random variables with a range of numbers, like a percentage. For a given set of variables, we use the **correlation coefficient** to give us the linear relationship between our variables. The correlation coefficient of two variables is defined in terms of their covariance and standard deviations (Section 4.3.3.1), denoted by σ_x, as seen below

$$\rho = \frac{\text{cov}(X, Y)}{\sigma_x \sigma_y} \tag{4.35}$$

where we will always have

$$-1 \leq \rho \leq 1$$

This provides us with a quick and easy way to view the correlation between our variables. If there is no relationship between the variables then the correlation coefficient will be zero and if there is a perfect positive match it will be one. If there is a perfect inverse relationship, where one set of variables increases while the other decreases, then the correlation coefficient will be negative one. This type of correlation is often referred to more specifically as the **Pearson's Correlation Coefficient**,or Pearson's Product Moment Correlation.

> NOTE: So far we have dealt with correlation simply as a number relating the relationship between any two variables. However, since our goal will be to relate random processes to each other, which deals with signals as a function of time, we will want to continue this study by looking at **correlation functions (Section 5.5)**.

4.4.3 Example

Now let us take just a second to look at a simple example that involves calculating the covariance and correlation of two sets of random numbers. We are given the following data sets:

$$x = \{3, 1, 6, 3, 4\}$$

$$y = \{1, 5, 3, 4, 3\}$$

To begin with, for the covariance we will need to find the expected value (Section 5.3), or mean, of x and y.

$$\overline{x} = \frac{1}{5} \left(3 + 1 + 6 + 3 + 4\right) = 3.4$$

$$\overline{y} = \frac{1}{5} \left(1 + 5 + 3 + 4 + 3\right) = 3.2$$

$$\overline{xy} = \frac{1}{5} \left(3 + 5 + 18 + 12 + 12\right) = 10$$

Next we will solve for the standard deviations of our two sets using the formula below (for a review click here (Section 4.3.3)).

$$\sigma = \sqrt{E\left[\left(X - E\left[X\right]\right)^2\right]}$$

$$\sigma_x = \sqrt{\frac{1}{5}\left(0.16 + 5.76 + 6.76 + 0.16 + 0.36\right)} = 1.625$$

$$\sigma_y = \sqrt{\frac{1}{6}\left(4.84 + 3.24 + 0.04 + 0.64 + 0.04\right)} = 1.327$$

Now we can finally calculate the covariance using one of the two formulas found above. Since we calculated the three means, we will use that formula (Equation 4.32) since it will be much simpler.

$$\sigma_{xy} = 10 - 3.4 \times 3.2 = -0.88$$

And for our last calculation, we will solve for the correlation coefficient, ρ.

$$\rho = \frac{-0.88}{1.625 \times 1.327} = -0.408$$

4.4.3.1 Matlab Code for Example

The above example can be easily calculated using Matlab. Below I have included the code to find all of the values above.

```
x = [3 1 6 3 4];
y = [1 5 3 4 3];

mx = mean(x)
my = mean(y)
mxy = mean(x.*y)

% Standard Dev. from built-in Matlab Functions
```

```
std(x,1)
std(y,1)

% Standard Dev. from Equation Above (same result as std(?,1))
sqrt( 1/5 * sum((x-mx).^2))
sqrt( 1/5 * sum((y-my).^2))

cov(x,y,1)

corrcoef(x,y)
```

4.5 Autocorrelation of Random Processes

Before diving into a more complex statistical analysis of random signals and processes (Section 5.1), let us quickly review the idea of correlation (Section 5.4). Recall that the correlation of two signals or variables is the expected value of the product of those two variables. Since our focus will be to discover more about a random process, a collection of random signals, then imagine us dealing with two samples of a random process, where each sample is taken at a different point in time. Also recall that the key property of these random processes is that they are now functions of time; imagine them as a collection of signals. The expected value (Section 5.3) of the product of these two variables (or samples) will now depend on how quickly they change in regards to *time*. For example, if the two variables are taken from almost the same time period, then we should expect them to have a high correlation. We will now look at a correlation function that relates a pair of random variables from the same process to the time separations between them, where the argument to this correlation function will be the time difference. For the correlation of signals from two different random process, look at the crosscorrelation function (Section 5.6).

4.5.1 Autocorrelation Function

The first of these correlation functions we will discuss is the **autocorrelation**, where each of the random variables we will deal with come from the same random process.

> **Definition 9: Autocorrelation**
> the expected value of the product of a random variable or signal realization with a time-shifted version of itself

With a simple calculation and analysis of the autocorrelation function, we can discover a few important characteristics about our random process. These include:

1. How quickly our random signal or processes changes with respect to the time function

2. Whether our process has a periodic component and what the expected frequency might be

As was mentioned above, the autocorrelation function is simply the expected value of a product. Assume we have a pair of random variables from the same process, $X_1 = X(t_1)$ and $X_2 = X(t_2)$, then the autocorrelation is often written as

$$
\begin{array}{rcl}
R_{\mathrm{xx}}(t_1, t_2) & = & E[X_1 X_2] \\
& = & \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x_1 x_2 f(x_1, x_2)\, dx_2 dx_1
\end{array}
\tag{4.36}
$$

The above equation is valid for stationary and nonstationary random processes. For stationary processes (Section 5.2), we can generalize this expression a little further. Given a wide-sense stationary processes, it can be proven that the expected values from our random process will be independent of the origin of our time function. Therefore, we can say that our autocorrelation function will depend on the time difference and not some absolute time. For this discussion, we will let $\tau = t_2 - t_1$, and thus we generalize our autocorrelation expression as

$$
\begin{aligned}
R_{\mathrm{xx}}\left(t, t+\tau\right) &= R_{\mathrm{xx}}\left(\tau\right) \\
&= E\left[X\left(t\right) X\left(t+\tau\right)\right]
\end{aligned}
\tag{4.37}
$$

for the continuous-time case. In most DSP course we will be more interested in dealing with real signal sequences, and thus we will want to look at the discrete-time case of the autocorrelation function. The formula below will prove to be more common and useful than Equation 4.36:

$$
R_{\mathrm{xx}}\left[n, n+m\right] = \sum_{n=-\infty}^{\infty} \left(x\left[n\right] x\left[n+m\right]\right)
\tag{4.38}
$$

And again we can generalize the notation for our autocorrelation function as

$$
\begin{aligned}
R_{\mathrm{xx}}\left[n, n+m\right] &= R_{\mathrm{xx}}\left[m\right] \\
&= E\left[X\left[n\right] X\left[n+m\right]\right]
\end{aligned}
\tag{4.39}
$$

4.5.1.1 Properties of Autocorrelation

Below we will look at several properties of the autocorrelation function that hold for *stationary* random processes.

- Autocorrelation is an even function for τ

$$
R_{\mathrm{xx}}\left(\tau\right) = R_{\mathrm{xx}}\left(-\tau\right)
$$

- The mean-square value can be found by evaluating the autocorrelation where $\tau = 0$, which gives us

$$
R_{\mathrm{xx}}\left(0\right) = \overline{X^2}
$$

- The autocorrelation function will have its largest value when $\tau = 0$. This value can appear again, for example in a periodic function at the values of the equivalent periodic points, but will never be exceeded.

$$
R_{\mathrm{xx}}\left(0\right) \geq \left|R_{\mathrm{xx}}\left(\tau\right)\right|
$$

- If we take the autocorrelation of a period function, then $R_{\mathrm{xx}}\left(\tau\right)$ will also be periodic with the same frequency.

4.5.1.2 Estimating the Autocorrleation with Time-Averaging

Sometimes the whole random process is not available to us. In these cases, we would still like to be able to find out some of the characteristics of the stationary random process, even if we just have part of one sample function. In order to do this we can *estimate* the autocorrelation from a given interval, 0 to T seconds, of the sample function.

$$
_{\mathrm{xx}}\left(\tau\right) = \frac{1}{T-\tau} \int_0^{T-\tau} x\left(t\right) x\left(t+\tau\right) dt
\tag{4.40}
$$

Figure 4.6: Gaussian density function. By examination, can easily see that the above statement is true - the mean equals zero.

However, a lot of times we will not have sufficient information to build a complete continuous-time function of one of our random signals for the above analysis. If this is the case, we can treat the information we do know about the function as a discrete signal and use the discrete-time formula for estimating the autocorrelation.

$$_{\text{xx}}[m] = \frac{1}{N-m} \sum_{n=0}^{N-m-1} \left(x[n] \, x[n+m] \right) \tag{4.41}$$

4.5.2 Examples

Below we will look at a variety of examples that use the autocorrelation function. We will begin with a simple example dealing with Gaussian White Noise (GWN) and a few basic statistical properties that will prove very useful in these and future calculations.

Example 4.1:

We will let $x[n]$ represent our GWN. For this problem, it is important to remember the following fact about the mean of a GWN function:

$$E[x[n]] = 0$$

Along with being *zero-mean*, recall that GWN is always *independent*. With these two facts, we are now ready to do the short calculations required to find the autocorrelation.

$$R_{\text{xx}}[n, n+m] = E[x[n] \, x[n+m]]$$

Since the function, $x[n]$, is independent, then we can take the product of the individual expected values of both functions.

$$R_{\text{xx}}[n, n+m] = E[x[n]] \, E[x[n+m]]$$

Now, looking at the above equation we see that we can break it up further into two conditions: one when m and n are equal and one when they are not equal. When they are equal we can combine the expected values. We are left with the following piecewise function to solve:

$$R_{\text{xx}}\left[n, n+m\right] = \left\{ \begin{array}{c} E\left[x\left[n\right]\right]E\left[x\left[n+m\right]\right] \text{ if } m \neq 0 \\ E\left[x^2\left[n\right]\right] \text{ if } m = 0 \end{array} \right.$$

We can now solve the two parts of the above equation. The first equation is easy to solve as we have already stated that the expected value of $x\left[n\right]$ will be zero. For the second part, you should recall from statistics that the expected value of the square of a function is equal to the variance. Thus we get the following results for the autocorrelation:

$$R_{\text{xx}}\left[n, n+m\right] = \left\{ \begin{array}{c} 0 \text{ if } m \neq 0 \\ \sigma^2 \text{ if } m = 0 \end{array} \right.$$

Or in a more concise way, we can represent the results as

$$R_{\text{xx}}\left[n, n+m\right] = \sigma^2 \delta\left[m\right]$$

4.6 Crosscorrelation of Random Processes

Before diving into a more complex statistical analysis of random signals and processes (Section 5.1), let us quickly review the idea of correlation (Section 5.4). Recall that the correlation of two signals or variables is the expected value of the product of those two variables. Since our main focus is to discover more about random processes, a collection of random signals, we will deal with two random processes in this discussion, where in this case we will deal with samples from two *different* random processes. We will analyze the expected value (Section 4.3.1) of the product of these two variables and how they correlate to one another, where the argument to this correlation function will be the time difference. For the correlation of signals from the same random process, look at the autocorrelation function (Section 5.5).

4.6.1 Crosscorrelation Function

When dealing with multiple random processes, it is also important to be able to describe the relationship, if any, between the processes. For example, this may occur if more than one random signal is applied to a system. In order to do this, we use the **crosscorrelation function**, where the variables are instances from two different wide sense stationary random processes.

Definition 10: Crosscorrelation
if two processes are wide sense stationary, the expected value of the product of a random variable from one random process with a time-shifted, random variable from a different random process

Looking at the generalized formula for the crosscorrelation, we will represent our two random processes by allowing $U = U\left(t\right)$ and $V = V\left(t - \tau\right)$. We will define the crosscorrelation function as

$$\begin{array}{rcl} R_{uv}\left(t, t-\tau\right) & = & E\left[UV\right] \\ & = & \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} uv f\left(u, v\right) dv du \end{array} \qquad (4.42)$$

Just as the case with the autocorrelation function, if our input and output, denoted as $U\left(t\right)$ and $V\left(t\right)$, are at least jointly wide sense stationary, then the crosscorrelation does not

depend on absolute time; it is just a function of the time difference. This means we can simplify our writing of the above function as

$$R_{uv}(\tau) = E[UV] \tag{4.43}$$

or if we deal with two real signal sequences, $x[n]$ and $y[n]$, then we arrive at a more commonly seen formula for the discrete crosscorrelation function. See the formula below and notice the similarities between it and the convolution (Section 2.5) of two signals:

$$\begin{aligned} R_{xy}(n, n-m) &= R_{xy}(m) \\ &= \sum_{n=-\infty}^{\infty}(x[n]\,y[n-m]) \end{aligned} \tag{4.44}$$

4.6.1.1 Properties of Crosscorrelation

Below we will look at several properties of the crosscorrelation function that hold for two *wide sense stationary (WSS)* random processes.

- Crosscorrelation is *not* an even function; however, it does have a unique symmetry property:
$$R_{xy}(-\tau) = R_{yx}(\tau) \tag{4.45}$$

- The maximum value of the crosscorrelation is not always when the shift equals zero; however, we can prove the following property revealing to us what value the maximum cannot exceed.
$$|R_{xy}(\tau)| \le \sqrt{R_{xx}(0)\,R_{yy}(0)} \tag{4.46}$$

- When two random processes are statistically independent then we have
$$R_{xy}(\tau) = R_{yx}(\tau) \tag{4.47}$$

4.6.2 Examples

Exercise 4.1:

Let us begin by looking at a simple example showing the relationship between two sequences. Using Equation 4.44, find the crosscorrelation of the sequences

$$x[n] = \{\ldots, 0, 0, 2, -3, 6, 1, 3, 0, 0, \ldots\}$$

$$y[n] = \{\ldots, 0, 0, 1, -2, 4, 1, -3, 0, 0, \ldots\}$$

for each of the following possible time shifts: $m = \{0, 3, -1\}$.

4.7 Introduction to Adaptive Filters

In many applications requiring filtering, the necessary frequency response may not be known beforehand, or it may vary with time. (Example; suppression of engine harmonics in a car stereo.) In such applications, an adaptive filter which can automatically design itself and which can track system variations in time is extremely useful. Adaptive filters are used extensively in a wide variety of applications, particularly in telecommunications.

Outline of adaptive filter material

Figure 4.7

1. **Wiener Filters** - L_2 optimal (FIR) filter design in a statistical context

2. **LMS algorithm** - simplest and by-far-the-most-commonly-used adaptive filter algorithm

3. **Stability and performance of the LMS algorithm** - When and how well it works

4. **Applications of adaptive filters** - Overview of important applications

5. **Introduction to advanced adaptive filter algorithms** - Techniques for special situations or faster convergence

4.8 Discrete-Time, Causal Wiener Filter

Stochastic L_2 optimal (least squares) FIR filter design problem: Given a wide-sense stationary (WSS) input signal x_k and desired signal d_k (WSS $\Leftrightarrow E[y_k] = E[y_{k+d}]$, $r_{yz}(l) = E[y_k z_{k+l}]$, $r_{yy}(0) < \infty$,) The Wiener filter is the linear, time-invariant filter minimizing $E[\epsilon^2]$, the variance of the error.

As posed, this problem seems slightly silly, since d_k is already available! However, this idea is useful in a wide cariety of applications.

Example 4.2:

active suspension system design

> NOTE: optimal system may change with different road conditions or mass in car, so an **adaptive** system might be desirable.

Example 4.3:

System identification (radar, non-destructive testing, adaptive control systems)

Exercise 4.2:

Usually one desires that the input signal x_k be "persistently exciting," which, among other things, implies non-zero energy in all frequency bands. Why is this desirable?

y_k = level of car body ... d_k = constant

W = suspension system

x_k = road level

Figure 4.8

X_k

Unknown system

Wiener filter

d_k

y_k

ε_k
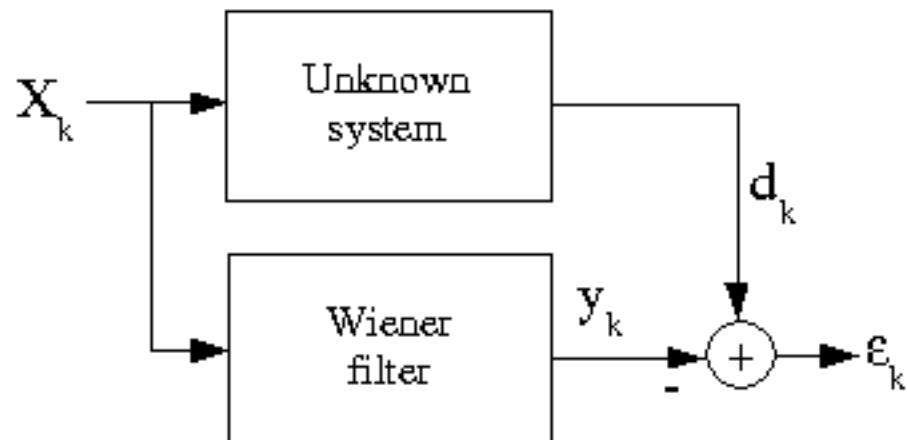
Figure 4.9

4.8.1 Determining the optimal length-N causal FIR Weiner filter

NOTE: for convenience, we will analyze only the causal, real-data case; extensions are straightforward.

$$y_k = \sum_{l=0_{10}}^{M-1_{10}} (w_l x_{k-l})$$

$$
\begin{aligned}
\underset{w_l}{argmin} E\left[\epsilon^{2_{10}}\right] &= E\left[(d_k - y_k)^{2_{10}}\right] \\
&= E\left[\left(d_k - \sum_{l=0_{10}}^{M-1_{10}} (w_l x_{k-l})\right)^{2_{10}}\right] \\
&= E\left[d_k^{2_{10}}\right] - 2_{10}\sum_l \\
&= 0_{10}^{M-1}(w_l E\left[d_k x_{k-l}\right]) + \sum_{l=0_{10}}^{M-1_{10}} \le ft(\sum_{m=0_{10}}^{M-1_{10}}((w_l w_m E\left[x_{k-l} x_{k-m}\right]))
\end{aligned}
$$

$$E\left[\epsilon^2\right] = r_{\mathrm{dd}}(0) - 2\sum_{l=0}^{M-1}(w_l r_{\mathrm{dx}}(l)) + \sum_{l=0}^{M-1}\left(\sum_{m=0}^{M-1}(w_l w_m r_{\mathrm{xx}}(l-m))\right)$$

where

$$r_{\mathrm{dd}}(0) = E\left[d_k^2\right]$$
$$r_{\mathrm{dx}}(l) = E\left[d_k X_{k-l}\right]$$
$$r_{\mathrm{xx}}(l-m) = E\left[x_k x_{k+l-m}\right]$$

This can be written in matrix form as

$$E\left[\epsilon^2\right] = r_{\mathrm{dd}}(0) - 2\mathbf{P}\mathbf{W}^T + \mathbf{W}^T R\mathbf{W}$$

where

$$\mathbf{P} = \begin{pmatrix} r_{\mathrm{dx}}(0) \\ r_{\mathrm{dx}}(1) \\ \vdots \\ r_{\mathrm{dx}}(M-1) \end{pmatrix}$$

$$R = \begin{pmatrix} r_{\mathrm{xx}}(0) & r_{\mathrm{xx}}(1) & \dots & \dots & r_{\mathrm{xx}}(M-1) \\ r_{\mathrm{xx}}(1) & r_{\mathrm{xx}}(0) & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & r_{\mathrm{xx}}(0) & r_{\mathrm{xx}}(1) \\ r_{\mathrm{xx}}(M-1) & \dots & \dots & r_{\mathrm{xx}}(1) & r_{\mathrm{xx}}(0) \end{pmatrix}$$

To solve for the optimum filter, compute the gradient with respect to the top weights vector W

$$\left(\nabla \doteq \begin{pmatrix} \frac{\partial}{\partial w_0}(\epsilon^2) \\ \frac{\partial}{\partial w_1}(\epsilon^2) \\ \vdots \\ \frac{\partial}{\partial w_{M-1}}(\epsilon^2) \end{pmatrix}\right)$$

$$\nabla = -(2\mathbf{P}) + 2R\mathbf{W}$$

(recall $\frac{d}{d\mathbf{W}}\left(\mathbf{A}^T W\right) = \mathbf{A}^T$, $\frac{d}{d\mathbf{W}}(\mathbf{W}M\mathbf{W}) = 2M\mathbf{W}$ for symmetric M) setting the gradient equal to zero $\Rightarrow$

$$\left(W_{\mathrm{opt}} R = \mathbf{P} \Rightarrow W_{\mathrm{opt}} = R^{-1}\mathbf{P}\right)$$

Since R is a correlation matrix, it must be non-negative definite, so this is a minimizer. For R positive definite, the minimizer is unique.

4.9 Practical Issues in Wiener Filter Implementation

The weiner-filter, $W_{\text{opt}} = R^{-1}\mathbf{P}$, is ideal for many applications. But several issues must be addressed to use it in practice.

Exercise 4.3:

In practice one usually won't know exactly the statistics of x_k and d_k (i.e. R and $\mathbf{P}$) needed to compute the Weiner filter.

How do we surmount this problem?

Exercise 4.4:

In many applications, the statistics of x_k, d_k vary slowly with time.

How does one develop an **adaptive** system which tracks these changes over time to keep the system near optimal at all times?

Exercise 4.5:

How can $\hat{r_{\text{xx}}^k}(l)$ be computed efficiently?

Exercise 4.6:

how does one choose N?

4.9.1 Tradeoffs

Larger $N \rightarrow$ more accurate estimates of the correlation values $\rightarrow$ better $\hat{W}_{\text{opt}}$. However, larger N leads to slower adaptation.

> NOTE: The success of adaptive systems depends on x, d being roughly stationary over at least N samples, $N > M$. That is, all adaptive filtering algorithms require that the underlying system varies slowly with respect to the sampling rate and the filter length (although they can tolerate occasional step discontinuities in the underlying system).

4.9.2 Computational Considerations

As presented here, an adaptive filter requires computing a matrix inverse at each sample. Actually, since the matrix R is Toeplitz, the linear system of equations can be sovled with $O\left(M^2\right)$ computations using Levinson's algorithm, where M is the filter length. However, in many applications this may be too expensive, especially since computing the filter output itself requires $O\left(M\right)$ computations. There are two main approaches to resolving the computation problem

1. Take advantage of the fact that R^{k+1} is only slightly changed from R^k to reduce the computation to $O\left(M\right)$; these algorithms are called Fast Recursive Least Squareds algorithms; all methods proposed so far have stability problems and are dangerous to use.

2. Find a different approach to solving the optimization problem that doesn't require explicit inversion of the correlation matrix.

> NOTE: Adaptive algorithms involving the correlation matrix are called **Recursive least Squares** (RLS) algorithms. Historically, they were developed after the LMS algorithm, which is the slimplest and most widely used approach $O\left(M\right)$. $O\left(M^2\right)$ RLS algorithms are used in applications requiring very fast adaptation.

Figure 4.10

4.10 Quadratic Minimization and Gradient Descent

4.10.1 Quadratic minimization problems

The least squares optimal filter design problem is quadratic in the filter coefficients:

$$E\left[\epsilon^2\right] = r_{\text{dd}}\left(0\right) - 2\mathbf{P}^T\mathbf{W} + \mathbf{W}^T R\mathbf{W}$$

If R is positive definite, the error surface $E\left[\epsilon^2\right]\left(w_0, w_1, \ldots, w_{M-1}\right)$ is a unimodal "bowl" in $\mathbb{R}^N$. The problem is to find the bottom of the bowl. In an adaptive filter context, the shape and bottom of the bowl may drift slowly with time; hopefully slow enough that the adaptive algorithm can track it.

For a quadratic error surface, the bottom of the bowl can be found in one step by computing $R^{-1}\mathbf{P}$. Most modern nonlinear optimization methods (which are used, for example, to solve the L^P optimal IIR filter design problem!) locally approximate a nonlinear function with a second-order (quadratic) Taylor series approximation and step to the bottom of this quadratic approximation on each iteration. However, an older and simpler appraoch to nonlinear optimaztion exists, based on **gradient descent**. The idea is to iteratively find the minimizer by computing the gradient of the error function: $E\nabla = \frac{\partial}{\partial w_i}\left(E\left[\epsilon^2\right]\right)$. The gradient is a vector in $\mathbb{R}^M$ pointing in the steepest uphill direction on the error surface at a given point $\mathbf{W}^i$, with ∇ having a magnitude proportional to the slope of the error surface in this steepest direction.

By updating the coefficient vector by taking a step *opposite* the gradient direction : $\mathbf{W}^{i+1} = \mathbf{W}^i - \mu\nabla^i$, we go (locally) "downhill" in the steepest direction, which seems to be a sensible way to iteratively solve a nonlinear optimization problem. The performance obviously depends on μ; if μ is too large, the iterations could bounce back and forth up out of the bowl. However, if μ is too small, it could take many iterations to approach the bottom. We will determine criteria for choosing μ later.

Contour plot of ε-squared

Figure 4.11

In summary, the gradient descent algorithm for solving the Weiner filter problem is:

$$\text{Guess } W^0$$

$$\text{do } i = 1, \infty$$

$$\nabla^i = -(2P) + 2RW^i$$

$$W^{i+1} = W^i - \mu\nabla^i$$

$$\text{repeat}$$

$$W_{\text{opt}} = W^\infty$$

The gradient descent idea is used in the LMS adaptive fitler algorithm. As presented, this alogrithm costs $O\left(M^2\right)$ computations per iteration and doesn't appear very attractive, but LMS only requires $O\left(M\right)$ computations and is stable, so it is very attractive when computation is an issue, even thought it converges more slowly then the RLS algorithms we have discussed so far.

4.11 The LMS Adaptive Filter Algorithm

Recall the Weiner filter problem $\{x_k\}$, $\{d_k\}$ jointly wide sense stationary
Find W minimizing $E\left[\epsilon_k{}^2\right]$

$$\epsilon_k = d_k - y_k = d_k - \sum_{i=0}^{M-1} (w_i x_{k-i}) = d_k - {X^k}^T W^k$$

$$X^k = \begin{pmatrix} x_k \\ x_{k-1} \\ \vdots \\ x_{k-M+1} \end{pmatrix}$$

Figure 4.12

$$W^k = \begin{pmatrix} w_0^k \\ w_1^k \\ \vdots \\ w_{M-1}^k \end{pmatrix}$$

The superscript denotes absolute time, and the subscript denotes time or a vector index.
the solution can be found by setting the gradient $= 0$

$$
\begin{aligned}
\nabla^k &= \frac{\partial}{\partial W}\left(E\left[\epsilon_k^2\right]\right) \\
&= E\left[2\epsilon_k\left(-X^k\right)\right] \\
&= E\left[-2\left(d_k - X^{k^T}W_k\right)X^k\right] \\
&= -\left(2E\left[d_k X^k\right]\right) + \left(E\left[X^{k^T}\right]\right)W \\
&= -2P + 2RW
\end{aligned}
\tag{4.48}
$$

$$\Rightarrow \left(W_{\text{opt}} = R^{-1}P\right)$$

Alternatively, W_{opt} can be found iteratively using a gradient descent technique

$$W^{k+1} = W^k - \mu\nabla^k$$

In practice, we don't know R and P exactly, and in an adaptive context they may be slowly varying with time.

To find the (approximate) Wiener filter, some approximations are necessary. As always, the key is to make the *right* approximations!

GOOD IDEA: Approximate R and P: $\Rightarrow$ RLS methods, as discussed last time.

BETTER IDEA: Approximate the gradient!

$$\nabla^k = \frac{\partial}{\partial W}\left(E\left[\epsilon_k^2\right]\right)$$

Note that ϵ_k^2 itself is a very noisy approximation to $E\left[\epsilon_k^2\right]$. We can get a noisy approximation to the gradient by finding the gradient of ϵ_k^2! Widrow and Hoff first published the LMS algorithm, based on this clever idea, in 1960.

$$\hat{\nabla}^k = \frac{\partial}{\partial W}\left(\epsilon_k^2\right) = 2\epsilon_k\frac{\partial}{\partial W}\left(d_k - W^{k^T}X^k\right) = 2\epsilon_k\left(-X^k\right) = -\left(2\epsilon_k X^k\right)$$

This yields the LMS adaptive filter algorithm

Example 4.4:   The LMS Adaptive Filter Algorithm

1. $y_k = {W^k}^T X^k = \sum_{i=0}^{M-1} \left(w_i^k x_{k-i} \right)$

2. $\epsilon_k = d_k - y_k$

3. $W^{k+1} = W^k - \mu \hat{\nabla}^k = W^k - \mu \left(-2\epsilon_k X^k \right) = W^k + 2\mu\epsilon_k X^k$ $\left(w_i^{k+1} = w_i^k + 2\mu\epsilon_k x_{k-i} \right)$

The LMS algorithm is often called a **stochastic gradient** algorithm, since $\hat{\nabla}^k$ is a noisy gradient. This is *by far* the most commonly used adaptive filtering algorithm, because

1. it was the first

2. it is very simple

3. in practice it works well (except that sometimes it converges slowly)

4. it requires relatively litle computation

5. it updates the tap weights every sample, so it continually adapts the filter

6. it tracks slow changes in the signal statistics well

4.11.1 Computational Cost of LMS

To Compute $\Rightarrow$	y_k	ϵ_k	W^{k+1}	$=$ **Total**
multiplies	M	0	$M+1$	$2M+1$
adds	$M-1$	1	M	$2M$

So the LMS algorithm is $O\left(M\right)$ per sample. In fact, it is nicely balanced in that the filter computation and the adaptation require the same amount of computation.

Note that the parameter μ plays a very important role in the LMS algorithm. It can also be varied with time, but usually a constant μ ("convergence weight facor") is used, chosen after experimentation for a given application.

4.11.1.1 Tradeoffs

large μ: fast convergence, fast adaptivity

small μ: accurate $W \rightarrow$ less misadjustment error, stability

4.12 First Order Convergence Analysis of the LMS Algorithm

4.12.1 Analysis of the LMS algorithm

It is important to analyze the LMS algorithm to determine under what conditions it is stable, whether or not it converges to the Wiener solution, to determine how quickly it converges, how much degradation is suffered due to the noisy gradient, etc. In particular, we need to know how to choose the parameter μ.

4.12.1.1 Mean of W

does W^k, $k \to \infty$ approach the Wiener solution? (since W^k is always somewhat random in the approximate gradient-based LMS algorithm, we ask whether the expected value of the filter coefficients converge to the Wiener solution)

$$
\begin{aligned}
E\left[W^{k+1}\right] &= W^{k+1*} \\
&= E\left[W^k + 2\mu\epsilon_k X^k\right] \\
&= W^{k*} + 2\mu E\left[d_k X^k\right] + 2\mu E\left[-\left(\left(W^{k^T}X^k\right)X^k\right)\right] \\
&= W^{k*} + 2\mu P + \left(-\left(2\mu E\left[\left(W^{k^T}X^k\right)X^k\right]\right)\right)
\end{aligned}
\tag{4.49}
$$

4.12.1.1.1 Patently False Assumption

X^k and X^{k-i}, X^k and d^{k-i}, and d_k and d_{k-i} are statistically independent, $i \neq 0$. This assumption is obviously false, since X^{k-1} is the same as X^k except for shifting down the vector elements one place and adding one new sample. We make this assumption because otherwise it becomes extremely difficult to analyze the LMS algorithm. (First good analysis not making this assumption: *Macchi and Eweda*[?]) Many simulations and much practical experience has shown that the results one obtains with analyses based on the patently false assumption above are quite accurate in most situations

With the independence assumption, W^k (which depends only on previous X^{k-i}, d^{k-i}) is statitically independent of X^k, and we can simplify $E\left[\left(W^{k^T}X^k\right)X^k\right]$

Now $\left(W^{k^T}X^k\right)X^k$ is a vector, and

$$
\begin{aligned}
E\left[\left(W^{k^T}X^k\right)X^k\right] &= E\left[\begin{pmatrix} \vdots \\ \left(\sum_{i=0}^{M-1}\left(w_i^k x_{k-i}\right)\right)x_{k-j} \\ \vdots \end{pmatrix}\right] \\
&= \begin{pmatrix} \vdots \\ \sum_{i=0}^{M-1}\left(E\left[w_i^k x_{k-i}x_{k-j}\right]\right) \\ \vdots \end{pmatrix} \\
&= \begin{pmatrix} \vdots \\ \sum_{i=0}^{M-1}\left(\left(w_i^k\right)E\left[x_{k-i}x_{k-j}\right]\right) \\ \vdots \end{pmatrix} \\
&= \begin{pmatrix} \vdots \\ \sum_{i=0}^{M-1}\left(w_i^{k*}r_{\mathrm{xx}}\left(i-j\right)\right) \\ \vdots \end{pmatrix} \\
&= RW^{k*}
\end{aligned}
\tag{4.50}
$$

where $R = E\left[X^k X^{k^T}\right]$ is the data correlation matrix.

Putting this back into our equation

$$
\begin{aligned}
\overline{W^{k+1}} &= \overline{W^k} + 2\mu P + \left(-\left(2\mu R\overline{W^k}\right)\right) \\
&= (I - 2\mu R)\overline{W^k} + 2\mu P
\end{aligned}
\tag{4.51}
$$

Now *if* $\overline{W^{k \to \infty}}$ converges to a vector of finite magnitude ("convergence in the mean"), what does it converge to?

If $\overline{W^k}$ converges, then as $k \to \infty$, $\overline{W^{k+1}} \approx \overline{W^k}$, and

$$\overline{W^\infty} = (I - 2\mu R)\,\overline{W^\infty} + 2\mu P$$

$$2\mu R \overline{W^\infty} = 2\mu P$$

$$R\overline{W^\infty} = P$$

or

$$\overline{W_{\text{opt}}} = R^{-1}P$$

the Wiener solution!

So the LMS algorithm, *if* it converges, gives filter coefficients which on average are the Wiener coefficients! This is, of course, a desirable result.

4.12.1.2 First-order stability

But does $\overline{W^k}$ converge, or under what conditions?

Let's rewrite the analysis in term of $\overline{V^k}$, the "mean coefficient error vector" $\overline{V^k} = \overline{W^k} - W_{\text{opt}}$, where W_{opt} is the Wiener filter

$$\overline{W^{k+1}} = \overline{W^k} - 2\mu R \overline{W^k} + 2\mu P$$

$$\overline{W^{k+1}} - W_{\text{opt}} = \overline{W^k} - W_{\text{opt}} + \left(- \left(2\mu R \overline{W^k} \right) \right) + 2\mu R W_{\text{opt}} - 2\mu R W_{\text{opt}} + 2\mu P$$

$$\overline{V^{k+1}} = \overline{V^k} - 2\mu R \overline{V^k} + \left(- \left(2\mu R W_{\text{opt}} \right) \right) + 2\mu P$$

Now $W_{\text{opt}} = R^{-1}$, so

$$\overline{V^{k+1}} = \overline{V^k} - 2\mu R \overline{V^k} + \left(- \left(2\mu R R^{-1}P \right) \right) + 2\mu P = (I - 2\mu R)\,\overline{V^k}$$

We wish to know under what conditions $\overline{V^{k \to \infty}} \to \overline{0}$?

4.12.1.2.1 Linear Algebra Fact

Since R is positive definite, real, and symmetric, all the eigenvalues are real and positive. Also, we can write R as $\left(Q^{-1}\Lambda Q \right)$, where Λ is a diagonal matrix with diagonal entries λ_i equal to the eigenvalues of R, and Q is a unitary matrix with rows equal to the eigenvectors corresponding to the eigenvalues of R.

Using this fact,

$$V^{k+1} = \left(I - 2\mu \left(Q^{-1}\Lambda Q \right) \right) V^k$$

multiplying both sides through on the left by Q: we get

$$Q\overline{V^{k+1}} = (Q - 2\mu\Lambda Q)\,\overline{V^k} = (1 - 2\mu\Lambda)\,Q\overline{V^k}$$

Let $V^{'} = QV$:

$$V^{'k+1} = (1 - 2\mu\Lambda)\,V^{'k}$$

Note that $V^{'}$ is simply V in a rotated coordinate set in $\mathbb{R}^m$, so convergence of $V^{'}$ implies convergence of V.

Since $1 - 2\mu\Lambda$ is diagonal, all elements of $V^{'}$ evolve independently of each other. Convergence (stability) bolis down to whether all M of these scalar, first-order difference equations are stable, and thus $\to 0$.

$$V_i^{'k+1} = (1 - 2\mu\lambda_i)\,V_i^{'k} \quad , \quad i = [1, 2, \ldots, M]$$

These equations converge to zero if $|1 - 2\mu\lambda_i| < 1$, or $|\mu\lambda_i| < 1$, μ and λ_i are positive, so we require $\mu < \frac{1}{\lambda_i}$, so for convergence in the mean of the LMS adaptive filter, we require

$$\mu < \frac{1}{\lambda_{\max}} \tag{4.52}$$

This is an elegant theoretical result, but in practice, we may not know $\lambda_{\max}$, it may be time-varying, and we certainly won't want to compute it. However, another useful mathematical fact comes to the rescue...

$$\operatorname{tr}(R) = \sum_{i=1}^{M} r_{\text{ii}} = \sum_{i=1}^{M} \lambda_i \geq \lambda_{\max}$$

Since the eigenvalues are all positive and real.

For a correlation matrix, $r_{\text{ii}} = r(0)$, $i \in \{1, M\}$. So $\operatorname{tr}(R) = Mr(0) = ME[x_k x_k]$. We can easily estimate $r(0)$ with $O(1)$ computations/sample, so in practice we might require

$$\mu < \frac{1}{M r(0)}$$

as a conservative bound, and perhaps adapt μ accordingly with time.

4.12.1.3 Rate of convergence

Each of the modes decays as

$$(1 - 2\mu\lambda_i)^k$$

GOOD NEWS: The *initial* rate of convergence is dominated by the fastest mode $1 - 2\mu\lambda_{\max}$. This is not surprising, since a dradient descent method goes "downhill" in the steepest direction

BAD NEWS: The *final* rate of convergence is dominated by the slowest mode $1 - 2\mu\lambda_{\min}$. For small $\lambda_{\min}$, it can take a long time for LMS to converge.

Note that the convergence behavior depends on the data (via R). LMS converges relatively quickly for roughly equal eigenvalues. Unequal eigenvalues slow LMS down a lot.

4.13 Adaptive Equalization

GOAL: Design an approximate inverse filter to cancel out as much distortion as possible.

In principle, $WH \approx z^{-\Delta}$, or $W \approx \frac{z^{-\Delta}}{H}$, so that the overall response of the top path is approximately $\delta(n - \Delta)$. However, limitations on the form of W (FIR) and the presence of noise cause the equalization to be imperfect.

4.13.1 Important Application

Channel equalization in a digital communication system. If the channel distorts the pulse shape, the matched filter will no longer be matched, intersymbol interference may increase, and the system performance will degrade.

An adaptive filter is often inserted in front of the matched filter to compensate for the channel. This is, of course, unrealizable, since we do not have access to the original transmitted signal, s_k.

There are two common solutions to this problem:

Figure 4.13

Figure 4.14

Figure 4.15

Decision-directed equalizer
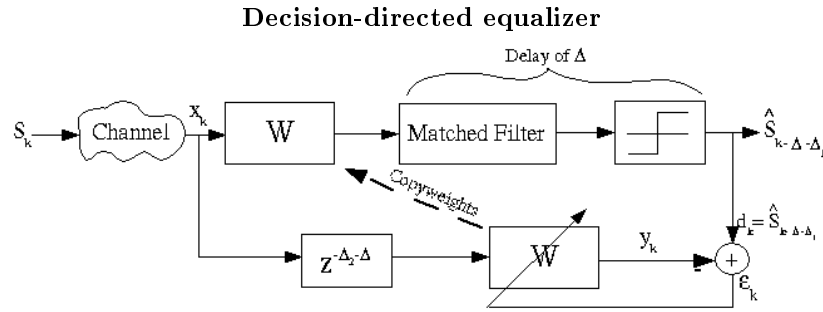
Figure 4.16

1. Periodically broadcast a known **training signal**. The adaptation is switched on only when the training signal is being broadcast and thus s_k is known.

2. Decision-directed feedback: If the overall system is working well, then the output $\hat{s}_{k-\Delta_0}$ should almost always equal $s_{k-\Delta_0}$. We can thus use our received digital communication signal as the desired signal, since it has been cleaned of noise (we hope) by the nonlinear threshold device! As long as the error rate in $\hat{s}_k$ is not too high (say $< 75\%$), this method works. Otherwise, d_k is so inaccurate that the adaptive filter can never find the Wiener solution. This method is widely used in the telephone system and other digital communication networks.

Solutions to Exercises in Chapter 4

Solution to Exercise 4.1:

1. For $m = 0$, we should begin by finding the product sequence $s[n] = x[n]y[n]$. Doing this we get the following sequence:

$$s[n] = \{\ldots, 0, 0, 2, 6, 24, 1, -9, 0, 0, \ldots\}$$

and so from the sum in our crosscorrelation function we arrive at the answer of

$$R_{xy}(0) = 22$$

2. For $m = 3$, we will approach it the same was we did above; however, we will now shift $y[n]$ to the right. Then we can find the product sequence $s[n] = x[n]y[n-3]$, which yields

$$s[n] = \{\ldots, 0, 0, 0, 0, 0, 1, -6, 0, 0, \ldots\}$$

and from the crosscorrelation function we arrive at the answer of

$$R_{xy}(3) = -6$$

3. For $m = -1$, we will again take the same approach; however, we will now shift $y[n]$ to the left. Then we can find the product sequence $s[n] = x[n]y[n+1]$, which yields

$$s[n] = \{\ldots, 0, 0, -4, -12, 6, -3, 0, 0, 0, \ldots\}$$

and from the crosscorrelation function we arrive at the answer of

$$R_{xy}(-1) = -13$$

Solution to Exercise 4.3:

Estimate the statistics

$$\hat{r_{\mathrm{xx}}}(l) \approx \frac{1}{N}\sum_{k=0}^{N-1}(x_k x_{k+l})$$

$$\hat{r_{\mathrm{xd}}}(l) \approx \frac{1}{N}\sum_{k=0}^{N-1}(d_k x_{k-l})$$

then solve $\hat{W}_{\mathrm{opt}} = \hat{R^{-1}} = \hat{P}$

Solution to Exercise 4.4:

Use short-time windowed estiamtes of the correlation functions.

EQUATION IN QUESTION:

$$\left(\hat{r_{\mathrm{xx}}}(l)\right)^k = \frac{1}{N}\sum_{m=0}^{N-1}(x_{k-m}x_{k-m-l})$$

$$\left(\hat{r_{\mathrm{dx}}}(l)\right)^k = \frac{1}{N}\sum_{m=0}^{N-1}(x_{k-m-l}d_{k-m})$$

and $W_{\mathrm{opt}}{}^k \approx \left(\hat{R}_k\right)^{-1}\hat{P}_k$

Solution to Exercise 4.5:

Recursively!

$$r_{\text{xx}}^k (l) = r_{\text{xx}}^{k-1} (l) + x_k x_{k-l} - x_{k-N} x_{k-N-l}$$

This is critically stable, so people usually do

$$(1 - \alpha) \, r_{\text{xx}}^{k} (l) = \alpha r_{\text{xx}}^{k-1} (l) + x_k x_{k-l}$$

Glossary

A Autocorrelation

1. the expected value of the product of a random variable or signal realization with a time-shifted version of itself

C Correlation

1. A measure of how much one random variable depends upon the other.

Covariance

1. A measure of how much the deviations of two or more variables or processes match.

Crosscorrelation

1. if two processes are wide sense stationary, the expected value of the product of a random variable from one random process with a time-shifted, random variable from a different random process

D difference equation

1. An equation that shows the relationship between consecutive values of a sequence and the differences among them. They are often rearranged as a recursive formula so that a systems output can be computed from the input signal and past outputs.

Example:

$$y[n] + 7y[n-1] + 2y[n-2]$$

$$= \mathrm{x}[n] - 4x[n-1] \ (4.1)$$

F FFT

1. (Fast Fourier Transform) An efficient **computational algorithm** for computing the DFT[3].

P poles

1. The value(s) for z where $Q(z) = 0$.

2. The complex frequencies that make the overall gain of the filter transfer function infinite.

R random process

1. A family or ensemble of signals that correspond to every possible outcome of a certain signal measurement. Each signal in this collection is referred to as a **realization** or **sample function** of the process.

   Example:

   As an example of a random process, let us look at the Random Sinusoidal Process below. We use $f[n] = A\sin(\omega n + \phi)$ to represent the sinusoid with a given amplitude and phase. Note that the phase and amplitude of each sinusoid is based on a random number, thus making this a random process.

S stationary process

1. a random process where all of its statistical properties do not vary with time

Z zeros

1. The value(s) for z where $P(z) = 0$.

2. The complex frequencies that make the overall gain of the filter transfer function zero.

---

[3]http://cnx.org/content/m10249/latest/

Index of Keywords and Terms

Keywords are listed by the section with that keyword (page numbers are in parentheses). Keywords do not necessarily appear in the text of the page. They are merely associated with that section. *Ex.* apples, § 1.1 (1) **Terms** are referenced by the page they appear on. *Ex.* apples, 1