HOMEWORK PROBLEMS FOR:
"THE SCIENTIST AND ENGINEER'S GUIDE TO DIGITAL SIGNAL PROCESSING"

Contents

CHAPTER 2: STATISTICS, PROBABILITY AND NOISE

1. A signal contains 100,000 samples, and each sample is represented by 10bits. Assume that addition and subtraction require 1 microsecond, multiplication and division require 3 microseconds, the square-root requires 10 microseconds, and other programming actions (such as array indexing and loop control) are negligible. Find the time required to calculate the mean and standard deviation of the signal using:
   a. The direct method, as shown in Table 2-1.
   b. Running statistics, as shown in Table 2-2.
   c. (You will need to modify the program in Table 2-2 to calculate the mean and standard deviation only of the entire signal, not after each point. That is, move the NEXT I% in line 320 to line 250).
   d. The histogram method, as shown in Table 2-3. Repeat (c) for the case that each sample is represented by 32 bits. How much memory is required to hold the histogram for this calculation, assuming each bin requires two bytes?
2. You are asked to evaluate a new device for detecting cancer in humans. When a healthy person is tested, the device produces a number that follows a normal distribution with a mean of 100 and a standard deviation of 10. When a person with cancer is examined, the resulting number is normally distributed with a mean of 120 and a standard deviation of 8. When the device is used on a person of unknown health, a threshold of 110 is used to
   make the diagnosis: if the reading is < 110, the person is classified as healthy; if the reading is $>$ 110, the person is considered to have cancer.
   a. Sketch the two pdfs, and indicate on your sketch the location of the threshold.
   b. For the healthy distribution, how many standard deviations above the mean is the threshold?
   c. For the sick distribution, how many standard deviations below the mean is the threshold?
   d. If a healthy person is tested with the system, what is the probability that the reading produced will be less than 110? Greater than 110?
   e. If a sick person is tested with the system, what is the probability that the reading produced will be less than 110? Greater than 110?
   f. What percentage of sick patients are incorrectly classified as healthy?
   g. What percentage of healthy patients are incorrectly classified as sick.
3. Using the same cancer detection system as in problem 2, where should the threshold be set to insure that 99% of all sick persons being examined are reported as sick? At this threshold, what fraction of healthy persons are incorrectly reported as being sick?

4. Twelve financial experts are asked to predict the stock market price 30 days in advance. The values they provide are: 996, 868, 855, 956, 867, 933, 866, 887, 936, 901, 818, 956. In 30 days, the true stock market price is
found to be 876.
    a. What is the mean of the predictions?
    b. What is the standard deviation of the predictions?
    c. What is the accuracy of the experts' prediction?
    d. What is the precision of an experts' prediction?
5. An astronomer measures the brightness of a star on 30 consecutive nights. Due to atmospheric turbulence and other random errors, the measurements have a coefficient-of-variation of 3.0%. One of the measurements is found to be 6.9% higher than the average.
    a. What is the signal-to-noise ratio of the 30 sample signal?
    b. What is the probability that any one measurement will be at least 6.9% higher than the mean from random error?
    c. Can the astronomer conclude that the 6.9% reading is a result of the star changing in brightness? Explain.
    d. Repeat (b) and (c) for the high measurement being 10.2% above the mean.
6. When two or more random signals are added, the resulting signal has a mean equal to the sum of the means of the component signals. Likewise, the variance of the resulting signal is equal to the sum of the variances of the component signals. Derive an equation showing how the standard deviation of the resulting signal is related to the standard deviations of the component signals. (This is often called "adding in quadrature").
7. Find the mean and SD of the signal that results from adding the random signals indicated.
    a. 1 volt mean, 20 mV SD; 1 volt mean, 20 mV SD.
    b. 1 volt mean, 20 mV SD; 1 volt mean, 2 mV SD.
    c. 1 volt mean, 20 mV SD; 1 volt mean, 0.2 mV SD.
    d. 3 volt mean, 10 mV SD; 5 volt mean, 15 mV SD.
    e. 10 volt mean, 10 mV SD; 5 volt mean, 15 mV SD; 0.2 Volt mean, 100 mV SD.
8. 8. Electronic systems contain many noise sources; however, the total noise of a system is usually dominated by only one of these noise sources. Based on your answers in the last two problems, explain why this is so.

CHAPTER 3: ADC AND DAC

1. Specify how many bits are needed to appropriately digitize each of the following signals. Choose from: 6 bits, 8 bits, 10 bits, 12 bits, 14 bits, or 16 bits.

    a. A signal where the maximum amplitude is 1 volt and the rms noise is 1.5 millivolts.

    b. A signal with a signal-to-noise ratio of 900 to 1.

    c. A signal with a coefficient-of-variation of 0.4%.

    d. A high-fidelity audio system (hint: a jack-hammer is about 50,000 times louder than a pin drop).

    e. A black and white digital image (hint: under the best conditions, the human eye can differentiate about 200 shades of gray between pure black and pure white).

2. A scientist evaluates two different digital thermometers. Each has a digital read-out to the nearest one degree, and updates once each second. When the temperature is held constant, the digital display of thermometer A does not change, but the display of thermometer B randomly toggles between three to four adjacent readings (i.e., +/- 2 degrees). For the questions below, assume that the temperature is held constant.

    a. What is the largest possible error in an individual reading from A?

    b. What is the largest possible error in an individual reading from B?

    c. In a single reading, which provides the "best" information? Explain.

    d. If 1000 readings are taken with A, what is the standard deviation?

    e. If 1000 readings are taken with B, what is the approximate standard deviation? (choose from 0.1, 0.5, 2.0 and 4.0).

    f. If 1000 readings are taken with thermometer B, what is the "typical error between the mean of the readings, and the mean of the underlying process?

    g. If 1000 readings are taken with each thermometer, which data set provides the "best" information? Explain.

    h. How many readings must be taken with thermometer A to reliably detect a temperature change of 0.15 degrees?
    (choose from 10, 100, 1000, 1 million, 1 billion, or "it cannot be done"). Explain.

    i. Repeat (h) for thermometer B.

3. An engineer designs a microprocessor controlled ADC board that can acquire an 8 bit sample every 10 microseconds. His boss walks in and says: "You'll get a raise if the system can be modified to acquire a 12 bit sample every 100 milliseconds- but it needs to be done by tomorrow!" The first thing the engineer does is to measure the noise on the analog signal entering the ADC chip. He then smiles, and plans how to spend the extra money. In detail, explain what the engineer was looking for in the noise measurement, and how he can make the modification.

4. An analog electronic signal is composed of three sine waves: 1 kHz @ 1 volt amplitude, 3 kHz @ 2 volts amplitude, and 4 kHz @ 5 volts amplitude (all voltage readings are peak-to-peak). The signal is digitized with 12 bits, spread over the range of -5 volts to +5 volts. For each sampling rate below, describe the frequency components that exist in the digital signal. Be sure to specify three things for each component: its digital frequency (a number between 0 and 0.5), its amplitude (in digital numbers, peak-to-peak), and its phase relative to the original analog signal (either 0 degrees or 180 degrees).
   a. Sampling rate = 100 kHz.
   b. Sampling rate = 10 kHz.
   c. Sampling rate = 7.5 kHz.
   d. Sampling rate = 5.5 kHz.
   e. Sampling rate = 5 kHz
   f. Sampling rate = 1.7 kHz.
5. A high-fidelity audio signal (containing frequencies between 20 Hz and 20 kHz) is contaminated with interference from a nearby switching power supply, operating at 32 kHz. To eliminate the interference, the analog signal is passed through an 8 pole Butterworth filter with a cutoff frequency of 24 kHz. The filtered signal is then digitized at 44 kHz.
   a. Sketch and label the frequency spectrum of the analog signal, showing: the audio frequency band, the interfering signal, the frequency response of the filter, one-half the sampling frequency, and the sampling frequency.
   b. What is the approximate attenuation of the filter at the frequency of the interference?
   c. What is the effect of the filter on the audio information?
   d. Sketch and label the frequency spectrum of the digital signal, showing: the audio frequency band, and the interfering signal.
   e. At what frequency does the aliased interference appear in the digitized signal?
   f. If an 8 pole Chebyshev filter (6% ripple) were used instead, how would the amplitude of the aliased interference change?
   g. Repeat (a)-(f) with the cutoff frequency of the filter set at 20 kHz.
6. A multirate technique is used to handle the interference in the last problem. The original analog signal (with interference) is digitized at a sampling rate of 176 kHz. A low-pass digital filter then removes all digital frequencies above 0.18 with less than a 0.02% residue, while passing all frequencies below 0.12 with less than a 0.02% passband ripple (an easy task for a digital filter). The digital signal is then resampled at 44 kHz, that is, every three out of four samples are discarded.
   a. Sketch and label the frequency spectrum of the digital signal before filtering, showing: the audio frequency band, the interfering signal, and the approximate frequency response of the digital filter.

b. Has aliasing occurred during the sampling? Explain.

c. In a fair comparison, should this digital filter be compared against the analog Butterworth filter, or the analog Chebyshev filter? Explain.

d. How much better is the digital filter performing compared to the analog filter you indicated in (c)? (It can be difficult to compare analog and digital filters; make some sort of quantitative comparison, and explain your method).

e. Sketch and label the frequency spectrum of the digital signal after resampling, showing: the audio frequency band, and any interfering signals.

7. On television, rotating objects such as wagon wheels and airplane propellers often appear to be moving very slowly or even backwards. This is a result of aliasing, caused by the sampling rate of the video (30 frames per second) being less than twice the frequency of the rotational motion. To understand this, imagine we paint one of the blades of an airplane propeller so that we can identify it from the other blades. We will then turn the propeller at 33 rotations per second, in a clockwise direction. In frame number 1 of our video sequence, the marked blade happens to be exactly at the top of the propeller.

   a. How many rotations does the marked blade make between two successive frames?

   b. Draw a sketch of how the propeller would appear in frames 1, 2, 3 and 4.

   c. How many frames does it take for the marked blade to again appear at the top?

   d. What rotational frequency is (c) in rotations per second?

   e. Is this apparent rotation clockwise or counterclockwise?

   f. Explain using Fig. 3-4 how the marked blade's actual frequency, the frame rate, and the marked blade's observed frequency are related.

   g. Repeat (a) to (f) when the propeller is turning at 57 rotations per second.

8. When viewed on television at 30 frames per second, what is the apparent rotational rate of a 4 blade propeller, turning at 44.7 rotations per second, if all the blades are identical? Is the apparent rotation clockwise or counterclockwise?

CHAPTER 4: DSP SOFTWARE

1. The following subroutine is used to calculate the function: y = exp(x), using an efficient implementation of Eq. 4-3:

```
y = 1
term = 1
for counter = 1 TO 150
 term = term * x / counter
 y = y + term
next counter
```

   a. What are the values of "y" and "term" at the end of the first, second, and third loops?
   b. At the end of all 150 loops, how many terms have been included in the calculation?

2. Calculate the numeric value of the first 14 terms of the series for exp(x) (Eq. 4-3), where x = 1.3. From this data, how many terms must be used to achieve an accuracy of one part in one-million?

3. To illustrate the size of quantization levels, imagine representing the heights of two buildings as digital numbers. Building A is exactly 100 meters in height, while Building B is 100.0001 meters. That is, Building B is about the thickness of a sheet of paper (0.1 mm) higher than building A. Indicate whether or not each of the following types of digitized numbers could show that the two buildings are different in height.
   a. Integer (8 bit, unsigned)
   b. Integer (16 bit, 2's complement)
   c. Single precision floating point
   d. Double precision floating point

4. Repeat problem 3 for Building B being one ten-millionth of a meter higher than Building A (about the diameter of a single atom). If double precision were used, how much smaller than the diameter of an atom are the quantization levels?

5. Find the decimal number that corresponds to the following floating point bit patterns.
   a. 10111000101010100000000000000000
   b. 10000000000000000000000000000000
   c. 01111001111110010000000000000000
   d. 01111111100000000000000000000000

6. Convert the following decimal numbers into their IEEE floating point bit patterns:
   a. 1
   b. 2
   c. 4
   d. -5
   e. 18

7. Imagine you are trying to represent the number: 4.0000003 in single precision.
   a. What bit pattern corresponds to the number 4?
   b. What bit pattern corresponds to the next largest number that can be represented?
   c. What decimal number corresponds to the bit pattern in (b)?
   d. Which of the above two binary patterns should be used to represent the number: 4.0000003? Why?
   e. What is the error introduced when this number is stored in single precision? Express your answer both as an absolute number, and as a fraction of the number being represented.
8. In an FIR digital filter, each sample in the output signal is found by multiplying M samples from the input signal by M predetermined coefficients, and adding the products. The characteristics of these filters (high-pass, low-pass, etc.) are determined by the coefficients used. For this problem, assume M = 5000, and that single precision floating point math is used.
   a. How many math operations (the number of multiplications plus the number of additions) need to be conducted to calculate each point in the output signal?
   b. If the output signal has an average amplitude of about one-hundred, what is the expected error on an individual output sample? Assume that the round-off errors combine by addition. Give your answer both as an absolute number, and as a fraction of the number being represented.
   c. Repeat (b) for the case that the round-off errors combine randomly.

CHAPTER 5: LINEAR SYSTEMS

1. Two discrete waveforms, x[n] and y[n], are each eight samples long, given by:

   x[n]: 1, 2, 3, 4,-4,-3,-2,-1
   y[n]: 0,-1, 0, 1, 0,-1, 0, 1

   For this problem, you can add additional samples with a value of zero on either side of the signals, as needed. Calculate, sketch and label the following signals:
   a. x[n]
   b. y[n]
   c. 5x[n]
   d. -7y[n]
   e. x[n-3]
   f. y[n+1]
   g. 2x[n+1]
   h. -y[n-1]
   i. x[n] + y[n]
   j. -2x[n-1] + 3y[n+2]
   k. 3x[n+2] - 2y[n+2]
   l. x[n] + x[n-2]
   m. 2x[n] -3x[n-2] + 3y[n+1]

2. Sketch the following discrete signals for -8 < n < 8:
   a. x[n] = sin(2 pi n / 8)
   b. x[n] = cos(2 pi n / 4)
   c. x[n] = sin(2 pi n / 2)
   d. x[n] = cos(2 pi n / 2)
   e. x[n] = n-3 if n > 2, 0 otherwise
   f. x[n] = 1 if n < -3, 0 if 0 < n < 4, 5 otherwise

3. Sketch the following continuous signals for -8 < t < 8:
   a. x(t) = sin(2 pi t / 8)
   b. x(t) = cos(2 pi t / 4)
   c. x(t) = sin(2 pi t / 2)
   d. x(t) = cos(2 pi t / 2)
   e. x(t) = n-3 if t > 2, 0 otherwise
   f. x(t) = 1 if t < -3, 0 if 0 < t < 4, 5 otherwise

4. Samples 0 to 11 of a signal have the value: 0, 2, 3, 4, 3, 2,-1, 0, -2, -3, 2, 1.
   Calculate, sketch and label:
   a. The even and odd parts.
   b. The interlaced decomposition.
   c. The step decomposition.

5. Two continuous waveforms, b(t) and x(t) are defined by:

   $b(t) = 1$ for $0 < t < 2$, 0 otherwise

   $x(t) = -1$ for $1 < t < 2$, 1 for $2 < t < 3$, 4 for $3 < t < 4$, 2 for $4 < t < 5$, 0 otherwise

   a. Sketch b(t) and x(t)
   b. Show that x(t) can be decomposed into three scaled and shifted versions of b(t). That is, find: a1, a2, a3, s1, s2, s3, such that: $x(t) = a1\ b(t-s1) + a2\ b(t-s2) + a3\ b(t-s3)$
   c. Sketch these three component signals.

6. Systems are proven to be linear by mathematically showing that they obey the properties of additivity and homogeneity. However, systems in the real world are often only understood by empirical measurements. That is, a scientist or engineer places a test signal into the input, and looks at what comes out.

   a. Is it possible to prove that a system is linear based on measurements of the input and output alone, without knowing the formal mathematical relationship between the input and output? Explain.

   b. Is it possible to prove that a system is nonlinear in this way? Explain.

   To help you answer these questions, think about an electronics technician testing a "black box" for being linear. He does this by placing signals into the input and observing the output. However, the technician has absolutely no information about what is "inside" the system. For instance, it might contain an evil demon trying to deceive the technician. Or in another case, it might contain a timer that scrambles the output once every ten million years of operation.

CHAPTER 6: CONVOLUTION

1. A system has an impulse response, h[n], given by: 1, 2, 2, 1, 0, -1, 0, 0, for the values of samples 0 to 7. Calculate the output of the system in response to the following input signals.
   a. 1, 0, 0, 0, 0, 0, 0, 0, 0
   b. -3, 0, 0, 0, 0, 0, 0, 0, 0
   c. 0, 0, 1, 0, 0, 0, 0, 0, 0
   d. 1, 0, 1, 0, 0, 0, 0, 1, 0
   e. 3, 0,-1, 0, 0, 2, 0, 0, 0
   f. 2,-1, 0, 0, 1, 0,-1, 0, 0

2. Adding zeros to the end of a signal is a common DSP technique called "padding with zeros." Use your results in the last problem to answer the following:

   a. How would the output signals be changed if 5 additional samples, all with a value of zero, were added to the end of the impulse response?

   b. How would the output signals be changed if 5 additional samples, all with a value of zero, were added to the end of the input signals?

   c. How would the output signals be changed if 5 additional samples, all with a value of zero, were added to the end of both the input signals and the impulse response?

   d. Complete the following statement summarizing how "padding with zeros" affects convolution: "When M zeros are added to either the input signal or the impulse response, the only change to the output signal is [fill in the blank].

3. Two signals, x[n] and h[n], are defined by:

   x[n]: 1, 0, 2, 3, 2, 1,-1,-2,-1, 0, 2, 3, 3, 2, 1, 1 (samples 0-15)
   h[n]: 1, 2, 3,-3,-2,-1 (samples 0-5)

   If y[n] = x[n]*h[n], use the input side algorithm to determine the contribution to y[n] from the indicated sample:
   a. x[2]
   b. x[6]
   c. x[9]

4. For the signals in the last problem, use the output side algorithm to calculate the value of the following samples:
   a. y[8]
   b. y[10]
   c. y[3]
   d. y[18]

5. Two signals, a[n] and b[n], are defined by:
   a[n]: 1, 0, 0, 2, 1, 0
   b[n]: 0,-1,-2, 0, 0, 1
   a. Calculate a[n]*b[n] by using an impulse decomposition on a[n], convolving each of the components with b[n], and synthesizing (adding) the resulting signals.
   b. Calculate a[n]*b[n] by using an impulse decomposition on b[n], convolving each of the components with a[n], and synthesizing (adding) the resulting signals.
   c. Do the results of these two methods agree? What property is demonstrated in this problem?
6. Calculate the convolution of the signal: h[n] = 1, 2, 3, 0, 0 with the indicated signals (assume each of the following run from sample 0 to 7).
   a. x[n] = delta[n]
   b. x[n] = -5 delta[n-2]
   c. x[n] = 2 delta[n+1] - delta[n+1]
   d. x[n] = 1, 2, 3, 0, 0 ...
   e. x[n] = -n for 0 < n < 5, and 0 otherwise
   f. x[n] = 2^(-n)
   g. x[n] = sin(2 π n)
   h. x[n] = cos(2 π n)
   i. x[n] = sin (π n)
7. Calculate the convolution of the following signals (your answer will be in the form of an equation):
   a. h[n] = delta[n], x[n] = delta[n]
   b. h[n] = delta[n], x[n] = delta[n-k]
   c. h[n] = delta[n-2], x[n] = delta[n-1] + delta[n+4]
   d. h[n] = delta[n-1] + delta[n+1], x[n] = delta[n-a] + delta[n+b]
   e. h[n] = delta[n], x[n] = exp(-n)
   f. h[n] = delta[n+2], x[n] = exp(n)
   g. h[n] = delta[n-2], x[n] = exp(-n)
   h. h[n] = exp(-n), x[n] = delta[n-2]
   i. h[n] = delta[n] - delta[n-1], x[n] = exp(-n)
8. A financial expert receives daily reports on the value of a particular stock. Each day he calculates the average value of the stock over the last 30 days. If this averaging were describe as a system:
   a. What are the input and output signals?
   b. Is this system linear?
   c. What is the impulse response of the system?
   d. What practical purpose would this system be serving?
   e. What would be the impulse response if the average was taken over M days?
9. If the signal, x[n], has a value of zero over the interval: A <= n <= B, and if signal, h[n], has a value of zero over the interval: C <= n <= D, then x[n]*h[n] must be zero over the interval, E <= n <= F. Express the variables, E and F, in terms of: A, B, C, and D.

10. Two signals, a[n] and b[n], each contain 6 points, as defined below.
Calculate a[n]*b[n]. (Where the "*" denotes convolution).

    a[n]: 1, 0, 0, 2, 1, 0
    b[n]: 0, -1, -2, 0, 0, 1

    a. Where both signals run from sample 0 to 5
    b. Where both signals run from sample 2 to 7
    c. Where a[n] runs from sample 0 to 5, and b[n] runs from sample -3 to 2
    d. Where a[n] runs from sample -10 to -5, and b[n] runs from sample -5 to 0

CHAPTER 7: PROPERTIES OF CONVOLUTION

1. Classify the following signals as either casual or noncausal.
    a. x[n] = delta[n]
    b. x[n] = delta[n-2]
    c. x[n] = delta[n-1] + delta[n+1]
    d. x[n] = delta[n] - 5 delta[n-5]
    e. x[n] = delta[n] + delta[n+5]
    f. x[n] = delta[n-1] - delta[n-4] + delta[n-7]
    g. x[n] = exp(-n)
    h. x[n] = exp(-abs(n)) (where "abs" is the absolute value function)
    i. x[n] = abs(n)
    j. x[n] = n + abs(n)
2. Classify the signals in the last problem as either zero phase, linear phase, or nonlinear phase.

3. The impulse responses of three linear systems are given below. Calculate the impulse response of the indicated combination

system A: 3, 2, 1, 0
system B: 0, 1,-1, 0
system C: 1, 1, 1, 1

a. The parallel combination of system A and system B.
b. The parallel combination of system A, system B, and system C.
c. The cascade of System A and system B.
d. The cascade of System B and system A.
e. The cascade of System A and system B, in parallel with system C.

4. System A is an "all pass" system, meaning that its output is identical to its input. System B is a low-pass filter that passes all frequencies below the cutoff frequency without change, and blocks all frequencies above. Call the impulse response of system B, b[n].

a. What is the impulse response of system A?
b. How would the impulse response of system B need to be changed to make the system have an inverted output (i.e., the same output, just changed in sign)?
c. If the two systems are arranged in parallel with added outputs, what is the impulse response of the combination?
d. If the two systems are arranged in parallel, with the output of system B subtracted from the output of system A, what is the impulse response of the

combination?
e. What kind of filter is the system in (d)?
f. Describe an algorithm for changing a low-pass filter kernel into a high-pass filter kernel.
g. Describe an algorithm for changing a high-pass filter kernel into a low-pass filter kernel. Is this exactly the same procedure as in (f)?
h. In this problem, system B has the ideal characteristic of passing certain frequencies "without change." How would the algorithm you described in (f) be affected if the low-pass filter delayed (i.e., shifted) the output signal by a small amount, relative to the input signal?

5. From calculus, you know that the derivative and integral are inverse operations; one undoes the effect of the other. Prove that the first difference and the running sum are also inverse operations. That is, show that the cascade of these two systems is identical to the delta function.

6. Echoes are added to audio signals to make the listener "feel" that they are in a particular size of room. Assume that an audio signal is sampled at 44 kHz, and that sound propagates at 332 meters/second. In a "small" room, a person stands about 3 meters from the walls; in a "large" room, the distance increases to about 10 meters.

a. In a small room, how long is the delay between a person making a sound and its echo from the walls.
b. How many samples does this correspond to in the digital signal?
c. What is the impulse response of a digital system simulating this echo, if the amplitude of the echo is 20%?
d. Repeat (a) to (c) for the large room.
e. In a real listening environment, each echo will also generate another echo. That is, each original sound will be heard over and over with diminishing amplitude. How would the impulse response in (c) be modified to account for these echoes of echoes?

CHAPTER 8: THE DISCRETE FOURIER TRANSFORM

1. A sinusoid at 1.7 kHz is digitized at 10,000 samples per second. The signal is passed through a 2048 point DFT, and converted to polar form. Draw four sketches of the magnitude, one for each of the four ways that the frequency domain's independent variable can be expressed. Be sure to indicate the frequency symbol used, the range of values, the units, and at what frequency the sinusoid appears.

2. A peak appears at index number 19 when a 256 point DFT is taken of a signal.

   a. What is the frequency of the peak expressed as a fraction of the sampling rate? Do you need to know the actual sampling rate to answer this question?
   b. What is the frequency of the peak expressed as a natural frequency?
   c. What is the sampling rate if the peak corresponds to 21.5 kHz in the analog signal?
   d. What is the frequency of the sinusoid (in hertz) if the sampling rate is 100 kHz?

3. Calculate, sketch and label the basis functions for an 8 point DFT.

4. An 8 sample signal is given by: 20,21,22,23,24,25,26,27. Its frequency domain is given by: R0,R1,R2,R3,R4 and I0,I1,I2,I3,I4. Write the simultaneous equations relating the frequency domain and the time domain. Use the numerical value for each point in the basis functions, for example, use 0.7071, not sin(pi/4). How difficult would it be to solve these equations?

5. Using the signals in the last problem, write 10 equations for calculating the 10 points in the frequency domain, using the correlation algorithm. Solve these equations.

6. The frequency domain of a signal is given by:

   real part: 1, 2, 3, 3, 1,-2,-1, 1, 2
   imag part: 0,-1,-2, 0, 0, 0, 2, 1, 0

   a. What length of DFT does this correspond to?
   b. Calculate the amplitudes of the sine and cosine waves that comprise the time domain signal.
   c. What is the mean (average value) of the time domain signal?

7. You are told that the following signals are the frequency domain of a 32 point real DFT. Give two reasons why this is not possible.

   real part: 1,2,3,4,5,6,7,8,7,6,5,4,3,2,1,0

   imag part: 8,7,6,5,4,3,2,1,0,1,2,3,4,5,6,7

8. Convert the following real and imaginary parts into polar form. Sketch a diagram of each, such as Fig. 8-9. In each case, state if the conversion equation: phase = arctan(IP/RP), provides the correct answer without additional steps:

   a. RP = 1, IP = 1

   b. RP = 1, IP = -1

   c. RP = -1, IP = 1

   d. RP = -1, IP = -1

   e. RP = 1, IP = 0

   f. RP = -1, IP = 0

   g. RP = 0, IP = 1

   h. RP = 0, IP = -1

CHAPTER 9: APPLICATIONS OF THE DFT

1. There are three signals involved in linear systems: the input signal (containing information we want to understand), the impulse response (controlling how the information is modified), and the output signal (a result of the other two signals). It is not a coincidence that the DFT has three main uses. Match each DFT techniques with its corresponding signal. Explain how each DFT technique provides a way of understood or dealing with the associated signal.

2. A scientist acquires 65,536 samples from an experiment at a sampling rate of 1 MHz. He knows that the signal contains a sinusoid at 100 kHz. He needs to determine is if there is a second sinusoid at 103 kHz. As a start, he takes a 65,536 point DFT of the signal. To his surprise, all he can see in the spectrum is noise. He estimates that the signal he is looking for is 4 times lower in amplitude than the random noise (i.e., SNR = 0.25). He also estimates that the SNR will need to be at least 3.0 for the signal to be detected, if present. To improve the SNR, he decides to break the signal into segments, and average their spectra. Arrange your answers to the following questions in a table.

   a. If he uses a segment length of 16,384 samples, what will be the frequency resolution (i.e., the spacing between data points) in the averaged spectrum? Give your answer in hertz. How many segments will be averaged? What is the SNR of the averaged spectrum? Does this have the required frequency resolution? Does this have the required SNR?

   b. Repeat for segment lengths of: 8192, 4096, 2048, 1024, 512, 256, and 128.

3. If the scientist in the last problem wanted to improve this experiment, what advice could you give?

4. A filter kernel (impulse response) consists of 250 samples, and is designed to pass all frequencies below 0.11, and block all frequencies above 0.12. To evaluate how this filter performs, you want to closely inspect its frequency response over this range. To do this, you pad the impulse response with zeros to make the total length 256 samples, and then take the DFT.
   a. How may data points are spread over the range of interest?
   b. Repeat using a DFT length of 2048.
   c. Repeat using a DFT length of 2 to the 50th power.
   d. Is there anything limiting how many samples can be placed over this range of interest? How does this relate to the DTFT? Explain.

5. A signal containing 1000 points is to be convolved with a signal containing 128 points.
   a. What is the length of the resulting signal?
   b. If frequency domain convolution is used, what length of DFT is appropriate?
   c. If a 1024 point DFT is used, how many samples are correct, and how many are corrupted by circular convolution?
   d. Repeat (a) to (c) for the two signals having 490 samples and 23 samples.

CHAPTER 10: PROPERTIES OF THE DFT

1. If x[n] has the frequency domain: Xreal[f] and Ximag[f], and y[n] has the frequency domain: Yreal[f] and Yimag[f], calculate the frequency domain of the following signals:
   a. x[n]/5
   b. 5.5y[n]
   c. x[n] + y[n]
   d. 3.14x[n] + y[n]/3.14
   e. ax[n] + by[n], where a and b are constants
2. If x[n] has the frequency domain: Xmag[f] and Xphase[f], and y[n] has the frequency domain: Ymag[f] and Yphase[f], calculate the frequency domain of the following signals:
   a. x[n-2]
   b. 1.2x[n-1]
   c. y[n+2]/10
   d. ax[n-b], where a and b are constants
3. Regarding the signals in problems 1 and 2:
   a. In problem 1, why would it be difficult to calculate the frequency domain of: x[n-2]?
   b. In problem 2, why would it be difficult to calculate the frequency domainof x[n] + y[n]?
   c. Complete the following statements describing this situation:
      i. "When time domain signals are added, the frequency domain is easiest to understand when in [fill in the blank] form."
      ii. "When time domain signals are shifted, the frequency domain is easiest to understand when in [fill in the blank] form."
      iii. "When time domain signals are scaled, the frequency domain is easiest to understand when in [fill in the blank] form."
4. Suppose that the frequency spectrum in Fig 9-2 represents a digital signal with a sampling rate of 160 Hz. Sketch the digital frequency spectrum for frequencies between -2.0 and 2.0.
5. A digital low-pass filter passes all frequencies below 0.1, and blocks all frequencies above.
   a. Sketch this frequency response, showing all frequencies between -1.5 and 1.5.
   b. If the filter kernel is multiplied by a sine wave with a frequency of 0.3, sketch the new frequency response.
   c. Based on the method in (b), describe an algorithm for converting a low-pass filter with cutoff frequency, fc, into a bandpass filter with a center frequency, fcenter, and a bandwidth, BW.
   d. If the low-pass filter kernel is multiplied by a cosine wave with a frequency of 0.5, sketch the new frequency response.
   e. Based on the method in (d), describe an algorithm for converting a low-pass filter with cutoff frequency, flp, into a high-pass filter with a cutoff frequency, fhp.
   f. Why must a cosine wave be used in (e) instead of a sine wave?

6. To represent the sound of a human voice, a signal only needs to contain frequencies between 100 Hz and 4 kHz (Land line phones only reproduce 300 Hz to 3300 Hz). In comparison, high-fidelity music must contain all the frequencies that humans can hear, i.e., 20 Hz to 20 kHz (young, good ears).
   a. Sketch and label the frequency spectra of these two signals, including the negative frequencies. Assume all the frequencies have the same amplitude.
   b. A DC bias is often added to audio signals in electronic circuits. This is to make the voltage representing the signal always have a positive value (see Fig. 10-14a). Repeat (a) assuming that each audio signal has such a DC bias.
   c. Sketch the frequency spectrum of a voice signal multiplied by a sine wave at 1 MHz. Repeat for a voice signal plus DC bias.
   d. Sketch the frequency spectrum of a music signal multiplied by a sine wave at 1.1 MHz. Repeat for a music signal plus DC bias.
   e. If the signal in (c) is transmitted from a radio station, what band of frequencies must be assigned by the government to avoid interference with other radio applications? Repeat for the signal in (d).
   f. If a frequency band of 50.1 to 50.2 MHz were available, how many voice signals could be simultaneously transmitted? How many high-fidelity music signals? Explain how the signals would be prepared for this simultaneous transmission.

7. The "no bias" signal in (c) of the last problem is filtered to remove all frequencies below 1 MHz. The result is called "single-sideband" (SSB) modulation.
   a. Sketch the single-sideband frequency spectrum.
   b. Does this contain the same information as the original voice signal?
   c. A disadvantage of SSB is that it requires more complicated modulation and demodulation electronics. What would be the advantage of using single-sideband modulation over AM?
   d. If the lower sideband were retained instead, would the information in the original signal still be preserved? Explain.

CHAPTER 11: FOURIER TRANSFORM PAIRS

1. Give the mathematical equation for the frequency domain corresponding to the following waveforms. You can state your answer in either rectangular or polar form. Example: x[n] = delta[n], answer: Mag X[f] = 1, Phase X[f] = 0.

   a. x[n] = 2delta[n-2]
   b. x[n] = sin(2 pi n 14 / 256)
   c. x[n] = cos(2 pi n 0.2)
   d. x[n] = 1 for 10 < n < 18, 0 otherwise
   e. The signal in (d) convolved with itself
   f. A Gaussian centered at sample 100, with a standard deviation of 20.
   g. The signal in (f) multiplied by a cosine wave of frequency 0.3.

2. Those experienced in DSP and electronics can approximate the highest frequency contained in a signal simply by look at its graph. This problem will help you master this useful skill. In general, the most rapidly changing sections of a signal will correspond to the highest frequency contained in the signal. As an example, the sinc function in Fig. 11-5a has a period of about 11 samples. This corresponds to the highest frequency present in the signal being about 1/11 = 0.09. In a more realistic example, Figs. (c) and (d) show waveforms that resemble one-half cycle of a sinusoid being completed over 16 cycles. This corresponds to one cycle every 32 samples, or a highest frequency of 1/32 = 0.03. Use this method to estimate the highest frequency component in the following signals:

   a. Fig. 7-14, y[n]
   b. Fig. 8-8b
   c. Fig. 9-7a
   d. Fig. 11-6b-d
   e. Fig. 3-5a (for "time" in seconds)
   f. Fig. 10-14a
   g. Fig. 13-8a
   h. A signal where each sample is obtained from a random number generator.
   i. In some of the above figures, the actual frequency spectra are also given. Using this information, approximately how accurate is this method?
   (Choose from 1%, 3%,10%, or 30%)

3. A signal contains 16 consecutive samples with a value of 1, with the remainder of the samples having a value of zero.

   a. What are the frequencies of the first four zeros crossings in the frequency domain?
   b. What are the periods of these frequencies?
   c. Make a sketch showing how sinusoids at the first two zero crossings fit evenly inside the rectangular pulse.

4. A discrete sinusoid of frequency 0.125 is half-wave rectified (that is, all the negative valued samples are set to zero).

    a. Sketch the original time domain signal and its frequency spectrum.
    b. Sketch the rectified time domain signal and its frequency spectrum (don't worry about the exact amplitude of the harmonics).
    c. Which harmonics of the rectified signal will be aliased?
    d. At what frequency does the 5th harmonic appear?
    e. At what frequency does the 10th harmonic appear?
    f. At what frequency does the 100th harmonic appear?

5. Referring to the chirp system in Fig. 11-10:

    a. What is the frequency domain magnitude of the signal in (a)? Of the signal in (b)?
    b. What is the frequency domain phase of the signal in (a)? Of the signal in (b)? (express these answers as equations using the constants alpha and beta).
    c. Using Parseval's relation, what is the relationship between the total energy contained in the waveforms in (a) and (b)? Explain.
    d. Recall from physics that power is equal to the total energy released divided by the time over which it is released. If the waveform in (a) is one sample long, and the waveform in (b) is 80 samples long, what power reduction is provided by this chirp system?

6. Figures 11-5 (a) and (e) illustrate a sinc function and a Gaussian, respectively. Both these functions decrease in amplitude toward the left and right. However, neither of these functions ever reach a constant value of zero.

    a. Write an equation describing how the amplitude of the sinc's oscillations decrease, moving away from the center of symmetry.
    b. Repeat (a) for the Gaussian (estimating the standard deviation from the figure).
    c. When the main lobes are about the same width [as illustrated in Figs. (a) and (e)], which of these two functions drops toward zero faster? To answer this, evaluate the amplitude of the two functions at 3, 10, 30, 100, and 300 samples from the center of symmetry.
    d. For each function, how many samples must you move away from the center of symmetry before the amplitude fall below the single precision round-off noise?
    e. Would a Gaussian or a rectangular pulse in the frequency domain be more likely to result in time domain aliasing? Explain.
    f. Would a Gaussian or a rectangular pulse in the time domain be more likely to result in frequency domain aliasing? Explain.

CHAPTER 12: THE FAST FOURIER TRANSFORM

1. The following spectrum was produced by the real DFT. Generate the frequency spectrum of the corresponding complex DFT.

   samples 0 to 8 of the real part: 1, 2,-1,-2, 0, 1, 2, 3, 2
   samples 0 to 8 of the imaginary part: 0, 2, 4, 5,-3,-2, 1, 1, 0

2. A time domain signal, consisting of a real part, rex[n], and an imaginary part, imx[n], has the following complex spectrum:

   samples 0 to 7 of the real part: 1, 2,-1,-2, 1, 0, 2, 3
   samples 0 to 7 of the imaginary part: 3, 2, 4, 5,-1,-2, 1, 1

   a. Separate this spectrum (both the real and imaginary parts) into even and odd parts.
   b. What would be the spectrum if the values in imx[n] were set to zero?
   c. What would be the spectrum if the values in rex[n] were set to zero?

3. Suppose you want to conduct a spectral analysis of a signal containing 1,003,520 samples, and the computer you are using has values of: Kdft = 1 microsecond, and Kfft = 1.5 microsecond.

   a. Calculate the execution time if the signal is broken into 64 sample segments, and the DFT by correlation algorithm is used. Repeat for segment lengths of: 256, 1024, and 4096. (Ignore the calculation time for averaging the frequency spectra).

   b. Repeat (a) using the FFT algorithm.

CHAPTER 13: CONTINUOUS SIGNAL PROCESSING

1. How short must a pulse be to act as an impulse to the following systems? (Give an "order-of- magnitude" approximation and justify your reasoning).

    a. A high-fidelity music system, designed to reproduce sounds between 20 Hz and 20 kHz.
    b. An automotive suspension (think about how fast it recovers after the vehicle hits a bump).
    c. An 8 pole Bessel low-pass filter with a cutoff frequency of 1 kHz. (hint: see Fig. 3-13)
    d. The disruption of a galaxy when struck by another galaxy (see the picture on the cover of the book, and the caption opposite the title page). Hint: The "time constant" of disruption is the time it takes the expanding gas to equal the diameter of the galaxy.

2. Calculate and sketch the convolution of a(t) and b(t). Also provide sketches showing how each region in the output signal is calculated.

    a. a(t) = 1 for 0 < t < 2, and 0 otherwise
    b(t) = 1 for 0 < t < 2, and 0 otherwise

    b. a(t) = 2 for 0 < t < 1, and 0 otherwise
    b(t) = 1 for 2 < t < 5, and 0 otherwise

    c. a(t) = 1 for 0 < t < 1, and 0 otherwise
    b(t) = t for 0 < t < 3, and otherwise

    d. a(t) = exp(-kt) for 0 < t, and 0 otherwise, and k is a constant
    b(t) = exp(-kt) for 0 < t, and 0 otherwise

    e. a(t) = exp(-kt) for 0 < t, and 0 otherwise
    b(t) = -2 delta(t-2)

3. Convolve the following signals by differentiating one to reduce it to impulses (such as in Fig. 13-7).

    a(t) = 1 for 0 < t < 4, and 0 otherwise
    b(t) = sin(2 pi t) for -1 < t < 1, and 0 otherwise

4. Calculate the Fourier transform of the following signals. Give your answer in rectangular form.

    a. x(t) = 1 for -1 < t < 1, and 0 otherwise.
    b. x(t) = t for -1 < t < 1, and 0 otherwise
    c. x(t) = delta(t)
    d. x(t) = exp(-abs(kt)), where abs is the absolute value and k is a constant

5. Which of the following signals has a larger 15th harmonic: a square wave, a triangle wave, or a sawtooth wave?

CHAPTER 14: INTRODUCTION TO DIGITAL FILTERS

1. Digital Filters

a. You can characterize a digital filter by:
   i. It's unit pulse response
   ii. It's complex frequency response
   iii. It's unit step response
   iv. All of the above

b. The inverse discrete Fourier transform of the filter complex frequency response is it's unit pulse response.
   i. True
   ii. False

c. A Finite Impulse Response filter (FIR) transfer function has:
   i. Only zeros
   ii. Only poles
   iii. Both zeros and poles

d. A Infinite Impulse Response filter (IIR) transfer function has:
   i. Only zeros
   ii. Only poles
   iii. Both zeros and poles

e. An "Ideal" low pass filter is not realizable since:
   i. The sharp transition of the magnitude response can only be approximated
   ii. It requires a time response that starts before the input occurs
   iii. Both of the above

f. A "Band-Pass" filter passes all frequencies within it's pass-band with equal amplitude and totally blocks all frequencies outside of the pass band.
   i. True
   ii. False

g. Pass-Band ripple describes the maximum dB variation in amplitude response in the pass band.
   i. True
   ii. False

h. Stop-Band attenuation is uniform for all frequencies outside of the pass band.
   i. True
   ii. False

i. The Transition Band is defined as the frequency range over which the frequency response goes from the specified maximum pass band attenuation to the minimum stop band attenuation.
    i. True
    ii. False

2. FIR Filters

a. The transient response of an FIR filter continues forever
    i. True
    ii. False
b. FIR filters are always stable
    i. True
    ii. False
c. The signal delay through an FIR filter is half of the filter length
    i. True
    ii. False
d. An FIR filter has an "All Zeros" transfer function in the z-domain.
    i. True
    ii. False
e. An FIR filter always operates on a "window" on the input data stream.
    i. True
    ii. False
f. An FIR filter has both "Feed Forward" and Feedback paths in its signal flow diagram
    i. True
    ii. False

3. IIR Filters
    a. The transient response of an IIR filter continues forever
        i. True
        ii. False
    b. IIR filters are always stable
        i. True
        ii. False
    c. An IIR filter has an “All Zeros” transfer function in the z-domain.
        i. True
        ii. False
    d. An IIR filter always operates on a “window” on the input data stream.
        i. True
        ii. False
    e. An IIR filter has both “Feed Forward” and Feedback paths in its signal flow diagram
        i. True
        ii. False

CHAPTER 15: MOVING AVERAGE FILTERS

1. The transient response (rise time due to a step input) of an n-point moving average filter is _____ samples.
   a. 0
   b. n/2
   c. n
   d. 2n
2. Moving Average filters are used in technical analysis of the stock market
   i. True
   ii. False
3. Moving Average filters are FIR filters whose pulse response terms are all 1's
   i. True
   ii. False
4. You can apply a Window function to improve the frequency response of a moving average filter, but:
   a. The processing will be more complex
   b. Noise in the output signal will increase
   c. The effective length of the filter will be reduced
   d. All of the above

CHAPTER 16: WINDOWED-SINC FILTERS

1. Computer Experiment: Rectangular Windows
   For the lengths N = 11, 41, 81, and 121 design a lowpass FIR filter with cutoff frequency $\omega_c = 0.3\pi$ using rectangular windows.
   a. Plot the impulse response, h(n), and the magnitude response, $|H(\omega)|$, of each of the filters.
   b. Estimate the transition width of each filter from your frequency plots
   c. What is the approximate minimum attenuation of each filter in the stop band.
2. Repeat question 1 using Hamming windows. Compare your results to those in question 1.
3. Repeat question 1 using Blackman windows. Compare your results to those in question 1.

CHAPTER 17: CUSTOM FILTERS

The method described in this chapter often go by the title “Frequency Sampling” since you start by developing a vector of samples of the desired frequency response and use the inverse FFT to get the required pulse response.

1. Use the frequency sampling method to design a 25-tap lowpass FIR filter with a cutoff frequency of 0.25π radians/sample. Plot both the pulse response and magnitude of the transfer function for this filter.
2. Apply a Hamming window function to the pulse response of problem 1 above. How did the transfer function improve? What parameter of the transfer function was degraded and by what factor? What would you change to fix the degraded performance while retaining the improvement?
3. Repeat problem 2 using a Blackman window function.

CHAPTER 18: FFT CONVOLUTION

CHAPTER 19: RECURSIVE FILTERS

CHAPTER 20: CHEBYSHEV FILTERS

CHAPTER 21: FILTER COMPARISON

CHAPTER 22: AUDIO PROCESSING

CHAPTER 23: IMAGE FORMATION & DISPLAY

CHAPTER 24: LINEAR IMAGE PROCESSING

CHAPTER 25: SPECIAL IMAGING TECHNIQUES

CHAPTER 26: NEURAL NETWORKS (AND MORE!)

CHAPTER 27: DATA COMPRESSION

CHAPTER 28: DIGITAL SIGNAL PROCESSORS

CHAPTER 29: GETTING STARTED WITH DSPS

CHAPTER 30: COMPLEX NUMBERS

CHAPTER 31: THE COMPLEX FOURIER TRANSFORM

CHAPTER 32: THE LAPLACE TRANSFORM

CHAPTER 33: THE Z-TRANSFORM

CHAPTER 34: EXPLAINING BENFORD’S LAW