EECS 452 Midterm Exam (solns)

Fall 2012

Name: _____ unique name: _____

Sign the honor code:

I have neither given nor received aid on this exam nor observed anyone else doing so.

_____

Scores:

#	Points
Section I	/40
Section II	/30
Section III	/30
Total	**/100**
Class mean	70
Class std	17
Max	100
Min	33

NOTES:

- Open book, open notes.
- There are **8** pages including this one.
- Calculators are allowed, but no PDAs, Portables, Cell phones, etc.
- You have 120 minutes for the exam.
- **Be sure to show work and explain what you've done when asked to do so.** You will not receive partial credit without showing work.
- Unless otherwise specified all signed numbers are two's complement numbers.

Section I -- Short answer 40 points

1) Convert the decimal value -1.25 into: [4]

   a) 4 bit binary Q2: __1011_____

   b) 8 bit binary Q2 : __11111011_

   c) 8 bit Hexadecimal Q2: _ 0xFB____

   d) 16 bit hexadecimal Q14 __0xB000___

2) Show all the steps involved for multiplying the 4 bit Q3 number 0x81 by the 8 bit Q4 number 0x24. How many bits does the product have and what is its Q-value? [5]

```
                10000001
                00100100
                --------
       11111110000001
       11110000001
-------------------------------
       1110111000100100
```

Or equivalently, 0xEE24. The resultant has 16 bits and is in Q7. The decimal value is -35.7188.

3) The numbers 0xA2 and 0x04 are multiplied together and added to the product of the numbers 0x51 and 0x02. Assume all numbers are in 16 bit twos complement Q7. What is the result of this MAC operation a) in hexadecimal and b) in decimal value? Is there overflow? [5]

Note that the problem says that the numbers are in 16 bit twos complement. Thus, there is room for interpretation. Depending on your interpretation the correct answers are 0x0000032A (0.0494 decimal) or 0xFFFFFF2A (-0.0131 in decimal) in 32 bit Q14.

Soln 1: If you interpreted 0xA2 to be shorthand for 0x00A2, 0x04 shorthand for 0x0004, 0x51 shorthand for 0x0051, and 0x02 shorthand for 0x0002 then all numbers are positive since they all have leading zeros in the MSB.
(0x00A2)*(0x0004)=0x00000288 (shift 0x00A2 left by two noting that multiplier accumulator register is 32 bit Q14).
(0x0051)*(0x0002)=0x000000A2 (shift 0x0051 left by one)

The sum is simply 0x288+0x0A2=0x32A in 32 bit Q14.
The twos complement value 0x32A is equal to decimal value 810*2^(-14)=0.0494.
There is no overflow.

Soln 2: If you interpreted 0xA2 as 0xFFA2 and the others as in Soln 1 then the first number is negative and the first multiply is (0xFFA2)*(0x0004)=0xFFFFFE88. In 32 bit Q14. The second multiply is the same as above, i.e. 0x000000A2 (with zero extension). Adding these together you obtain 0xFFFFFF2A in 32 bit Q14. Again there is no overflow.

4) Explain the role of an equalizer filter in sampling and reconstruction. What is the utility of such a filter, what does it seek to compensate, where is it used in the embedded processing system, and how does it differ from an anti-alias filter? [5]

This answer can be taken right out of lecture 7, p. 18: "Apply equalization filter to invert the shading frequency function". More specifically, in sampling and reconstruction the equalizer filter is used in a CODEC and compensates for distortion created by a sample-and-hold reconstruction device. The device causes the spectrum of the ideal (cardinal series) reconstruction to be shaded by a sinc function. The filter is typically implemented after the DA (sample-and-hold) as an analog filter. It can also be implemented as a digital filter after the AD or in combination with the analog filter.

5) An A/D converter circuit is shown below.

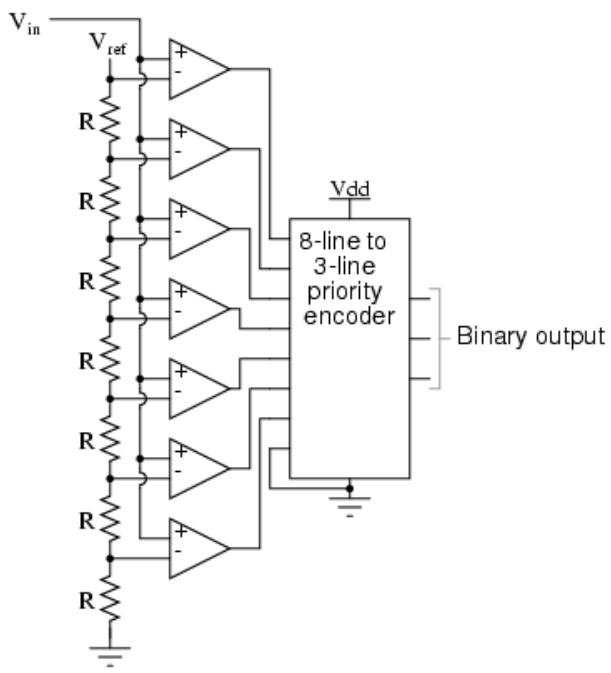

What is the number of quantization levels of this A/D converter? Sketch the transfer function of the A/D converter. What is the maximum A/D conversion error? Is the A/D converter compensated or uncompensated? If uncompensated how would you compensate it? [5]

The A/D converter has 8 levels that correspond to the following input voltage ranges ($\Delta = Vref/7$).

000	001	010	011	100	101	110	111	binary	output
0 to Δ	Δ to 2Δ	2Δ to 3Δ	3Δ to 4Δ	4Δ to 5Δ	5Δ to 6Δ	6Δ to 7Δ	7Δ to 8Δ	input	ranges

The maximum conversion error is Δ. The A/D is uncompensated. To compensate it you could replace the lowest resistor (connected to ground) by a resistor of resistance value ½ R and add another resistor of resistance value ½ R at top (connected to Vref).

6) Sketch the block diagram of a second order FIR implementation of an FIR filter in Direct Form and in Transpose Form. Under what circumstances might you prefer the Direct Form over the Transpose Form? [5]

Just copy these block diagrams from the notes (Lect 7 or 8). The direct form is preferred when you wish to limit the number of 32 bit multiply registers (saves power, memory, and FPGA fabric) since you can replace the M multiply-adds in the FIR direct form diagram by a single 32 bit MAC operation.

7) A 1024-pt FFT is implemented to compute the spectrum of samples of a real valued signal sampled at 48kHz. What is the frequency spacing of the 1024-pt FFT in Hz? If the sinusoidal signal consisted of a sinusoid with zero phase and unit amplitude at frequencies at 8kHz what is the magnitude of the FFT at this frequency? What is the FFT magnitude at frequency 8010Hz and at 7090Hz? Is there spectral leakage? [6]

The spacing is ΔF =48,000/1024=46.875Hz. As 8000/ ΔF = 170.6667 the sinusoidal frequency 8kHz is not an integer multiple of ΔF and the closest FFT frequency bins will be bin 170 and 171 so there is spectral leakage. The value of the magnitude FFT at these frequencies will be |sin(pi*(f_{170}-fc)N)/sin(pi*(f_{170}-fc))| and |sin(pi*(f_{171}-fc)N)/sin(pi*(f_{171}-fc))| N=1024, fc=8000/48000. f_{170}=170/N and f_{171} =171/N (recall lect 11 p. 3). The magnitude of the FFT near the frequencies 8010Hz and 7090Hz will be given by the same formulas with $f_{170,}$ f_{171} replaced by the closest FFT analysis frequencies near 8010/48000 and 7090/48000, respectively.

8) An N-point implementation of the FFT requires N/2 log N complex arithmetic operations. If each operation consumes 20 clock cycles, the CPU speed is 2GHz and the sampling rate is 48kHz what is the maximum size FFT (maximum number of points N) that can be implemented for real-time audio spectrum analysis? What is the corresponding best attainable frequency resolution in Hz? [5]

The time to do one complete FFT is N/2 log2(N)*20/(2*10^9). This must be less than the sampling period T=1/48000 for the FFT to be implemented in real time on disjoint blocks of data (we assume no buffering is available). Therefore:
N/2 log2(N)*20/(2*10^9)<T or N/2 log2(N)<T*10^8=2083.3.
We will restrict to a radix 2 FFT, i.e. N has to be an integer power of 2. Note that if N=256 then log2 (N)=8 and N/2 log2(N)=1024 which is the largest such N that satisfies the inequality (N=512 is too large as it gives N/2 log2(N)=2304). Therefore N=256 is the right answer. The frequency resolution for this value of N is Fs/N=48000/256= 187.5Hz.

Section II – longer answers 30 points

1) The following is the block diagram of a digital filter. [15]

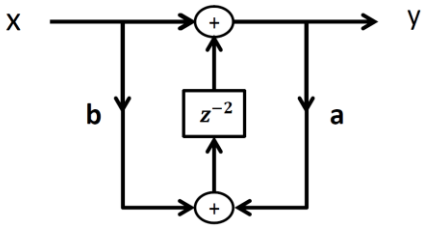

a) What is the filter's transfer function in the Z domain? [6]

$Y(z) = (1 + bz^{-2})X(z) + az^{-2}Y(z)$ so the transfer function is

$$H(z) = \frac{Y(z)}{X(z)} = (1 + bz^{-2})/(1 - az^{-2})$$

b) Assume that a=3/4 and b=-1. Plot the pole zero constellation of this filter and sketch the magnitude of its frequency response. [6]

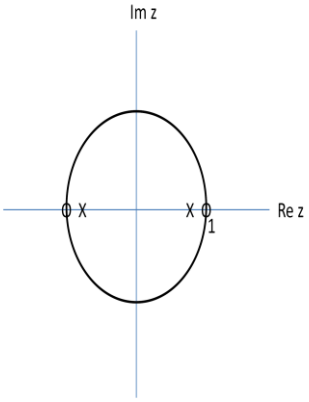

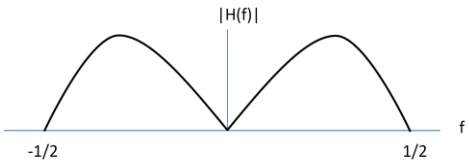

c) Classify the filter as lowpass, highpass, bandpass, bandstop, or none of the above. [3]

This could be interpreted as a bandpass filter with a wide bandwidth.

9) The magnitude DFT X[k] of a sequence {x[0],…,x[N-1]} is given below. [15]

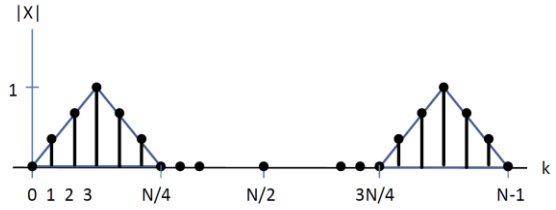

a) Plot the magnitude DFT of the length 2N sequence{0,…,0, x[0],…,x[N-1]} (zero padded in front). [5]

This is simply a zero padding of the sequence before the first point (you have seen this in class and in EECS451). Let {y[n]} denote this sequence. The sequence y[n] has DFT equal to: $Y_{DFT}[k] = X_{DTFT}\left(\frac{k}{2N}\right)e^{-j\pi k}$, k=0,1, …,2N-1. Hence $|Y_{DFT}[k]|$ (denoted |X| in the figure) is interpolated version of |X[k]|:

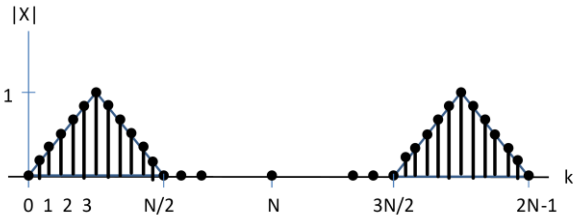

b) Plot the magnitude DFT of the length 2N sequence {x[0],x[1],…,x[N-1], 0,…,0} (zero padded in back). [5]

This is simply a zero pading of the sequence after the last point (you have seen this in class and in EECS451). Let {u[n]} denote this sequence. The sequence u[n] has DFT equal to: $U_{DFT}[k] = X_{DTFT}\left(\frac{k}{2N}\right)$, k=0,1, …,2N-1. Hence $|U_{DFT}[k]|$ (denoted |X| in the figure) is the interpolated version of |X[k]| shown in part (a).

c) Plot the magnitude DFT of the length 2N sequence {x[0],0,x[1],0,…,x[N-2],0,x[N-1],0} (zero interleaved). Comment on the differences between (a), (b), (c). [5]

This is simply the decimation of the sequence (you have seen this in EECS451). Let {v[n]} denote this sequence. The sequence v[n] has DFT equal to: $V_{DFT}[k] = X_{DTFT}\left(\frac{k}{N}\right)$, k=0,1, …,2N-1. Hence $|V_{DFT}[k]|$ (denoted |X| in the figure) takes the form

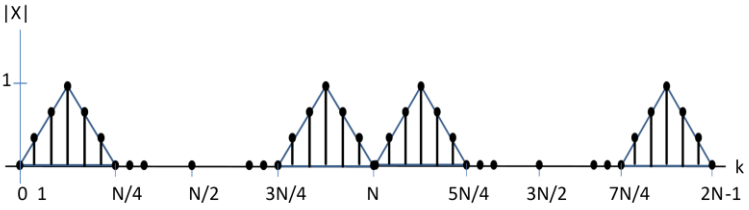

Section III Longer answers 30 points

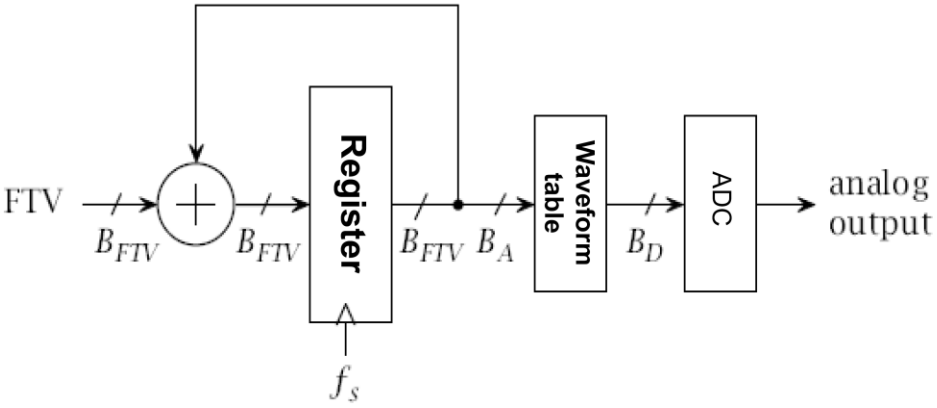

1) An arbitrary waveform synthesizer uses the DDS system above to generate a periodic waveform. The waveform table contains samples of a single period of a waveform s[0],…,s[2^B-1] where B=B_A. Assume that s[n] is a ramp s[n]=n 2^{-B} and that the ADC is ideal in the following. [15]

   a) Assume a clock rate fs = 20MHz, B_{FTV} =16, B_A =8, and FTV=1. What is the period and frequency of the output signal? Plot three or four periods of the analog output over time. Be careful to label your axes. [7]

The frequency is F=$FTV\ Fs/2^B$ which is equal to $20 * \frac{10^6}{2^{16}} = 305.2Hz$. The period is T=1/F=3.28 msecs.

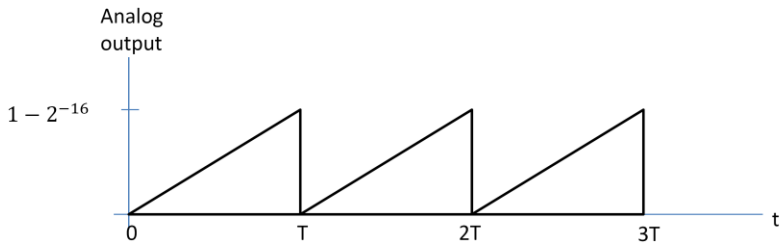

   b) Now we let FTV= FTV[n] be a time varying, where n increments with each uptick of the clock f_s . Specifically, assume that FTV[n] is equal to zero for n even valued and equal to one for n odd valued. Assuming the other parameters are the same as in (a) . What is the period and

frequency of the analog output signal? Draw three or four periods of the analog output as in (a).[8]

The only difference is that the counter upticks only ½ as fast. Thus the frequency is F/2 which is equal to $152.6Hz$. The period is twice as long: 6.56 msecs.

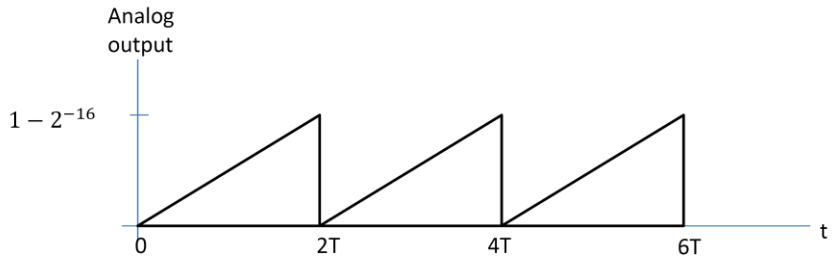

2) A pole zero constellation of a filter is shown below. [15]

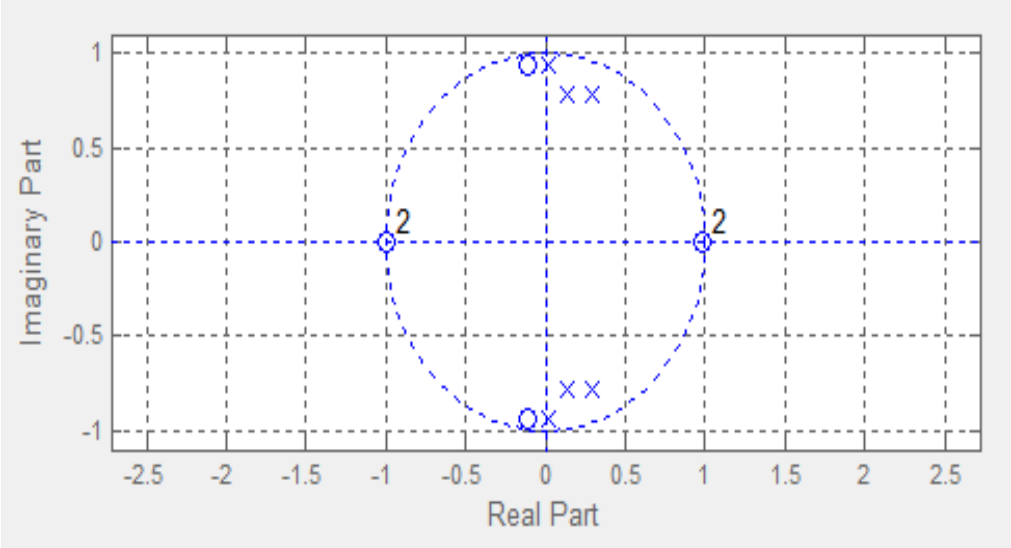

a) What kind of filter (Low Pass, High Pass, Band Pass, Band Stop) best describes it? [3]
This is really none of the above bandpass (although I would accept bandpass with center frequency near f=1/4 or bandstop with frequency at f=1/4).

b) It is desired to implement this filter using a cascade of biquad's as you did in Labs 5 and 6. How many biquads are required to implement it? [4]
3 biquads needed

c) Circle the pole-zero combinations that you would use for for each biquad in order to minimize the likelihood of overflow in your finite precision filter implementation. Label each combination 1,2,… corresponding to the order (from left to right) in which you would cascade the biquads in your filter implementation. [4]

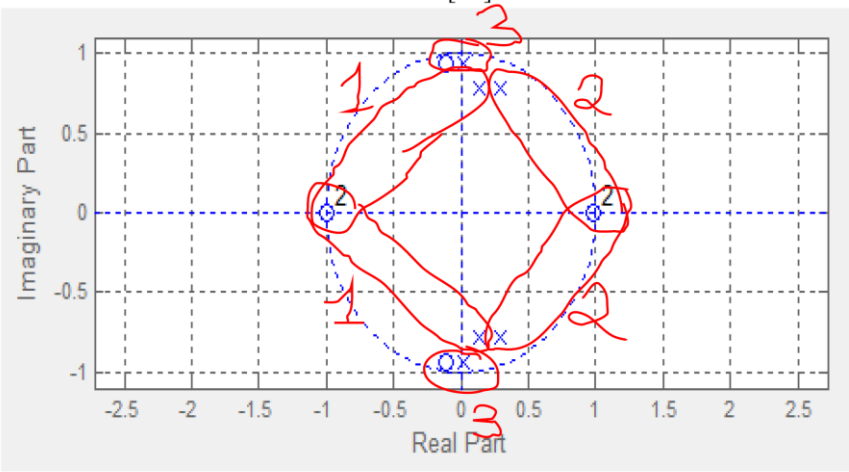

d) Would you need to normalize any of the denominator coefficients of any of the biquads in order to implement this filter in 16 bit Q15? If so which biquad(s) would you need to normalize in this manner? [4]

This is similar to the solution of the last problem on the practice exam. None of the poles have angles in the ranges $\left[-\frac{\pi}{3},\frac{\pi}{3}\right]$ *or* $[-\frac{2\pi}{3},\frac{2\pi}{3}]$ where the denominator of the biquad transfer function would have coefficient a1=2 r cos(θ) with magnitude greater than one.