Part I: fill in the blanks (25 points)

1. (25 pts total, 1 pt per blank) Fill in the blanks.

   (a) If we multiply a 16-bit $Q5$ value by a 16-bit $Q9$ value the result is a __32__-bit Q __14__ value.

   (b) A B-bit unsigned Q(B-1) number can represent all numbers from __0__ to $2 - 2^{-(B-1)}$.

   (c) A B-bit two's complement Q(B-1) number can represent all numbers from __-1__ to $1 - 2^{-(B-1)}$.

   (d) A 7-bit Q5 two's complement number has a range from __-2__ to $2 - 2^{-5}$ with a resolution of 2^{-5}.

   (e) The 6-bit two's complement Q2 number 100111 is __-6.25__ in decimal.

   (f) Two's complement multiplication of 0111 and 1000 gives: __11001000__.

   (g) We want to use an A/D converter to quantize a signal whose magnitude ranges from -1 to 2 V. To completely cover this range and use a step size no greater than 0.1V, the ADC needs to have at least $\lceil \log_2(30) \rceil$ bits. The absolute error (between any input and output value pair) of this A/D converter is at most __0.05__ V.

   (h) When sampled at 10KHz, a sinusoid of 23KHz would appear at __3__ KHz, a 8KHz sinusoid would appear at __-2__ KHz, and a 10KHz sinusoid appears at __0__ KHz. (Limit your answers to [-5, 5]KHz)

   (i) The 6-bit Q2 two's complement number 011011 is converted to a 4-bit Q0 number using regular rounding: __0111__, and using convergent rounding: __0111__.

   (j) An FIR filter is always stable, yes or no? __yes__

   (k) There exists IIR filters with linear phase response, yes or no? __no__

   (l) For an IIR filter to be stable, all zeros and poles must lie within the unity circle, yes or no? __no__

   (m) In order to have a frequency spacing of 0.5 Hz in the DFT with a sample rate of 40KHz, you will need __80000__ samples.

   (n) Can we use two's complement add/sub hardware to compute the unsigned add/sub, yes or no? __yes__

   (o) Of the four types of IIR filters discussed in class, the __Butterworth__ requires an excessive number of poles, and __Elliptic__ has the smallest transition band for a given order.

Part II: short questions (45 points)

2. (4 pts) In the following decimation-in-time FFT data flow diagram please fill in the required index values on the input x values.

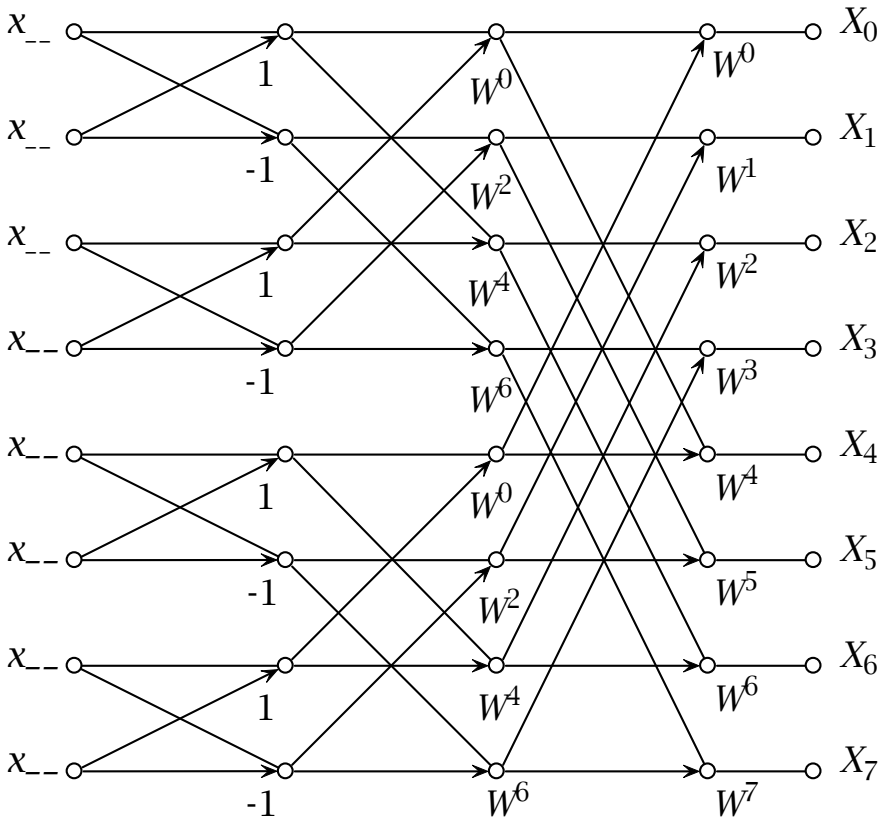

Going from top to bottom: 0, 4, 2, 6, 1, 5, 3, 7

3. (5 points) Consider a 3-tap FIR filter with coefficients $h[3] = \{28672, -16384, 28672\}$, all expressed as 16-bits Q15 numbers.

   (a) (3 pts) Find the decimal values for the filter coefficients.

   $h[0] = h[2] = 28627/2^{15} = 0.875$; $h[1] = (-16384)/2^{15} = -0.5$

   (b) (2 pts) If the input to our filter is between -1 and 1, what would be the maximum value at the output?

   The maximum value at the output is $|h[0]| + |h[1]| + |h[2]| = 2.25$.

4. (9 pts) Answer questions on the figures shown below, assuming $V_{ref} = 12$V. (If you don't recall a circuit's name just briefly state its intended purpose).

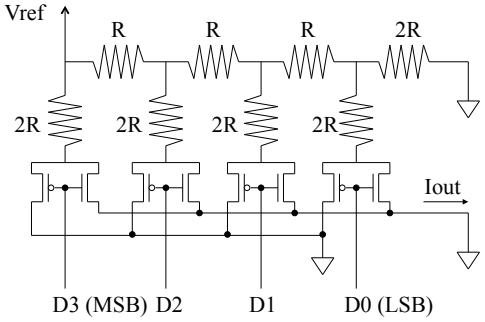

(a)

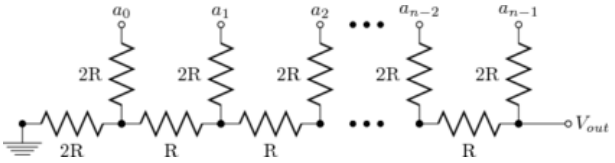

(b)

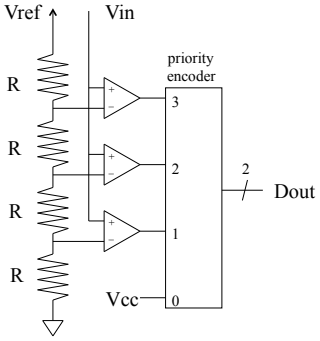

(c)

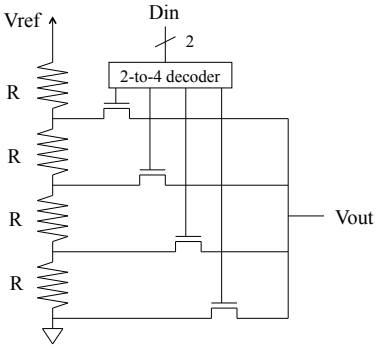

(d)

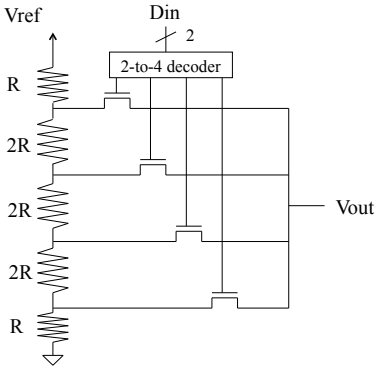

(e)

(a) (1 pt) What is this circuit called?
This is a 4-bit R-2R ladder DAC.

(b) (2 pts) Assume a bit a_i is attached to V_{ref} if its value is 1 and grounded if its value is 0. Take $n = 4$. If the bits are (from MSB to LSB) 1001, what is V_{out}?

This is an n-bit R-2R ladder DAC. V_{out} is $1001_b \times \frac{1}{2^4} \times V_{ref} = \frac{9 \times 12}{16} = 6.75V$.

(c) (2 pts) What is this circuit called? For $V_{in} = 10.4V$, what is D_{out}?

This is a voltage divider serving as a 2-bit ADC. $D_{out} = 11$.

(d) (2 pts) What is this circuit called? For $D_{in} = 10$, what is V_{out}?

This is a 2-bit DAC. $V_{out} = 6$V.

(e) (2 pts) What is this circuit called? For $D_{in} = 10$, what is V_{out}?

This is a 2-bit DAC. $V_{out} = 7.5$V.

5. (5 pts) The counter below uses 3 D flip-flops.

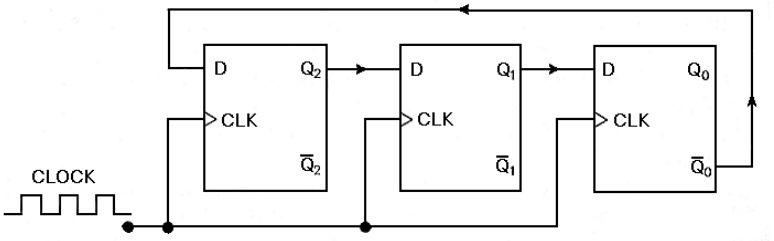

(a) (3 pts) Draw the timing table indicating the states of the counter over the first 8 clock cycles (assume initial state is 000).

Clk	Q_2	Q_1	Q_0
1	1	0	0
2	1	1	0
3	1	1	1
4	0	1	1
5	0	0	1
6	0	0	0
7	1	0	0
8	1	1	0

(b) (2 pts) What is the cycle of this counter?

The cycle is 6 clock ticks/cycles.

6. (5 pts) Consider two 16-bit Q15 format numbers x and y. We multiply $x \cdot y$ and need the result ("output") back to 16-bit Q15 format after regular rounding. Write down the C code that accomplishes this by filling out the blanks below.

```
Int16 x, y, output;
Int32 temp;
temp = (Int32)x*(Int32)y;
output = _____________________________________;
```

(Int16)(temp+ 0x00004000) » 15

7. (7 pts) We wish to implement the following filter: $y[n] = x[n] + x[n-1] + x[n-2] + \cdots + x[n-k]$.

   (a) (1 pt) Of what order is this filter?

   Order k; k past samples

   (b) (1 pt) Of what type is this filter?

   FIR; or averaging filter

   (c) (2 pts) Does this filter have constant group delay? If not why? If so what is the group delay for $k = 21$ and a sample rate of 100KHz?

   Yes it does, because all its coefficients are equal. 21 is the order of the filter. Using known results on linear phase FIR filters, the group delay of the filter is $21/2$ samples. This gives a group delay of $\frac{21}{2} \cdot \frac{1}{100000} = 0.105$ ms.

   (d) (3 pts) Please draw an implementation diagram for the above filter using as few adders (each adder limited to two inputs) as possible.

   You should only need 2 (actually you can have a single adder with multiple inputs, but 2 is good enough for this problem). Here are two ways to implement it, each transpose of the other.

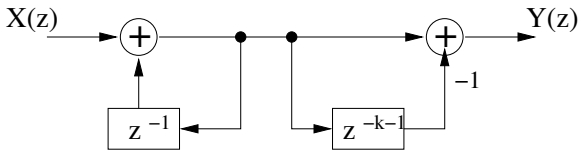

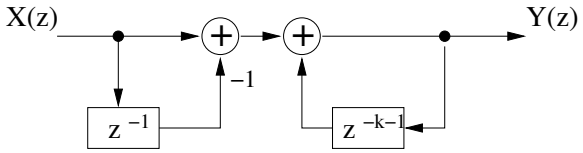

8. (10 pts) Suppose a 32-point DFT sampled at frequency $f_s = 3200$Hz gives you the following result for indices $k = 0, 1, \cdots, 16$:

$$\begin{aligned} X[0] &= 64\angle 0; \\ X[1] &= 0; \\ X[2] &= 32\angle 180; \\ X[3] &= 48\angle -90; \\ X[k] &= 0, \quad k = 4, 5, \cdots 16. \end{aligned}$$

(a) (2 pts) Please write down the DFT values for $k = 17, 18, \cdots, 31$.

$$\begin{aligned} X[30] &= 32\angle -180; \\ X[29] &= 48\angle 90, \end{aligned}$$

and zero everywhere else.

(b) (3 pts) Please write down the signal whose samples generated the above DFT.

The signal would be
$3\sin(2\pi(300)t) + 2\cos(2\pi(200)t + \pi) + 2$.

(c) (2 pts) For this signal and at sample rate $f_s = 3200$Hz, what is the smallest number of samples required in the DFT to avoid leakage?

$N = 32$

(d) (1 pt) For this signal, what is the smallest sampling rate to use in order to avoid aliasing?

$f_s = 600$Hz

(e) (2 pts) If we take a $N = 16$-point DFT, what is the smallest sampling rate to avoid aliasing and leakage?

$f_s = 800$Hz

Part III: longer questions (30 points)

9. (8 pts) Let $x(t) = 3\sin(2\pi f_1 t) + 2\sin(2\pi f_2 t) + \sin(2\pi f_3 t)$ where $f_1 = 800$ Hz, $f_2 = 1200$ Hz, and $f_3 = 1300$ Hz.

   (a) (4 pts) Sample the signal at $f_s = 2500$ Hz and take 100 samples and perform DFT using the "non-standard" definition by multiplying the forward DFT with $1/N$. Please sketch the resulting spectrum – you can either write down the DFT coefficients or plot them on a figure. In either case please specify the magnitude, phase, and the resolution in frequency.

   They key thing to note here is that a sample rate of 2500 is not sufficient for the $f_3 = 1300$ element, and you will get aliasing. In particular, it will appear at -1200Hz if you use the range $[-1250, 1250]$Hz, or 1300Hz if you use the range $[0, 2500]$. Everything else is standard DFT calculation by inspection. This is shown in the figure below.

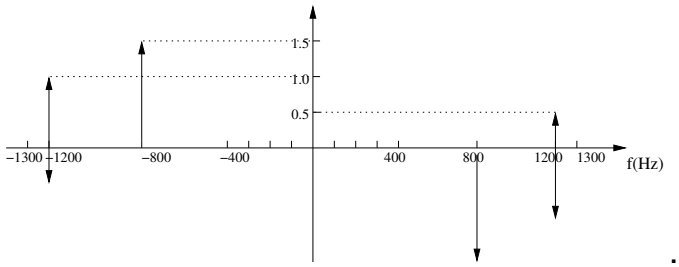

   (b) (4 pts) Repeat (a), but this time suppose that the signal is first passed through an anti-alias filter whose response is shown below, then sampled at $f_s = 2500$ Hz, with 100 samples taken and DFT performed. Please sketch the resulting spectrum.

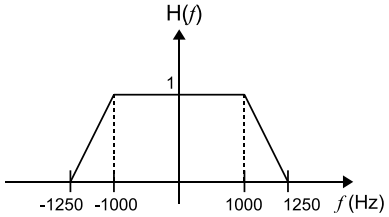

   Once through this filter, the f_3 component is gone, and the f_2 component is attenuated by a gain of 0.2, resulting in a 0.2 amplitude (this was mistakenly labeled as 0.4 on the figure).

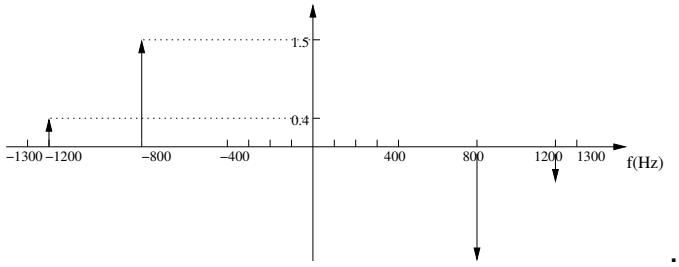

10. (9 pts) Consider the DDS diagram shown below.

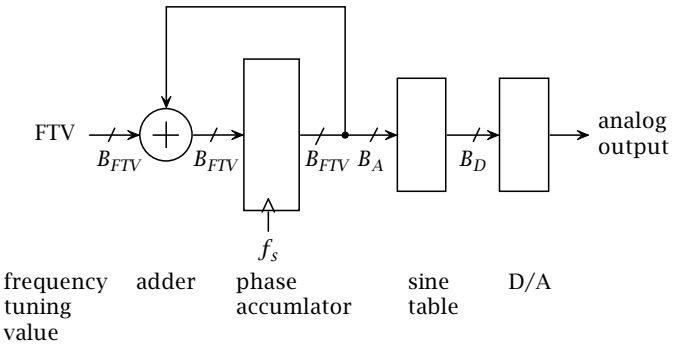

(a) (2 pts) Suppose $FTV = 4$, $f_s = 10$MHz, $B_{FTV} = 12$, and the sine table contains 256 values denoted as $x(0), x(1), \cdots, x(255)$. Assume the system is initialized to zero. Please write down the first 10 values coming out of the sine table, each upon a clock tick.

Since $FTV = 4$, the values coming out of the accumulator are $B_{FTV} = 0, 4, 8, 12, 16, 20, \cdots$. The sine table is addressed with 8 bits ($256 = 2^8$); these are the high 8 bits of B_{FTV}. The sine output value changes every time the low 4 bits overflows (or carries). So the output sequence is

$$x(0), x(0), x(0), x(0), x(1), x(1), x(1), x(1), x(2), x(2)$$

(b) (2 pts) Repeat (a) but with a sine table containing 512 values, denoted as $x'(0), x'(1), \cdots, x'(511)$.

Now the sine table is addressed with 9 bits ($512 = 2^9$). The sine output value changes every time the low 3 bits overflows. So the output sequence is

$$x'(0), x'(0), x'(1), x'(1), x'(2), x'(2), x'(3), x'(3), x'(4), x'(4)$$

(c) (2 pt) What is the output waveform frequency of (a) and (b)? Which one is higher?

Neither. The two are of the same frequency, as changing the resolution of the sine table has nothing to do with the output frequency, only the resolution or signal quality.

(d) (3 pts) Suppose that we need the output to have a range of $[5, 20]$MHz, tunable by the value FTV. Assume the FTV value is limited within the range of $[1, 10]$. What should be your choice of B_{FTV}?

To achieve the range, we need $1 \cdot \frac{f_s}{2^B} \leq 5$MHz and $10 \cdot \frac{f_s}{2^B} \geq 20$MHz. This means we need

$$\frac{10}{5} \leq 2^B \leq 10\frac{10}{20}, \quad \text{or} \quad \log_2(2) \leq B \leq \log_2(5) \ .$$

So B may be set to 1 or 2. However, note that this problem is fundamentally ill-posed because one cannot get output frequency higher than f_s in principle.

11. (13 pts) While designing an IIR filter, Matlab gave the following conjugate pairs of zeros:

$$A: \quad -0.9517 + j0.3071; \quad -0.9517 - j0.3071$$

$$B: \quad -0.5451 + j0.8384; \quad -0.5451 - j0.8384$$

$$C: \quad -0.7087 + j0.7055; \quad -0.7087 - j0.7055$$

and the following conjugate pairs of poles:

$$D: \quad 0.0031 + j0.9495; \quad 0.0031 - j0.9495$$

$$E: \quad 0.2024 + j0.7916; \quad 0.2024 - j0.7916$$

$$F: \quad 0.5157 + j0.3565; \quad 0.5157 - j0.3565$$

(a) (4 pts) Implement this filter with a cascade of biquads. Using rules discussed in class, indicate how you would order the biquads and what poles and zeros to include in each biquad. Please use the labels (A-F) given above. A diagram is provided below for you to work on.

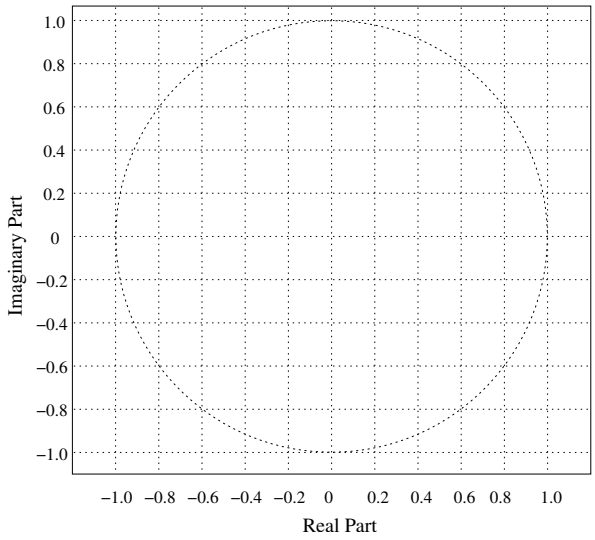

The sequence would be A-F, followed by C-E, followed by B-D. Also note that all zeros are basically exactly on the unit circle along the negative half, while the poles are all within the positive half of the circle.

(b) (3 pts) Give a rough sketch of the frequency response of this filter assuming $f_s = 12$kHz. What type of filter (lowpass, highpass, bandpass, bandstop) is this?

From the earlier description of the locations of the zeros and poles, we see that this is a lowpass filter, with the first zero crossing (induced by B, which is at an angle of approximately $\frac{2}{3}\pi$) at close to $f_s/2 \times 2/3 = 4$kHz.

(c) (2 pts) Take the first biquad of your design, and calculate the coefficients of the biquad transfer function $H(z)$. Show expressions if you don't know how to compute the precise coefficient values.

The coefficient values would be: $b_o = 1; b_1 = 1.903; b_2 = 1; a_o = 1; a_1 = -1.0314; a_2 = 0.343$.

(d) (4 pts) Draw the direct form I (DF1) diagram as well as the transposed direct form II (TDF2) diagram of this first biquad, labeled with all the coefficients.

Diagrams are not shown here, but you can find it easily in the lecture notes. The coefficient values would be: $b_o = 1; b_1 = 1.903; b_2 = 1; a_o = 1; a_1 = -1.0314; a_2 = 0.343$.