473 Fall 2023 Homework 2 answers

Due on Gradescope by 7**pm** on December 8th. ~~We'll take only 5% off if turned in my 5pm on the 6th.~~ 140 points. Corrections are in red.

Review—because September was a long time ago. [25 points]

1) Capacitors and power integrity **[8 points]**

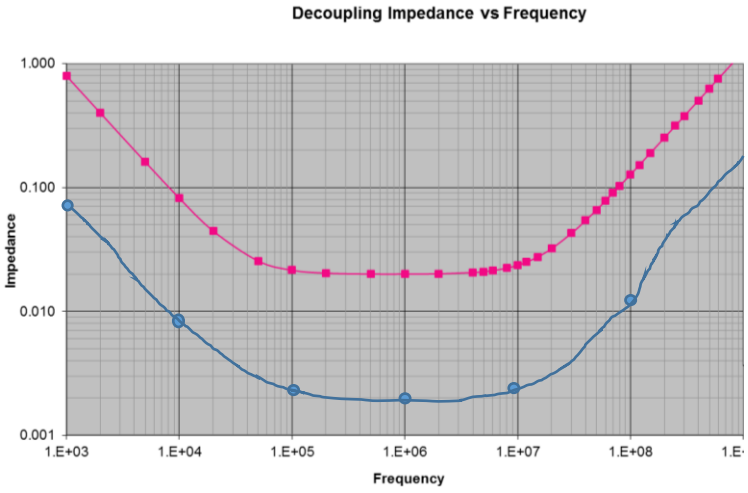

   a) In your own words, explain why we are concerned about the frequency response of the power/ground lines. **[4]**

      We are looking to be able to hold the difference of those two lines constant for each device connected. The devices we are powering have a variable current draw—at different times they may draw different amounts of current. As their current draw numbers vary, at fairly high frequencies, the power supply will not be able to react quickly enough in varying the amount of power supplied. We need to be able to react to very quickly changing current needs (high frequency) as well as longer-term (lower frequency) changes.

   b) In the context of power integrity, why might it be preferable to have two capacitors in parallel rather than one capacitor of double the capacitance? **[2]**

      In general, two capacitors in parallel will add their capacitance but have half the parasitic (resistance and inductance) values and so be closer to an ideal capacitor than the one capacitor of double the capacitance.

   c) Consider the capacitor described on the right. Redraw the curve to represent ten of the same capacitors in parallel. **[2]**

      Drawing should be about as shown above with the line shifted down by a factor of 10 in all places.

2) True/False **[9 points, -2 per wrong answer, minimum 0]**
For each statement, indicate if it is true or false. If it is false, explain what makes it false. **[9 points]**

a) **A LIPO battery is being used with a constant current draw of "5C". You would expect that the battery would last at least 12 minutes, but probably not much longer**. It would last *at most* 12 minutes.

b) Linux user-space programs are generally expected to use memory-mapped I/O addresses to talk with I/O devices. They use a file I/O model instead.

c) When designing a power distribution network, very low frequency noise (say 1KHz) is primarily handled by the power supply/voltage regulator.

d) Given our standard assumptions about scheduling, anything that can be scheduled using EDF scheduling can also be scheduled by RM scheduling. Anything that can be scheduled by RM can also be scheduled by EDF as EDF is optimal. But RM isn't optimal, so no.

e) An ideal switching power supply would waste no power. An ideal linear regular would.

f) The "lifetime" of a primary cell has to do with the number of times you can charge it before its capacity drops significantly. The lifetime of a primary cell is related to shelf life.

g) On a PCB, a 20 mil-wide trace has less resistance than a 10mm-wide trace of the same length. 10mm is much wider than 20 mils. So, the 10mm-wide trace has less resistance.

3) Take a look at problem 6 from our midterm. In your own words, explain the answer to problem 6b **[8 points]**

Basic theme: We put 0,1,2,3 in the buffer, change f_pos to 3 and return 2. So cat sees just 0 and 1 returned. On the next pass we put 3, 4, 5, 6 into the buffer, set f_pos to be 6 and return 2. So cat sees just 3 and 4. Next pass is the same but starting at 6, so cat sees 6 and 7. After that we return a 0, so cat has read everything and stops asking. So we get 013467.

Digital Signal Processing and other specialized processors [45 points]

4) Fixed point **[16 points]**

   a) What are the primary reasons we might use fixed point rather than floating point? **[3]**
   Chips with hardware floating point support typically cost more and draw more power. As such, we sometimes don't have that hardware support. Using floating point instructions without hardware support (so software emulation of floating point) generally involves a huge number of integer instructions per floating point instruction. (Do note that a C compiler compiling for a machine without floating point support will generate the integer instructions needed to do floating point, so it isn't that doing this is difficult.)

   b) Say you multiply two int8_t Q7 numbers as standard integers. What size/form will the result be assuming that the size of an "int" on your machine is 16 bits and no casting was used? Will there be any loss of data? **[2]**
   You will get a 16-bit Q14. There will be no loss of data (the fact that it's Q7 has no impact on that, we're must multiplying an 8-bit number with an 8-bit number and getting a 16 bit number, which loses no data).

   c) As b but the numbers re int16_t Q7 numbers. **[3]**
   This is a bit more tricky. In C or C++ operations (including multiplication) on A and B are done as the larger of A, B, or the int size. In this case, those are all 16 bits, so we'll get a 16-bit Q14. And yes, we may lose data as the answer may need 32 bits.

   d) Write a C function "Q15mult" which takes two values, A and B, each declared as an int16_t number and returns an int16_t number as the product of A and B. All three numbers are treated as Q15. There will be a fair bit of casting. Round to the nearest value (you can break a tie as you wish). Assume that the default "int" size on your computer is 16 bits (it does matter). **[8]**

```
int16_t Q15mult (int16_t A, int16_t B)
{
// It is a 32-bit Q30 after multiplication.  So, we want to shift it by 15 to
// get a Q15. The casting is needed because the return value of an operation
// (such as *) is the larger of the largest argument or the int size
// (all of which are 16 without casting).  The +0x4000 is done so the
// truncation after shifting will cause it to round to the nearest value.

return ((int32_t)A*B + 0x4000) >> 15;
```

5) One issue in modern computer architecture is that for a novel design we tend to find that the relationship between power (P) and performance (perf) is such that $P=perf^x$ where x is probably between 2 and 3. Thus, if I want to double the performance of a core with a "from scratch, redesign" we are looking at the power being 4 to 8 times as high. While this really harms high-end processor performance (too much power means it all melts), this relationship can be very helpful with respect to IoT devices. Why? **[8 points]**
   For many (probably most) IoT devices, we are very concerned about power consumption and have fairly low performance needs. So by reducing performance by, say, a factor of 10, we would be reducing power by a factor 100 to 1000. That's a trade-off that we're very happy to take.

6) Embedded AI overview **[13 points]**
Read https://www.edgeimpulse.com/blog/the-embedded-machine-learning-revolution-the-basics-you-need-to-know.

a) The author suggests that it's possible to train AI models in the cloud and execute them on embedded devices. **[4]**
   i) What is the difference between training and executing a model in this context?
   ii) Why would the training be done in the cloud rather than on the device?

This type of AI is all about getting the tool to be able to find things. It is generally done by some form of reinforcement training where you do things like give it a bunch of pictures of apples and a bunch of pictures of not apples and have it "learn" hot to distinguish which is which. Training is about teaching it, executing is about having it do the task based on that training.

Training can take a lot of time, computation, and data (for example to store all the training data). An embedded system is often not the ideal place to do that training (though it often can be done).

b) The author provides a list of reasons why one might choose to execute the model on an embedded system rather than in the cloud. Consider two applications: A) detecting bats and classifying them by species based on their call (should sound familiar) and B) counting the number of apples on a tree to figure out if they are ready for harvesting. For each application pick one or two of his benefits you think are most relevant to the application and justify why you picked those. **[9]**
Both are really quite similar. Privacy probably isn't a big issue for either of these, though pictures or recordings of people might be acquired during such a process and thus could be an issue. Both probably care about **bandwidth** quite a bit as pictures and audio are both pretty heavy. Neither probably cares much about latency—it seems unlikely we'd have an immediate use for the data—in fact taking a day or so probably would be okay. So probably bandwith.

7) List three different types of "specialized compute" processors other than a DSP and provide a link to an example processor that is that type of specialized compute. **[8]**
We'll take anything reasonable. GPUs are perhaps the most obvious, but there are plenty. Some examples:

- AI: Jetson Xavier: https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-xavier-series/
- Network processor: IXP1200 (from the 90s!) https://en.wikipedia.org/wiki/IXP1200
- A processor to run LISP: the LISP machine https://en.wikipedia.org/wiki/Lisp_machine, a really old machine.

Wireless communications [50 points]

8) We generally think of sourcing encoding and channel encoding as very different things. **[5 points]**
   a) Define those terms in your own words. **[2]**
   Source codes compress naturally redundant messages for efficient storage/transmission. Channel codes systematically add redundancy to enable error correction and thus high-rate transmissions because we can correct errors that might occur due to sending at such a high speed.
   b) One of them generally adds bits to the data and another one removes them. Explain how doing each of those things is useful/needed. **[3]**
   Source encoding removes bits and this is helpful as it makes the data that needs to be sent smaller while keeping the overall data intact. Channel encoding adds bits to allow the receiver to error correct in the event data is lost during the transfer.

9) Using the example on slide 29 of lecture 13-15, do the following **[15 points]**
   a) Find what parity bits would be sent with the data "1011". **[2]**
   P1 = (1,0,1) = 0 P2 = (1,1,1) = 1 P3 = (0,1,1) = 0

   b) Say our receiver received the following 7-bit packets (d_1, d_2, d_3, d_4, p_1, p_2, p_3 in that order) and performed error correction. For each of the data packets, show what the 4 bits of data would be after correction (if any) **[4]**
      i) 1111000 All parity bits are wrong, so d4 flipped, so data is 1110
      ii) 0000001 Only P3 is wrong, so only P3 flipped. Data is 0000
      iii) 1100001 Only P2 is wrong, so only P2 flipped. Data is 1100
      iv) 1111001 P1 and P2 are wrong, so D1 flipped. Data is 0111
   c) Using logic gates, draw a circuit that takes d[1:4] and p[1:3] as input and generates a corrected version of d[2] as output, which we will call c[2]. **[9]**
      Hint: Break this into three parts: a) figure out which parity bits don't match (if any), b) figure out if d[2] should be flipped and c) flip it if needed. Broken down that way you should be able to manage a solution with less than 8 gates (not including negation bubbles).

      Doing it in a code notation.

      X1=XOR(d1, d2, d3, p1) -- true if P1 is "wrong"
      X2=XOR(d1, d3, d4, p2) -- true if P2 is "wrong"
      X3=XOR(d2, d3, d4, p3) -- true if P3 is "wrong"
      E2=AND(X1, not(X2), X3) – true if P1 and P3 are "wrong" and P2 is "right"
      C2=XOR(D2, E2) – flips D2 iff D2 is "wrong".

10) Define the following wireless terms in your own words: **[4 points, 1 each]**

a) Keying

Keying is a family of modulations where we allow only a predetermined set of discrete values.

b) FSK

Use different frequencies to represent different discrete values

c) QAM

Vary both the phase and amplitude to represent different discrete values.

d) PSK

Change the phase to represent different discrete values.

11) Consider the 4-ASK scheme shown on slide 39 of lecture 13. **[14 points]**

a) Show how the message "1000110101" would be communicated as a wave. Draw all sinusoids that have a positive amplitude as starting at max value and ending at the max value while those with a negative amplitude start at the minimum value and end at that same value (just as is done in the example). Have each symbol last for 3 periods. **[5]**
(Answer coming a bit later)

b) In the part above, we had each symbol last for 3 periods. If you are communicating at 1 Mbit and using 4-ASK on a 2.4GHz carrier, about how many periods would there really be? Ignore all overhead (including overhead from synchronization, error correction, and protocol). **[4]**
Tricky. Notice it is 1 Mbit per second, not 1 Msymbols per second. So we've got 500,000 symbols per second. That means a given symbol will be there for 4800 cycles ("periods") per second.

c) Most schemes that rely on phase or amplitude require a message start with a fixed signal to allow for synchronization. Why do you suppose this is needed? Think about both phase and amplitude issues. **[5]**
For phase, we need to agree exactly what makes for a 0 phase. This is arbitrary so is required. For amplitude, we are doing wireless communication, so the receiver is getting a much weaker signal than the sender is sending. And so the receiver needs to scale the amplitude to figure out what counts as each value (high, low, in-between).

12) Consider the 4-PSK scheme as shown in the figure on the right. Show how the message "00110111" would be communicated as a wave. Draw each sinusoid as lasting for exactly 3 periods and assume a sinusoid with a phase of zero degrees would be exactly a sine wave starting at zero and rising (as shown).

[6 points]

(Answer coming a bit later)

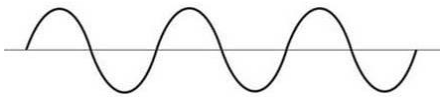

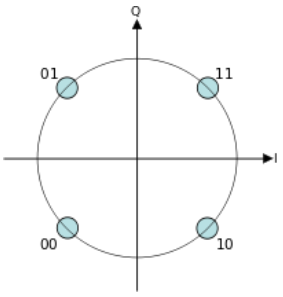

13) Consider the XBee modules as described at https://www.sparkfun.com/datasheets/Wireless/Zigbee/XBee-Datasheet.pdf **[11 points]**

a) What are the primary differences between the Zigbee PRO and Zigbee in terms of transmit power and receiver sensitivity? **[3]**

XBEE pro use far more transmit power than a regular XBee. Receiver sensitivity also differs by 8 dBm

b) Assuming they are used at the 2.4 GHz range and using a Monopole (Integrated whip) antenna on both sides

i) What is the theoretic range at which a Zigbee PRO could send to a Zigbee PRO? Show your work. **[6]**

$$r = \frac{10^{(p_t+g_t+g_r-p_r)/20}}{41.88*f} \qquad r = \frac{10^{(18+1.5+1.5+100)/20}}{41.88*2400} \qquad r = 11.16 \text{ km}$$

ii) How does your answer to i) compare to the Zigbee specification? **[2]**

The spec says outdoor with line of site is about 1.6km making the theoretical significantly higher. As mentioned in class, we generally expect this formula to be off by about a factor of 4 to 10. And that is the case.

Switching supplies [15 points]

14) Switching supplies **[7 points]**

a) What are the pros and cons of switching power supplies compared to linear regulators? **[4]**

Switching:

- **Pros: Often more power efficient (especially when dropping large voltages). Topologies are capable of stepping up a voltage or inverting it**
- **Cons: Switching induces noise on the PDN, requires tuning the capacitors and switching frequency and is therefore more difficult to get correct, and it is typically more expensive.**

Linear regulator:

- **Pros: Inefficient in comparison to switching regulators if there is a large voltage drop; only capable of stepping down a voltage**
- **Cons: Voltages are more stable and noiseless, less external components and easier to set up**

b) Explain the terms boost, buck and boost-buck as they apply to switching regulators **[3]**

Boost: Increase the output voltage of the regulator
Buck: Decrease the output voltage of the regulator
Boost-Buck: Can increase or decrease input voltage to specified output voltage (or invert it)

15) Answer the following **[8 points, -3 per wrong/blank answer, minimum 0]**

a. A(n) ***LDO / buck / boost / buck-boost*** converter which converts 10V to 7V can never be more than 70% efficient.

b. Figure #1 is a(n)
LDO / buck / boost / buck-boost converter.

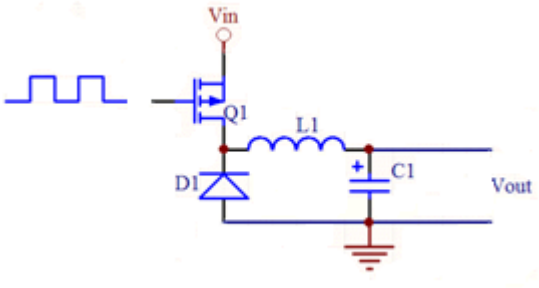

Figure 1: A voltage converter

c. What is the purpose of the diode in figure 1?
- ***It prevents the capacitor from getting too high of a voltage.***
- ***It gives the inductor someplace to draw current from when the transistor is off.***
- ***It stores energy to provide power when the transistor is off.***
- ***It controls the transistor's switching rate.***

d. An otherwise ideal buck converter which has a quiescent current of 5mA and otherwise has an efficiency of 90% will pull _____mA from its 12V input if it is generating 100mA at 5 volts. (Round to the nearest mA)

5mA+(100mA*5V/12V)/.9=51.29 or 51mA.