EE273 Digital Systems Engineering
Midterm Exam

October 26th, 1998

(Total time = 120 minutes, Total Points = 100)

Name: (please print)_____

In recognition of and in the spirit of the Stanford University Honor Code, I certify that I will neither give nor receive unpermitted aid on this exam.

Signature:_____

This examination is open notes open book. You may not, however collaborate in any manner on this exam. You have two hours to complete the exam. Please do all of your work on the exam itself. Attach any additional pages as necessary.

Before starting, please check to make sure that you have all 7 pages.

1		**24**
2		**20**
3		**20**
4		**16**
5		**20**
Total		**100**

Problem 1: Short Answer (24 Points: 6 questions, 4 points each)

A. Suppose you have a *parallel plate* transmission line (a pair of flat conductors that are wide enough that you can ignore the fringing fields – their capacitance is well approximated by the parallel plate component) with a characteristic impedance of 100Ω and a velocity of 1.5×10^8m/s. If you double the spacing between the two lines. What happens to the velocity of the line?

The easy way to solve this problem is to remember that velocity depends only on the relative permitivity of the insulator, v=c/sqrt(ε_r), and thus does not change with the geometry of the line. A harder method is to solve independently for the capacitance and inductance of the new line.

Since the capacitance is parallel plate, by doubling the spacing between the lines the capacitance per unit length C, is halved. In a homogeneous medium, we know that LC is a constant, thus if C is halved, L is doubled. We can also se that the inductance is doubled because the area of the "loop" is doubled and L = BA/I increases with A. So, since C is halved and L is doubled, the velocity of the line $v = \frac{1}{\sqrt{LC}}$ ***stays the same****.*

B. For the transmission line of (A), what happens to the impedance when the spacing between the lines is doubled?

From the last question we saw that L is doubled and C is halved. So the impedance of the line $Z = \sqrt{\frac{L}{C}}$ *is* ***doubled.***

C. Consider a 2Gb/s signaling system. At 1GHz the transmission line used by the system has an attenuation of A=0.2. There is essentially no attenuation, A=1, at DC. To overcome the high-frequency attenuation, the system uses a two-tap FIR filter for equalization. The filter takes an input stream with values +1 or –1 (for logic 1 and 0) and generates an equalized waveform on the line. To first approximation, what should the value of the second tap of the filter be if the value of the first tap is 1.0?

Here the attenuation of 0.2 means the value of the signal that is left is 0.2.

To properly equalize the line, the response to a long string of "1" bits (which won't be attenuated at all) should be 0.2 times the response to a single "1" bit (which will be attenuated to 0.2 times the DC value) so the equalizing filter equalizes the high-frequency attenuation by attenuating the low-frequency signal components by an equal amount. Let the value of the second tap be -w. Then the response to a single pulse (last bit different from current bit) is $x = (1 + w)$. *Thus we want the response to a repeated pulse (last bit the same as current bit to be 0.2x. This response is given by* $0.2x = (1 - w)$. *Solving these two equations gives us w* =0.67

D. In an on-chip RC transmission line using optimal repeater spacing, if the resistance per unit length of the wire is increased by a factor of 4 (and the repeater spacing adjusted to still be optimal), by what amount does the delay per unit length change?

The velocity $v = 1.3 * (t_b * RC)^{-0.5}$ *where R and C are the resistance and capacitance per unit length. If R goes up by a factor of 4, then the velocity is halved and the delay per unit length is* ***doubled.***

E. You wish to connect a 50Ω transmission line to a 75 Ω transmission line. Sketch a resistor network that you could place between the two lines that will allow waves to propagate from one line to the other in just one direction (from the 50Ω line to the 75Ω line) without reflections.

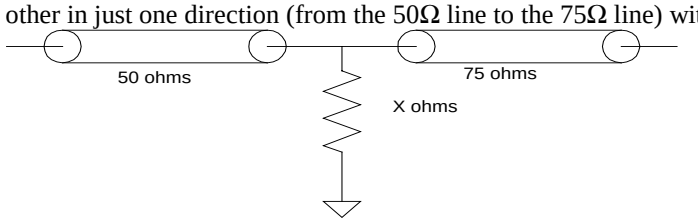

In going from left to right we write the following equation for no reflections, $50 = \frac{75x}{(x+75)}$ *from which we get x* = 150Ω.

F. Your noise analysis of a signaling system shows that with 10mV of Gaussian noise it has a BER of 10^{-9}. If you add a second 10mV Gaussian noise source, what will the BER be?

With one noise source, we know that the BER is given by.

$$\text{BER} = 10^{-9} = \exp\left(\frac{V_M}{10mV}\right)^2 = \exp\left(\frac{V_M{}^2}{100mV^2}\right)$$

where V_M *is the net noise margin of the system.*
Doubling the noise (in the RMS sense) gives us a new BER of

$$BER = \exp\left(\frac{V_M{}^2}{10^2 mV + 10^2 mV}\right) = \left[\exp\left(\frac{V_M{}^2}{200mV^2}\right)\right] = \exp\left(\frac{V_M{}^2}{100mV^2} * 0.5\right) = \left[\exp\left(\frac{V_M{}^2}{100mV^2}\right)\right]^{0.5} = 10^{-4.5}$$

so: $BER = 10^{-4.5} = 3.16 * 10^{-5}$

Problem 2: Transmission Lines (20 Points)

Consider the pair of coupled transmission lines shown below. The coupled section of each line has a capacitance of 20pF/m to the other line and 80pF/m to a shared ground plane. Each line also has a self inductance of 400nH/m and a mutual inductance to the other line of 80nH/m. The aggressor line is driven by a 1V step source with an output impedance matched to the line and a rise time of 100ps. The far end of the aggressor is shorted to the shared ground plane. The victim line is terminated with matched impedances at both ends.

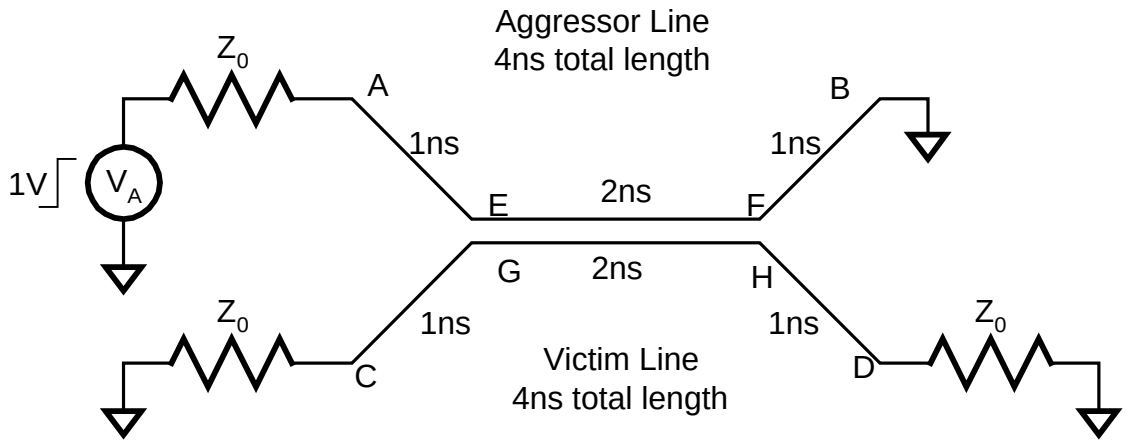

From the data above, we can calculate the following:

C_o	8.0E-11
C_d	2.0E-11
L	4.0E-7
M	8.0E-8
C	1.0E-10
Z_0	63.2
Z_even	77.5
Z_odd	51.6
k_cx	0.20
k_lx	0.20
k_rx	0.10
k_fx	0.00

Using this information, sketch and dimension the voltage waveforms at the near-end and far-end of the aggressor and victim lines.

The aggressor first ***causes near-end cross-talk at C and G****. After being reflected at B (short) the negative wave causes* ***near end cross talk at H and D****. Since k_fx is zero, there is no far-end cross talk anywhere. From voltage division we get the magnitude of the launched wave at A as 0.5V. Since k_rx is 0.10 the near end cross talk at C and G is* ***0.05V****. Similarly the reflected negative wave at B is –0.5V and causes near end cross talk of* ***–0.05V*** *at D and H. In each case the near end cross talk lasts for 4ns or twice the time of flight of the 2ns coupled section. The waveforms below show the voltages at points A, B, C, D, E, F, G and H from 0ns to 10ns. Each voltage has a 100ps rise /fall time.*

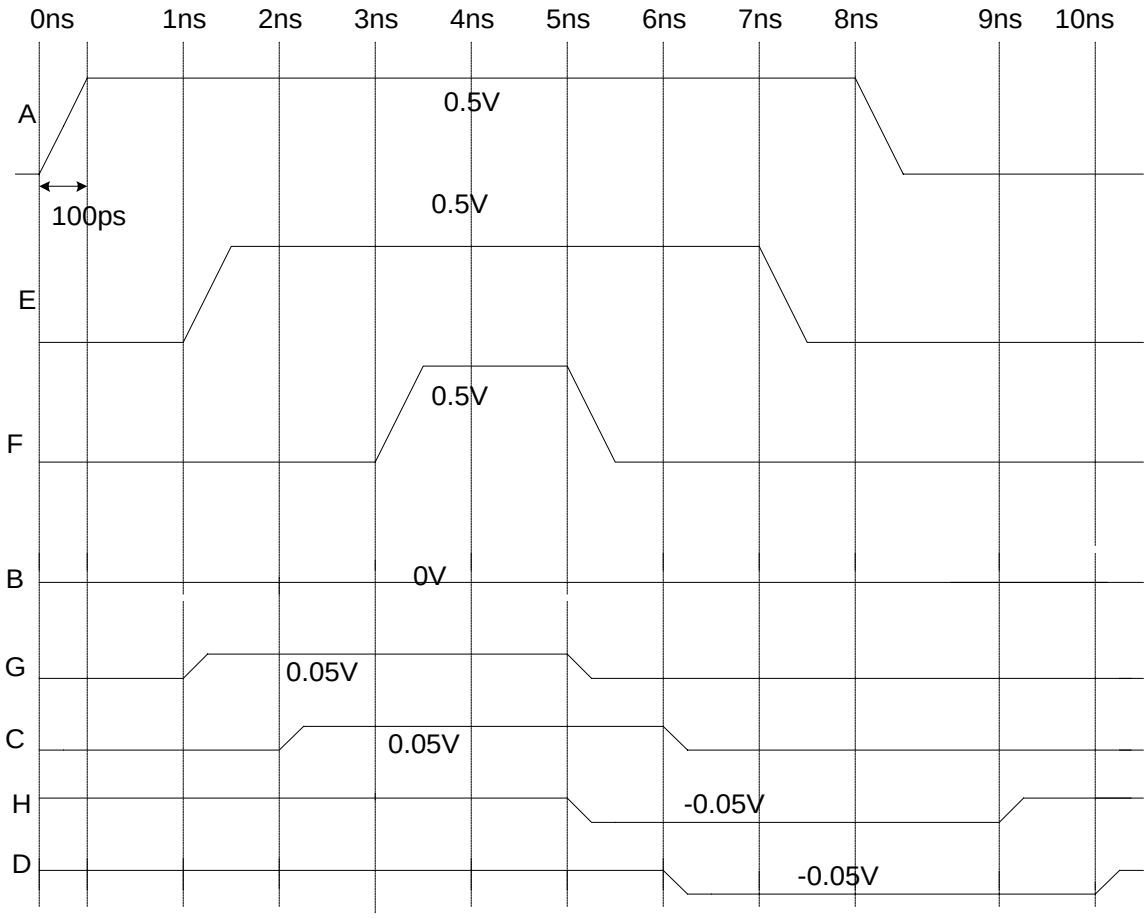

Problem 3: Signaling and Noise Analysis (20 points total)

Consider the bipolar current-mode signaling system shown below. A logic "1" is represented with 1mA of current drive and a logic 0 is represented with -1mA of drive. At present the line is terminated at both ends with matched impedances. The question is whether to (a) leave the system as is, (b) to remove the source termination, or (c) to remove the receiver termination.

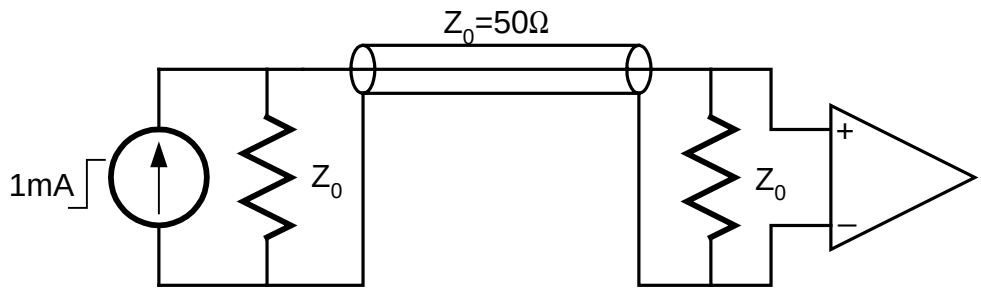

To help in deciding between these three options you have the following facts (in addition to the information above):

1. The near-end crosstalk coefficient is k_{rx}= 10%. The far-end crosstalk coefficient is zero.
2. The termination impedances are perfectly matched.
3. The attenuation of the line at the frequency of interest is 0.95.
4. The receiver offset and sensitivity combined are ± 20mV.

A. [5 Points each (15 Points total)] Compute the ratio of net margin to gross margin for each of the three termination alternatives.

For this kind of problems, it's best to construct a table:

	Option a)	Option b)	Option c)
ΔI	1mA	2mA	2mA
ΔV	50mV	100mV	100mV
K_{rx}	0	0.1	0.1
Attenuation	0.05	0.05	0.05
K_N	0.05	0.15	0.15
V_{NI}	20mV	20mV	20mV
Gross Margin	25mV	50mV	50mV
V_N	22.5mV	35mV	35mV
Margin Ratio	0.1	0.3	0.3

Note that for option c), even though the magnitude of the forward wave is 1mA, the net effect at the termination end is 2mA, since it's open circuited. So, ΔI and ΔV are identical for options b and c. Attenuation factor of 0.05 directly figures in K_N. V_N is simply $V_{NI} + K_N \Delta V$.

Many people was confused with attenuation factor. The table above is absolute worst case analysis, but another approach would be to see how the signal swing gets attenuated at the end of the line, and get the gross margin from that figure. This is also a valid approach, but is more tricky to be consistent. Partial credits were given to those who took this approach, depending on how valid their assumptions are.

Many people also had trouble with the crosstalk factor, k_{rx}. The problem states that there's no far-end crosstalk; therefore, only near end crosstalk has to be considered. Also, what's important is how this near end crosstalk shows up at the termination end. For option a, the crosstalk shows up only at the source end, and no reflection results from this crosstalk. Therefore, there's no crosstalk

term for option a). For option b), where there's no R_S, k_{rx} of the step voltage moves toward the source end, just like in option a). But unlike option a), all of this crosstalk voltage gets bounced back toward the termination end. Therefore, we have k_{rx} factor for option b). For option c), it's a bit more tricky. As shown on the table, we decided to use 100mV as ΔV, even though the initial forward incident wave is only 50mV. This forward wave of 50mV doesn't cause any crosstalk, since all the near end crosstalk voltage gets absorbed at the source end. But the reverse traveling wave of 50mV does cause crosstalk voltage of k_{rx}(50mV) to move toward the termination end. BUT, this voltage also doubles at the termination end, since it's open circuit, and the total crosstalk voltage is k_{rx} * (100mV). Therefore, k_{rx} is added to K_N factor. Some used $2k_{rx}$ for option c), but that's mistaken, since if you're using $2k_{rx}$, you're using ΔV = 50mV.*

B. [2 Points] Which of the three alternatives would you choose?

Full points to those who chose whatever the option with highest margin ratio found in part a).

C. [3 Points] Which alternative would you choose if you had a perfect receiver (0 offset and 0mV sensitivity)?

If you recalculate the margin ratio for the three options with V_{NI} = 0,

	Margin Ratio
Option a)	*0.9*
Option b)	*0.7*
Option c)	*0.7*

Therefore, Option a) is the best choice when independent noise is zero. This conceptually makes sense. Option b) and c) were ways to overcome fixed noise by amplifying the signal magnitude, at the expense of increased K_N. When V_{NI} is gone, whichever has lower K_N wins, and option a) has the lowest K_N because of no crosstalk term.

Problem 4: Basic Signaling (16 Points Total)

In an attempt to reduce signal return crosstalk, Ben Bitdiddle has devised the driver circuit, shown in the drawing below. Ben claims that a large number of these drivers can share a single pair of return I/O pads without excessive signal return crosstalk. The driver is a ±2mA bipolar current-mode driver. Ben has split the 2mA current source into two 1mA current sources. The top side of both sources are connected to the signal line, **sig**. One source takes its return current through the **ret_a** return I/O pad and is also connected to the termination resistor. The other source takes its return current through the **ret_b** I/O pad. The **ret_a** and **ret_b** pads are shared by a number of drivers. These two returns are tied together and to the transmission line return off chip as shown.

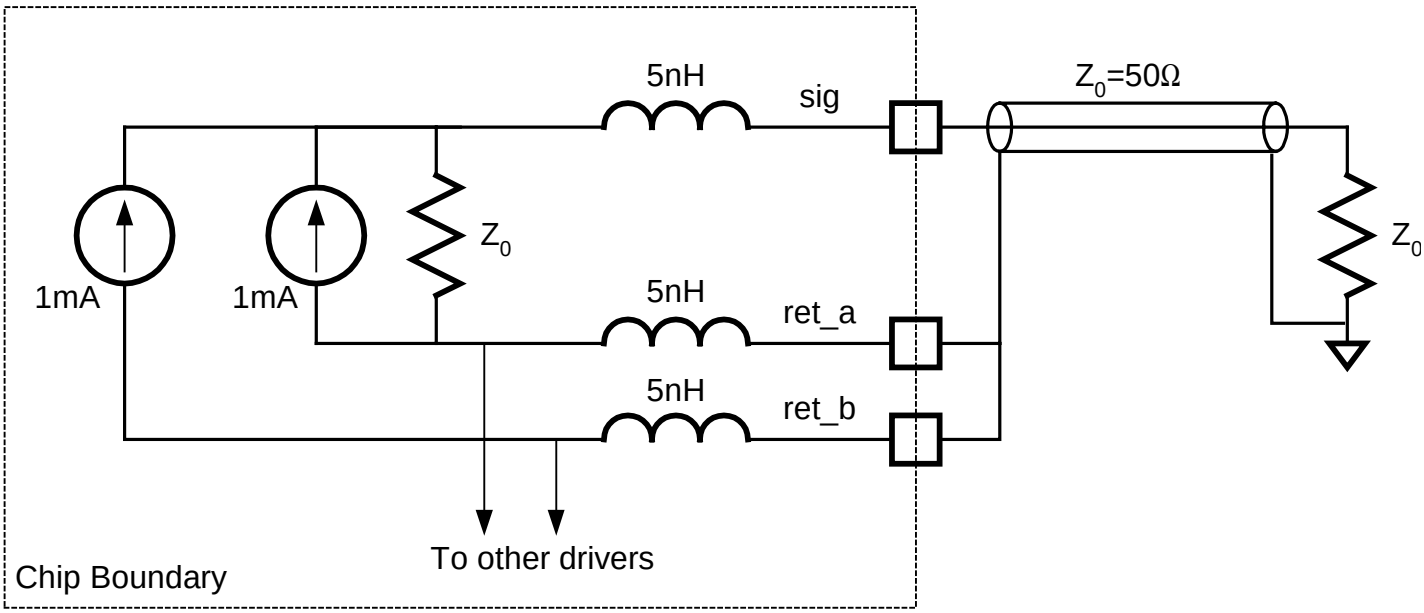

Assume the following:

1. The current sources switch in t_r=200ps with a linear current ramp.
2. The current sources have infinite output impedance but cannot operate with more than a 1V drop across their terminals.

- *General Comment about Grading:*
  This problem tended to yield very wide variation on people's scores, since part A), which is worth 8 points, was graded pretty much all or nothing. If you didn't get part A) right, part B) and C) got considerably more complicated. I tried to see what you did and give full score if your approach is correct.

A. [8 Points] What fraction of the return current flows through each of the return leads?

Combined current of 2mA flow out of the current sources, but since Z_0 is matched to the line, half of the current(1mA) flows through Z_0. So, basically, no current flows through ret_a, and all the remaining current (1mA) flows through ret_b. If you saw this, you got full 8 points. If not, you got 0. Many people got confused about 5nH inductor on the signal path. It's true that at the beginning of the transition, some current flows through the ret_a. But since the rise time of the step is 200psec and ramps up linearly, and the time constant for the inductor is 5nH/10Ohm = 50psec < 200psec, this inductor doesn't affect things too much.

B. [4 Points] What is the signal-return crosstalk coefficient for this arrangement?

Since all the current flows through ret_b, we only need to worry about voltage across bottom inductor. But since this one is in series with current source with infinite output resistance, it doesn't have any effect on the output and crosstalk coefficient is 0.

Some partial credits are given to those who realized that voltage across inductor on ret_b has no effect and the fact that the actual crosstalk coefficient depends on the ratio of current going through ret_a.

C. [4 Points] Ben's system has 100 such I/O drivers. How many total return pins does he need?

*The voltage across the inductor on ret_b has to be smaller than 950mV. It's not 1V since the signal line can move up by 1mA*50Ohm = 50mV with respect to ret_b. First, solving for the number of signal wires supportable by one ret_b, Note that since we're using bipolar current source, the ΔI is 2mA.*

$$L\frac{di}{dt} = 5nH \times \frac{2mA}{200\,ps} = 50\,mV \qquad \rightarrow \qquad \frac{950mV}{50mV} = 19\,signal\,lines/ret\_b$$

$$\frac{100\,signal\,lines}{19 signal\,lines/ret\_b} = 5.26\,ret\_b \quad \Rightarrow \begin{cases} 6\,ret\_b \\ 1\,ret\_a \end{cases}$$

Therefore, we need total of 7 return pins.

Problem 5: Advanced Signaling (20 Points)

Betty Bitdiddle suggests that simultaneous bidirectional signaling can be implemented without the need to subtract the outgoing signal. Instead she suggests using a receiver that detects three signal levels (H,M, or L) corresponding to the three voltage levels on the line. The H state always decodes to a received 1, and the L state always decodes to a received 0. The M state decodes to the opposite of the data being sent – viz. if the unit is sending a 0, the M state decodes to 1 and vice-versa.

Suppose you are building a simultaneous bidirectional signaling system with a 5mA bipolar current drive at either end. Compare Betty's proposed receiver with the one discussed in the book and lectures. Assume the following:

1. The transmission line is 50Ω and is terminated at both ends.
2. The dual-threshold receiver for Betty's scheme has the thresholds set at the optimal points – midway between L and M, and midway between M and H.
3. All resistors are ±10%.
4. Crosstalk is negligible.
5. The receivers have ±10mV offset voltage and ±10mV (20mV total) sensitivity.
6. The subtractor has an error of ±10%.

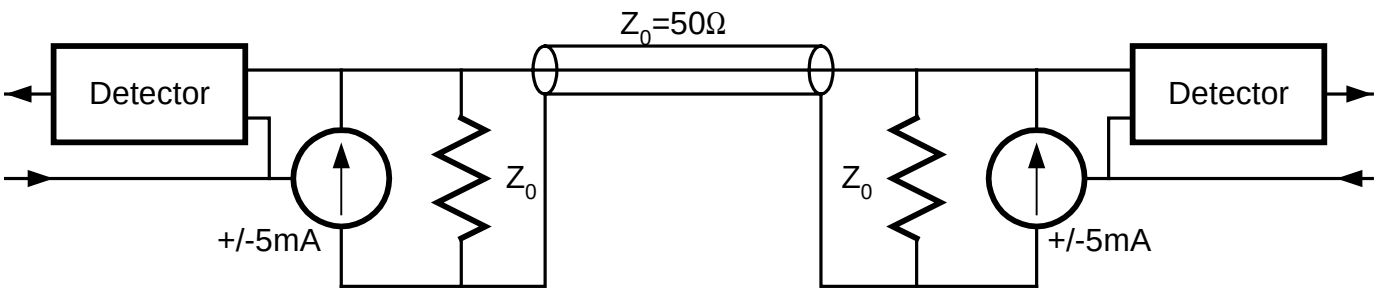

The important thing to realize here is that the two systems are nearly identical. Both are described by the figure above. They have identical transmitters and differ only slightly in the box labeled "detector". They are both ***binary*** *signaling systems.*

Consider each transmitter independently. When signaling a "1" the transmitter drives 5mA into 25Ω (a 50Ω line in parallel with a 50Ω terminator) for a voltage of +125mV. The termination at the far end absorbs all of the signal (ignoring the mismatch for now!). Similarly a "0" results in a –125mV level.

Now when we consider both transmitters operating together we get three possible signal levels. If they both drive 1s or 0s, the signals add and we get +/–250mV. If they both drive 0s the signals cancel and we get 0V.

The detector compares the received signal to a threshold which is 0 (the threshold for the bipolar signaling system) plus the magnitude of the transmitted wave (+/–125mV). The systems differ only in that the "book" system generates this reference on the fly while "Betty's" system compares the received signal against both possibilities and then selects the correct one based on the signal being transmitted. The two alternatives are depicted below.

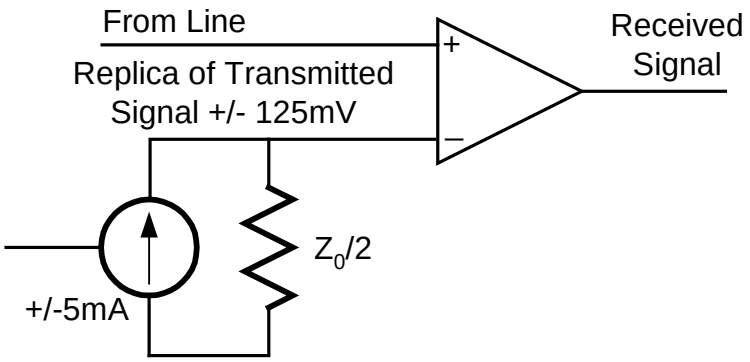

Book System

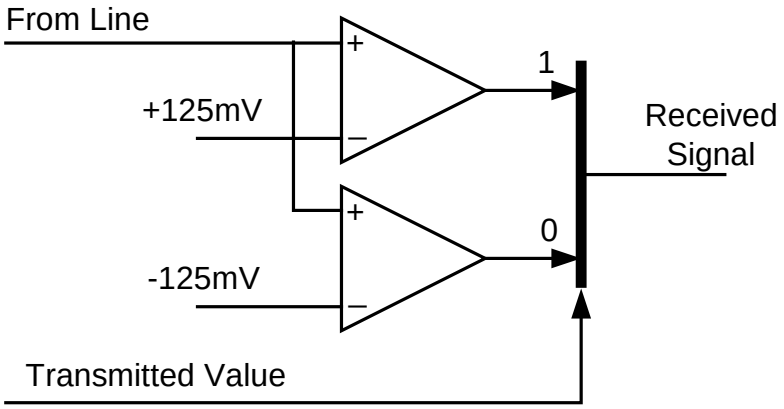

Betty's System

A. [4 points] How many comparators does Betty's system require at each end and at what voltages should their thresholds be set?

Betty's system needs two comparators, one midway between L and M, and one midway between M and H.

The three signal levels are H=250mV, M=0V, and L= –250mV. To see this note that when both drivers drive a "1" each puts 5mA into a system (10mA total) with a resistance of 25Ω to ground. Thus the comparators are set at +125mV and –250mV.

B. [4 Points] What is the gross noise margin of each system?

Recall that the ***gross*** *noise margin is the amount of noise that it takes to cause an error.*

In Betty's system each state (H, M, and L) is separated from the nearest comparator threshold by 125mV, so 125mV of noise in the wrong direction will cause an error. Hence the gross margin for this sytem is 125mV.

For the system in the book the situation is exactly the same. The same three states exist and the single threshold is always 125mV from the nominal state so the gross margin of this system is the same, 125mV.

C. [4 Points] What is the net noise margin for each system given the parameters above?

	Betty	Book
Forward Match	0.05	0.05
Reverse Match	0.05	0.05
ISI From Reflection	0.05	0.05
Subtractor		0.1
K_N	0.15	0.25
Delta_V	0.25	0.25
KN*DeltaV	0.0375	0.0625
V_ni	0.02	0.02
V_N	0.0575	0.0825
V_gm	0.125	0.125
V_nm	0.0675	0.0425
V_nm/V_gm	0.54	0.34

The spreadsheet above calculates the proportional and independent noise sources for each system. The systems are the same in that the termination mismatch causes ***three*** *separate noise effects. First, the ±10% variation in the termination at the transmit end causes the wave launched into the line to vary by ±5% (this is labeled forward match above). Second, the variation in the terminator at the transmit end of the line also causes a ±5% variation in the level of the reverse channel signal that sums with our transmitted signa (reverse match)l. Finally, the variation in the resistor at the far end of the line reflects ±5% of a previously transmitted signal back to us (ISI from reflection).*

The only other proportional error is the 10% variation in the subtractor. This is due to the resistor that generates the replica (some people counted it twice, once for the resistor and once for the comparator, with no loss of credit). This error affects only the "book" system, not "Betty's" system.

Thus, our total proportional noise is 15% for Betty and 25% for the book. Note that all of these proportional noise sources are caused by only one wave (the forward or the reverse), so they get multiplied by the single-wave swing, 250mV, not 500mV.

The only source of independent noise is the receiver offset and sensitivity which is 20mV.

D. [4 Points] Which system would you choose?

Betty's has the better noise margin because it avoids the 10% error in the subtractor but is otherwise identical.

E. [4 Points] Is there any advantage to the other system? If so, explain in no more than 2 sentences.

The book system has several practical advantages including: (a) it does not need two precise voltage references (that would probably be 10% in error anyway). (b) it only needs one comparator and thus puts less capacitive load on the line and uses less area, (c) the book design is less susceptible to glitches caused by timing skew between the transmitted signal on the line and the signal switching the multiplexer, (d) the book design doesn't need a multiplexer and hence has less delay. In practice the design in the book is almost always better, but not for the somewhat contrived circumstances of this problem .