EE273 Digital Systems Engineering
Final Exam

March 23, 2001
(version 0.3γ)

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

Name: (please print)__**SOLUTION**_____

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 11 pages.

1		**30**
2		**15**
3		**15**
4		**20**
5		**20**
Total		**100**

SITN Information here:

Problem 1: Short Answer (30 Points: 10 questions, 3 points each)

A. If a transmission line is terminated at both ends with resistors that are mismatched to the line by *x*%, what is the percentage of proportional noise due to reflections off of these mismatched resistors?

The reflection coefficient at both ends is x/2. A wave with magnitude x/2 reflects off the far end. This wave then reflects off the near end giving a wave with magnitude $x^2/4$ arriving at the receiver.

B. A signaling system has a bit-error rate (BER) of 10^{-15} and a net margin of 30mV. If you double the net margin to 60mV what will the new BER be?

BER ~ exp$(-(V_{NM}/V_G)^2/2)$. *Thus, doubling the net margin, quadruples the exponent of the BER from 10^{-15} to 10^{-60}.*

C. A brute-force synchronizer has a waiting time of 2ns and a failure rate of 10^{-4} s^{-1}. If you double the waiting time to 4ns, what will the new failure rate be? Assume $t_a f_a = 1$.

Failure rate is $t_a f_a exp(-t_w/\tau)$ or just $exp(-t_w/\tau)$ since $t_a f_a = 1$. Thus, doubling the waiting time doubles the exponent from 10^{-4} to 10^{-8} s^{-1}.

D. In a power distribution network, the effective inductance of the second rank of bypass capacitors must be kept less than 50pH to keep the voltage variation during a current transient within limits with an average current of 10A. If the average current is doubled to 20A, what is the new limit on the inductance of this rank of capacitors?

We know $L < C(\Delta V/I)^2$ so doubling I requires that L be reduced fourfold to 12.5pH

E. A bundled closed loop timing system cancels the skew between each data line and the reference clock line (true or false).

False – a per-line closed-loop timing system is required to cancel this skew.

F. If a well-designed mesochronous synchronizer is used to move a signal from one clock domain to a domain with an identical frequency but arbitrary phase, when is the only time this synchronizer should be vulnerable to synchronization failure?

The only time a mesochronous synchronizer performs a real synchronization, and hence the only time it is vulnerable is at system startup.

G. For a well-designed mesochronous synchronizer, what is the minimum data delay required to achieve a failure rate of less than 10^{-40} s^{-1}?

0

H. A plesiochronous three-register FIFO synchronizer moves a signal between two clock domains where the maximum frequency difference is 25ppm. By convention, the system puts a NULL symbol on the line at least once every 2,000 symbols. Ignoring flip-flop setup, hold, and delay times, at what relative phases must the receive pointer be adjusted when a NULL is present to ensure that a FIFO register never changes while it is selected by the output MUX?

In the 2×10^3 clocks before the next NULL arrives, the phase can drift by $2 \times 10^3(25/10^6) = 0.05$ unit circles or 18 degrees. Thus, the phase must be advanced when it falls below 18 degrees while there is a NULL on the line and it must be retarded when it goes above 240-18 = 222 degrees when there is a NULL on the line.

I. In a three-level signaling system the gross margin is what fraction of the maximum signal swing?

1/4 – there are three signal levels and two thresholds, giving four equally spaced intervals between the extremes of signal swing.

J. In an on-chip power distribution system, if there are very large local bypass capacitors near all sources of load, the on-chip distribution system must be sized for what type of current? Peak or average? (circle one)

Average – the local bypass capacitors provide the peak current.

Problem 2: Transmission Lines [15 Points]

A pair of 50-Ohm transmission lines are terminated and coupled as illustrated in the drawing below. The aggressor line is 4ns long and the victim line is 5ns long. The two lines are coupled over 2ns of their length with a near-end crosstalk coefficient of $k_{rx} = 0.1$ and a far-end crosstalk coefficient of $k_{fx} = 0$. The aggressor is driven with a 1V step with a 500ps rise time through a matched impedance at time zero. Sketch the waveform at point H.

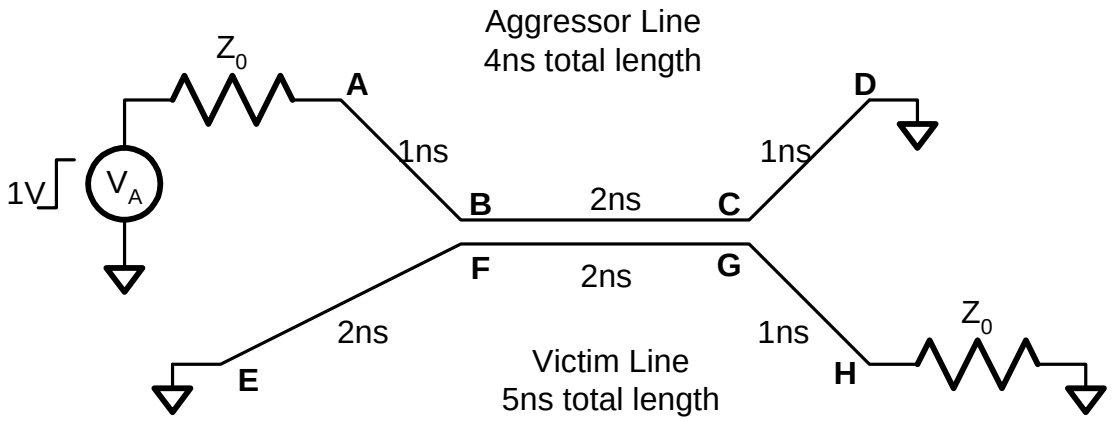

The following SPICE plots show the response of the line at each of these points:

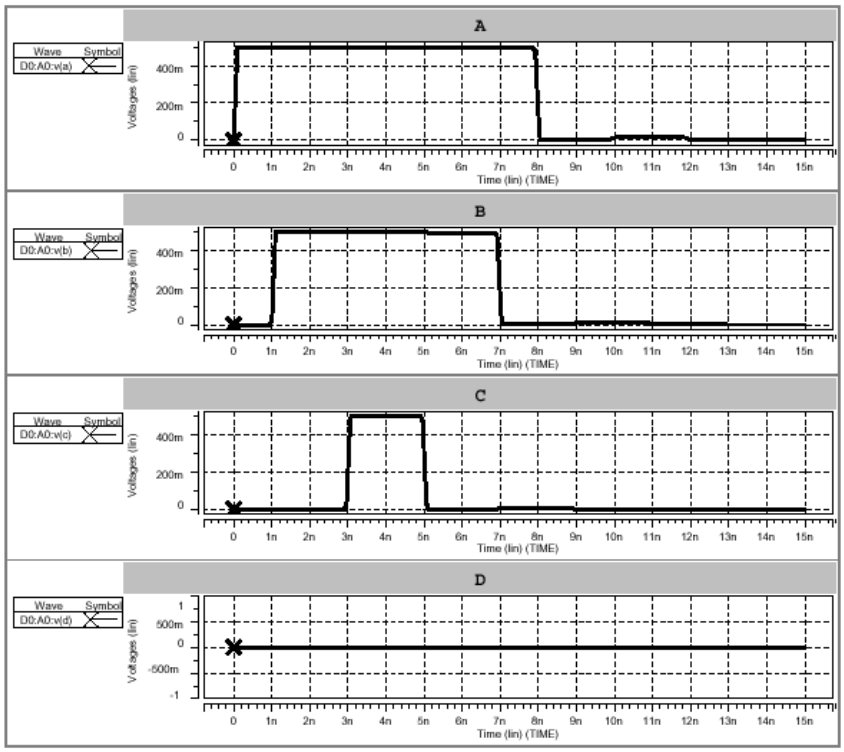

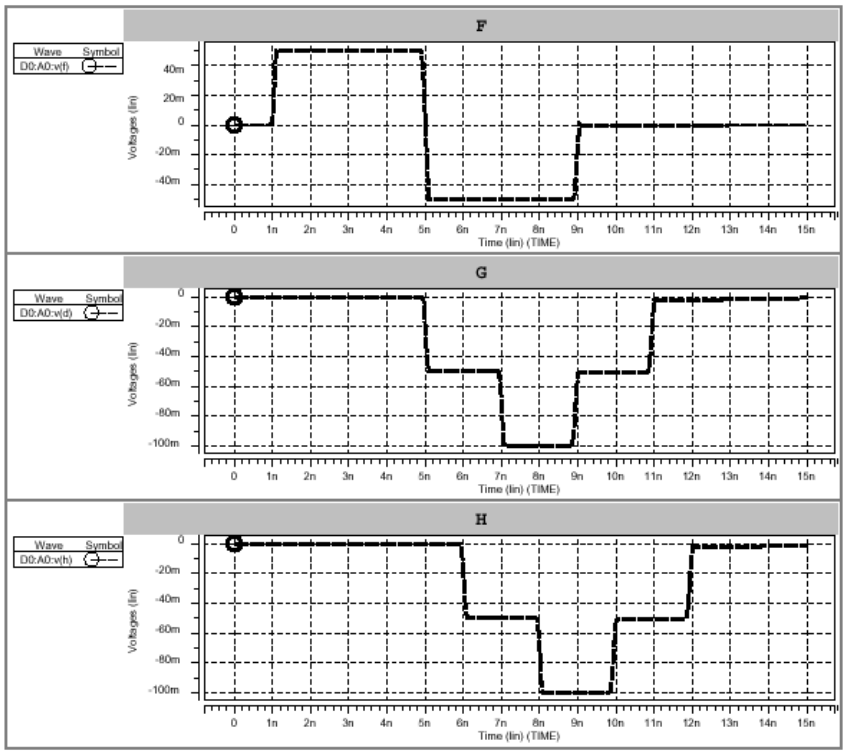
F
Wave Symbol
D0:A0:v(f)
40m
20m
0
-20m
-40m
Voltages (lin)
0 1n 2n 3n 4n 5n 6n 7n 8n 9n 10n 11n 12n 13n 14n 15n
Time (lin) (TIME)
G
Wave Symbol
D0:A0:v(d)
0
-20m
-40m
-60m
-80m
-100m
Voltages (lin)
0 1n 2n 3n 4n 5n 6n 7n 8n 9n 10n 11n 12n 13n 14n 15n
Time (lin) (TIME)
H
Wave Symbol
D0:A0:v(h)
0
-20m
-40m
-60m
-80m
-100m
Voltages (lin)
0 1n 2n 3n 4n 5n 6n 7n 8n 9n 10n 11n 12n 13n 14n 15n
Time (lin) (TIME)

Problem 3: Timing [15 Points total]

Consider the per-line closed loop timing system shown below. Signal din is clocked on both edges of a transmit oscillator and driven onto the line. At the receiver a local oscillator is phase shifted to center the clock on the eye and then used to clock data off the line to generate output dout.

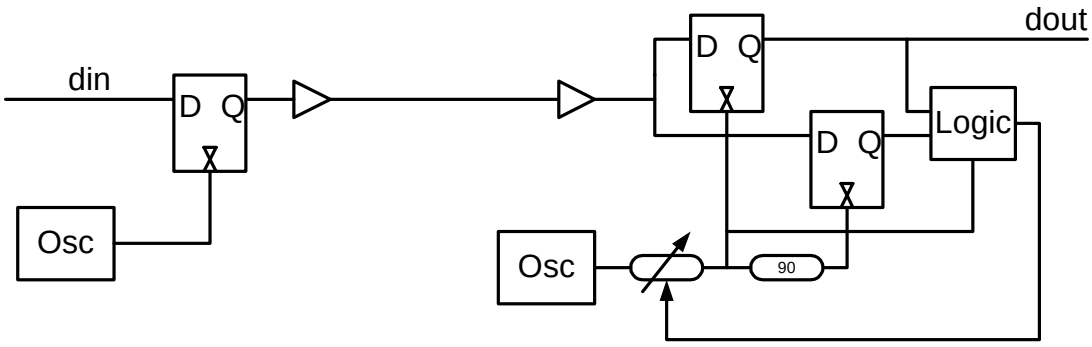

Part A [8 pts]: The oscillators both have 50ps p-p jitter. The transmission line is assumed to be ideal, and all flip flops, logic, and delay elements have skew (magnitude) that is 5% of their delay and jitter (p-p) that is 5% of their delay. The triangular symbols on the transmission line represent the transmitter and receiver. The timing properties of these elements are listed in the table below. The delay of the two delay lines is determined by the rate at which you operate the system. The transmitter rise/fall time is 100ps. You may assume that the setup and hold times of all flip-flops are exactly matched. What is the fastest rate at which you can operate this link?

Flip-flop clock-to-Q delay	t_{dCQ}	200	ps
Flip-flop setup time	t_s	100	ps
Flip-flop hold time	t_h	0	ps
Transmitter delay	t_{dTx}	200	ps
Receiver delay	t_{dRx}	200	ps

The only skew that affects the system is that of the 90-degree delay line. The +/-5% variation here consumes 5% of the bit cell (0.05UI) regardless of frequency.

Jitter sources include all elements in the path from the transmit oscillator to the receive flip flop – osc 50ps + flop 10ps + Tx 10ps + Rx 10ps = 80ps – plus all jitter in the path from the Rx oscillator to the receive flip-flop – osc 50ps + delay line 0.05UI (worst case). Thus the total timing uncertainty is 130ps + 0.1UI.

The minimum bit time 1UI is set by aperture 100ps + rise/fall 100ps + uncertainty 130ps + 0.1UI. Solving we get 0.9UI > 330ps or 1UI > 366ps which gives a frequency of 2.73Gb/s.

Part B [7pts]: The system is modified so the transmitter forwards its clock to the receiver over a line that is within +/- 100ps of the length of each data line as shown below (instead of the receiver using an oscillator). All other parts of the system are unchanged. Now what is the maximum rate at which you can operate the link?

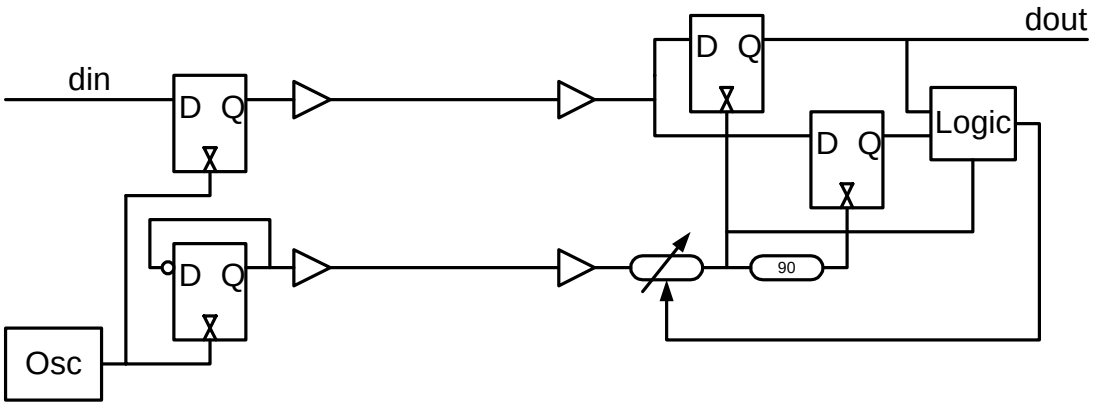

The system here is identical except that jitter in the transmit oscillator is common mode between the path to the data input of the receive flop and the clock input to the receive flop. This cancels the jitter of the leading edge of the eye – its correlated with the sampling point – but does ***not*** *cancel the jitter of the trailing edge of the eye – both the sample point and the eye end are independent samples of edge placement that can each vary by +/- 25ps for a total of +/- 50ps. The net result is that the total uncertainty due to the oscillator(s) is reduced from 100ps in part (a) to 50ps here. This is partly offset by 30ps of flop, Tx, and Rx jitter added to the clock path.*

With this change, the total uncertainty becomes 110ps + 0.1UI.

The minimum bit time is then 1UI = aperture 100ps + rise/fall 100ps + uncertainty 110ps + 0.1UI. So we get 0.9UI = 310ps or 1UI = 344ps, for a frequency of 2.9Gb/s.

Problem 4: Synchronization [20 points total]

A two-register plesiochronous synchronizer moves a signal A from the domain of tclk to the domain of rclk. A schematic for this synchronizer is shown below. A is clocked into one flip-flop by the positive edge of tclk and into a second flip-flop by the negative edge of tclk. A multiplexer controlled by the signal "select" selects which of these flip-flops to sample by the output flip-flop which is clocked by rclk.

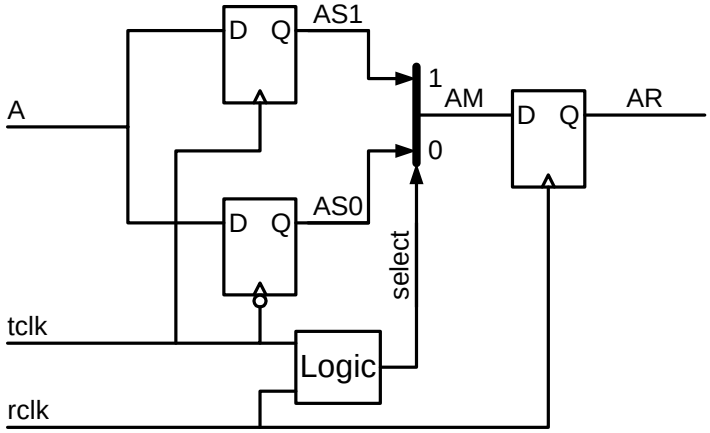

A detail of the logic block that generates select is shown below. This block samples a delayed version of tclk with rclk to decide which input of the multiplexer to select.

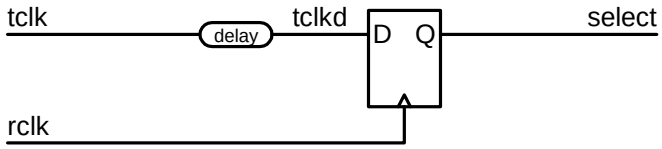

Part A [6 points]: There is something fundamentally wrong with the logic block above. What is it? Sketch a corrected logic block below. (Note, for parts A and B, assume that the signal select can switch on any symbol. It does not have to wait for a NULL on the line. We will revisit this in part C.)

This circuit is subject to synchronization failure and thus must be resampled before using the select signal to clock the multiplexer. A revised circuit is shown below. (We will assume that one rclk cycle is a long enough wait period).

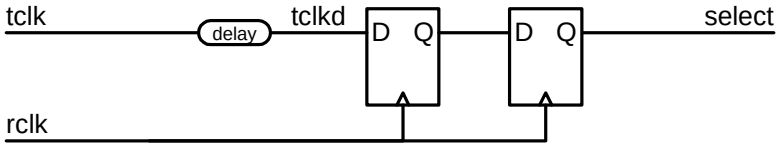

Part B [6 points]: What is the proper value for the delay element in the logic block above (as a fraction of a clock cycle)? Give the value that maximizes timing margins. To keep things simple, assume that clock-to-Q delay is zero and that $t_s = t_h$.

90-degrees (or pi/2 or 0.25UI) – This value ensures that the synchronizer switches to the other register whenever the sampling edge gets within 90 degrees of the center of the symmetrical keep-out region – from either side.

Part C [8 points]: For this part you must modify the logic block above so that no non-NULL symbol is replicated or dropped. You should restrict the switching of select as little as possible within this constraint. You may assume you have a signal available called ISNULL that indicates when the value on A is a NULL and a second signal called RSLOW that indicates when rclk is slower than tclk. You may assume that both of these signals are in the timing domain of rclk. (Hint: first identify which transitions of select are safe (don't drop or replicate a value) and which are unsafe. Then sketch a circuit that inhibits the unsafe transitions).

The key thing to realize here is that transitions that involve rclk crossing the delayed falling edge of tclk are safe. They involve switching the mux between the two registers when they are guaranteed to have the same value. Only when rclk crosses the delayed rising edge of tclk is the transition unsafe and must wait for a NULL. The rising edge is identified by a 0 to 1 transition on pending select if RSLOW is true and a 1 to 0 transition if RSLOW is false.

The logic that implements all of this is shown below. The selection is allowed to change whenever ISNULL is true – this is always safe – or when the transition is safe as defined by (old & ~new & rslow)|(~old & new & ~rslow).

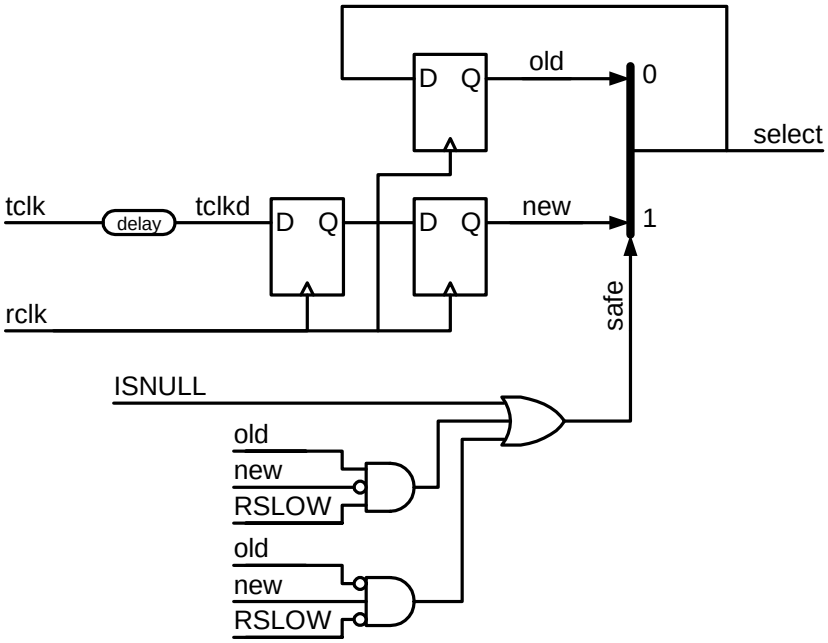
D Q
old
0
select
tclk
delay
tclkd
D Q
D Q
new
1
rclk
safe
ISNULL
old
new
RSLOW
old
new
RSLOW

Problem 5: Power Distribution [20 points total]

A single chip processor draws the following current profile:

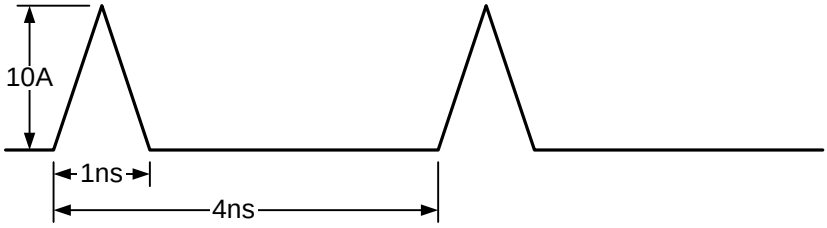

Part A [7 points]: The on-chip portion of the distribution network is modeled by the circuit shown below with 10 milli-Ohm series resistance on power and ground feeding a current source with the profile shown above. Given this load and resistive network, what is the minimum value for the local bypass capacitor, C_1, that will keep the total drop (the difference between V_A and V_L) on the supply voltage across the load less than 0.1V?

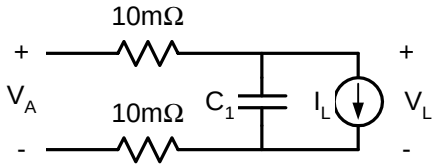

With a maximum drop of $\Delta V = 0.1$, the resistors supply a current of $\Delta V/R = 5A$. If we approximate the ramp in resistor current as linear, then the charge that must be supplied by the capacitor is the area of the shaded region in the figure below

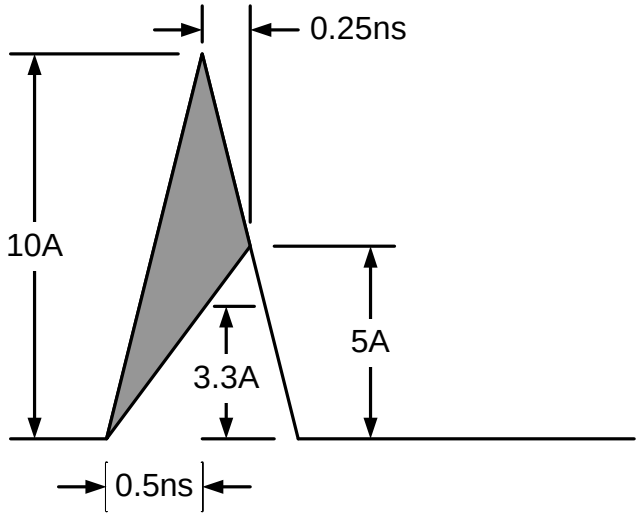

The charge is thus $Q = 0.5(6.7A)(0.75ns) = 2.5nC$ (Alternatively you can just note that this is half the area under the curve.) To keep the voltage drop to 0.1V, the Capacitor must be sized $C_1 = Q/\Delta V = 2.5nC/0.1V = 25nF$.

Part B [7 points]: Assume C_1 in the figure above is 50nF (this is **not** the answer to Part A). Ignoring all inductance except for that of the bypass capacitors, what set of bypass capacitors (types and number) would you use for the next level of bypassing? You have the following types available:

Type	Capacitance	Inductance	Resistance	Type of Cap
A	0.01μF	1nH	0	Ceramic chip
B	0.1μF	1nH	0	Ceramic chip
C	1.0μF	1nH	0	Ceramic chip
D	10μF	5nH	0.5Ω	Tantalum
E	100μF	10nH	0.5Ω	Tantalum
F	1000μF	10nH	0.5Ω	Aluminum Electrolytic

We have $L_2 < C_1(\Delta V/I)^2 = 50nF(0.1V/1.25A)^2 = 50nF(0.0064) = 0.32nH$

We choose 4 type "C" capacitors to give 4μF of capacitance and 0.25nH of inductance.

Part C [6 points]: The inductance back to the main power supply is 1μH. Can the second stage bypass you designed in Part B be directly connected to this main power supply through this 1 μH distribution inductance or is another stage of bypass capacitors required? If another stage is required, specify the type and number of capacitors to be used in this stage.

We need at least $C_x = L(I/\Delta V)^2 = 1\mu H(1.25A/0.1V)^2 = 156\mu F$ to handle a current transient out of the supply inductance. What we have from part B is not enough to handle this, so we will need a third rank. For this rank we compute

$L_3 < C_2(\Delta V/I)^2 = 4\mu F(0.1V/1.25A)^2 = 4\mu F\ (0.0064) = 25.6nH$

We also need to worry about resistance. With an average current of 1.25A a resistance of R = V/I = 0.1/1.25 = 0.08Ω uses up all of our voltage margin. Thus, with half-ohm capacitors, we need at least 0.5/0.08 = 6.25 caps to get the ESR down to this point. Thus, we choose 7 type "F" capacitors for this rank (7 type "E" would work too). The 7 "Fs" give us 7mF of capacitance (more than the 156μFminimum) and 10/7 = 1.43nH of inductance (less than the 25.6 minimum) and 0.5/7 = 0.07Ω of resistance (just less than the 0.08Ω minimum).