Full Length Cable & 15 Loads, 3 Point Plan

1.0 Introduction

1.1 What I will present is a plan to allow us to get to the goal of still having 15 loads and a cable length of 12 meters, 25 meters for point-to-point. We do not have to build really accurate implementations of these circuits to get to full load operation. Simple to moderate accuracy of these circuits will serve us fine.

The point about these techniques is that we can achieve a predictable or quantifiable amount of improvement over how we are did things for Fast-40. We can ensure that data phase will always operate at the 80 mega-transfers per second rate.

So don't worry that all of this will be too complicated, because it can be accomplished by only a fairly moderate effort. The ideas and means presented here are not intended for SCSI if we decide to limit the loads to 5, and cable length to 3 meters. The ability to do these techniques now at a simple implementation level, will give us valuable experience. Then, if later we need to do a more accurate implementation for FAST-160, we will be pro's at it.

2.0 Overview of the 3 point plan

2.1 ISI compensation

2.2 Auto-zero bias cancellation on receivers

2.3 Timing margining for both edges of ACK/REQ clocks at receiving data end

3.0 Receiver bias cancellation turn-on/off timings

3.1 Review of what receiver bias cancellation means:

Leave the voltage bias in the terminators alone, use symmetrical drivers for DATA PHASE, and have it be the responsibility of the receiver to cope with the offset voltage from the biased terminators. As one suggested solution, I recommend that you add a bias cancellation driver to your receiver to cancel-out the terminator voltage bias. Other solutions might be possible, but this is an easy one.

3.2 From page 109 of SPI-2 rev 20a, third complete sentence from the top.

That sentence leads me to believe that ACK & REQ must always be negated for 400ns after an information transfer phase change.

3.3 The timing rules for receiver bias cancellation turn-on/off are then:

3.3.1 15-50ns after the information phase changes to data phase, any ACK/REQ driver that is using bias cancellation at the driver, should turn off it's bias cancellation. If the ACK/REQ driver are asymmetrical, then there is no issue, except that you need to ensure that next time you drive ACK or REQ asserted, you must being driving in symmetrical mode.

3.3.2 75-100ns after the phase changes to data phase, enable the bias cancellation at the receiving ends of the initiator or target.

3.3.3 125-150ns after the phase changes to data phase, the ACK & REQ drivers should start their ISI compensation schedules.

3.3.4 125-150ns after the phase changes to data phase, the data drivers should be enabled at their minimum ISI driving strength. These last two steps ensure that when these lines change again, there will not be any ISI related timing problems from the 1st pulse.

3.3.5 15-50ns after the last ACK of the data phase, disable all the bias cancellation drivers on any receiver and prepare for the next phase. Or if the last ACK doesn't come, then turn off when normal error procedures begin.

4.0 General overview of timing & error sources

4.1 Instead of trying to talk about what the timing budget is here, we will discuss what kind of system tolerances and issues that cause timing skews and what ones are addressed by the 3 point plan and what ones are not.

4.2 1st the error sources that are addressed by the 3 point plan

4.2.1 ISI or inter-symbol interference. This happens when the signals in a cable are at a high enough frequency that there is measurable attenuation in the cable and the signals cannot reach their full level before they switch again. Because of this, some signals (lower frequency) reach higher levels than others (higher frequency). Because of this amplitude difference and the fact that there is a non-zero slew rate, higher levels are going to cross the zero axis reference later than lower level signals. This is what I call "zero crossing shift" and it can add anywhere from 0-3ns to the timing skew budget. The 3ns is for cables with very high attenuation and many distributed loads. 2ns is a more typical worst case number, and that is what we will use here. And I think I remember that 2ns is presently budgeted for fast-40. Some examples of this are illustrated in figures 2A-2C. Figures 2A-2C have various combinations of cable attenuation or capacitance loading.

4.2.1.1 Because of ISI, there is a double whammy effect. ISI can cause the receiver to be switched from a very large overdrive level to very small overdrive level. This effect increases the skew in the receiver beyond what the skew would be if it had equal amplitude inputs. This problem can be overcome somewhat with costly bi-cmos inputs which are more slew rate limited than cmos inputs (bipolar has higher gain). But this would be a costly solution. Figure 3 illustrates the combined effect of ISI and receiver skew at about 4.5ns of zero crossing skew. This is for a long cable with very little capacitance.

4.2.2 Signal offset is similar to ISI in effect, in that it also causes a zero crossing shift. Because it is difficult for a driver to exactly cancel-out the voltage bias from the terminators, there is sometimes an offset in the output levels due to this. Whenever there is a non-zero slew rate on a signal edge, zero crossing skew will be introduced by this offset. Figure 4 tries to illustrate this point.

Just as in 4.2.1.1 above, receiver skew is also affected by this in the double whammy kind of way.

4.2.3 REQ/ACK drivers, receivers, cable and cable capacitance load skews.

No matter how hard you try, you can not create a perfect square wave inside an IC, transmit it down a cable, receive it and buffer it without creating some duty-cycle distortion.You will always end up with some duty cycle distortion at the receiving end.

If we examine the clock path for REQ & ACK, we will see that there are many places where duty-cycle distortion can occur. The driver will not always have exactly the same low-high and high-low delays. The board traces and package wires may be slightly different in length and capacitance. The cable may have slightly longer or shorter wires for the differential signals. Differences in capacitance for the multi-drop loads. On the receiver board, again there are trace and package wire differences. Offsets in the IC receiver and differences in low-high and high-low delays for the clock buffer and receiver.

All of these add up to create duty cycle distortion and clock skew with respect to the data. And, by with respect to the data, we need to consider the data lines with both the least and most skew with respect to the clock.

This is especially important when you are doing dual-edge clocking.

4.3 These are the error sources that I’m NOT trying to address.

4.3.1 Clock jitter. Basically, the clock edge that changes the data isn’t necessarily the same edge of the clock that generates the ack/req that is used to receive the data. It could occur at least 1 clock later.

The may have been a slight change in clock period, so this could be one source of jitter. There maybe many others, and this one may not be.

4.3.2 For the 16-18 data lines, there is skew among them. Such that some data will arrive earlier or later than others.

This is due to the same path as describe above in 4.2.3.

4.3.3 Reflections! With the test chip we built a couple of years ago. we could clearly see reflections that happen only once every few thousand clocks or so. The reflection would cause a shift in the data timing due to the vector sums of the data and the reflection. This is similar to the problem from glitches for case 4 hot-plugging.

4.3.4 Any others I missed.

5.0 Block diagram of 3 point plan

5.1 Figure 5 shows the goals of the 3 point plan. That is, reduced the total amount of timing skew by letting the receiver skew be higher for short cables and lower for long cables. Reduce ISI zero crossing skew ~75%, from 1.5-2ns to ~400ps.

5.2 A block diagram in figure 6 shows how only 1 D/A driver is needed to do both ISI compensation and bias cancellation auto-zero at the receiver. Using a simple D/A, (variable strength driver) signal offset can be reduced ~75%.

6.0 Description of an ISI compensation plan

6.1 Previously shown in figure 2 were ISI zero crossing shifts for multi-frequency signals all lined up. Each step to a lower frequency has a greater amplitude than the last. The increase in amplitude appears to slow, as each new step to a lower frequency is at a smaller fractional decrease in period. The greatest amount of zero crossing shift comes at the beginning.

If the driver current is reduced in the same proportion as the amplitude increases, then any actual amplitude increase is nullified, and the amplitude will always be at a constant value when it goes to change states. And the zero crossing times for all the different frequency signals will be the same. This is because the signal state changes all have the same swing and starting point. ISI problems are eliminated.

This compares to cable equalization, where a LC network is added to a cable to attenuate the low frequencies, such that at the end of the cable low & high frequencies have the same attenuation. This is a digital form of cable equalization.

Therefore, we have to use the same assumptions as they do with cable equalization. All drivers must ALWAYS switch with the SAME drive current, and outputs that are stable for a long time will be cut-back in drive strength to simulate low frequency attenuation. (This implies that boosting doesn't work).

6.2 Since the exact amount of cable attenuation isn't always known, we have to design the compensation amount for the maximum amount of cable system attenuation and loading.

6.2.1 One possibility to investigate is to make the compensation programmable, and adjust it during timing margining for each device, but this maybe be over-kill and have a lot of overhead.

6.3 Since we are only trying to reduce the variation in amplitude by 75%, only a maximum of 4 cut-backs should be needed in the case for 30-40% cable attenuation. This implies only a very simple 2 bit D/A is needed to accomplish this.

6.4 The most important thing to note about this, is that we are dealing with percentages, and absolute cut-back values are not important, only percentage amounts. Therefore this only a matching issue and easy to implement.

7.0 Description for auto-zero bias-cancellation

7.1 Using the D/A (variable strength driver) that was used for ISI compensation in driver mode, when in receiver mode, drive the input asserted starting at your highest value that will guarantee an asserted level and continue to cut-back until your receiver just switches back to negated. Stop at this point, and use this setting for bias cancellation on the receiver.

7.2 This may call for a 3 bit D/A. The range of output current needed to cancel out the voltage bias in the terminators is 1.8-2.5mA, plus you might need a little more or less to cancel out your input receiver's offset voltage. To cover a range from 1.4-2.8mA in 0.2mA steps, you will need 8 steps. 0.2mA steps translate into 10mv resolution, which should be good enough for reaching the goal of a combine bias & input receiver offset of 20mV.

7.3 The best time to auto-zero all your req/ack, parity and data receivers is 1 time or some time or any time you have won arbitration. When you assert both BSY and SEL stop there and don't assert ATN yet.

At this point you have the bus with nobody on it and nobody expecting anything on it yet. Do your auto-zero and then either go back to BUS FREE or put the SCSI ID's on the data bus and assert ATN and continue.

7.3.1 See SPI-2 rev 20a page 107, section 11.1.3

8.0 How timing margining works

8.1 The idea of timing margining is for the data receiver to be able to null out any duty-cycle distortion and set the REQ/ACK edges in the middle of the best range of largest and smallest data skews.

8.2 One way to do this is add a fixed amount of extra delay to all the data and parity inputs. About 1ns to 4ns should be enough. This is about a 4-to-1 range which is about the tolerance of inverter delays.

For each edge of the clock, if you are doing dual-edge clocking, add an adjustable delay element that can go from 0ns to 2x of the delay amount on the data lines. Make these delay elements have from 8-16 steps.

This should allow you to trim at least 1ns out of the skew budget.

8.3 You will notice that for fast delays, you will have less adjustment range, but since the whole IC is at fastest case, all other things are enhanced. Flip-flops will need shorter setup & hold times, your receivers will be fast and have less skew. This is good tracking.

8.4 The rest of this will have to be expanded upon at a later time.

9.0 Output driver currents

9.1 With ISI compensation and bias-cancellation, drive strengths can be reduced and power along with it. Starting backwards, the minimum output voltage level needed to maintain a high level is 150mV. This is because we already said that 125mV is good enough for the voltage bias in the terminators, and this is higher than that.

9.1.1 If the minimum level is 150mV or 3mA and we plan to allow for 40% attenuation, for example, then 250mV or 5mA is the highest current needed.

9.1.2 There is really no need to put an upper limit, but if we want one.

9.1.2.1 Let the minimum current for full ISI reduction be from 3-6mA, and the starting current will be 6-12ma, where the current always decreases as a percentage of the starting current.

10.0 Example specification for ISI compensation

10.1 ISI compensation is specified by placing a tolerance on the amplitude variations into a standard attenuation load. As soon as the system cable signal attenuation is defined, an exact specifications can be written. A preliminary can be drawn-up in the mean time. You can see from figures 2A-2C that we need to decide the ratio of cable attenuation to cable load capacitance for the test load.

11.0 Example specification for auto-zero bias-cancellation

11.1 Auto-zeroing can be specified for the bias cancellation technique, other techniques for the receivers dealing with the voltage bias from the terminators will have to specify their own measurement standards.

11.2 For bias-cancellation:

Measure the offset of the receiver. Then calibrate the receiver to the upper and lower limits and 10 places in-between. Measure the voltage offset for each load setting. Offset on the terminator load plus the offset of the receiver shall be less than 20mV.

12.0 Timing training

12.1 Later.

13.0 Mandatory or optional

13.1 If we make these items mandatory for fast-80, then there is no need to modify any commands to negotiate about whether this is supported of not. But, as a back up plan for our design, we would like to add the ability to negotiate using fast-40 style drivers/receivers when in fast-80 mode.

message → data

command → data

data → data

Figure 1

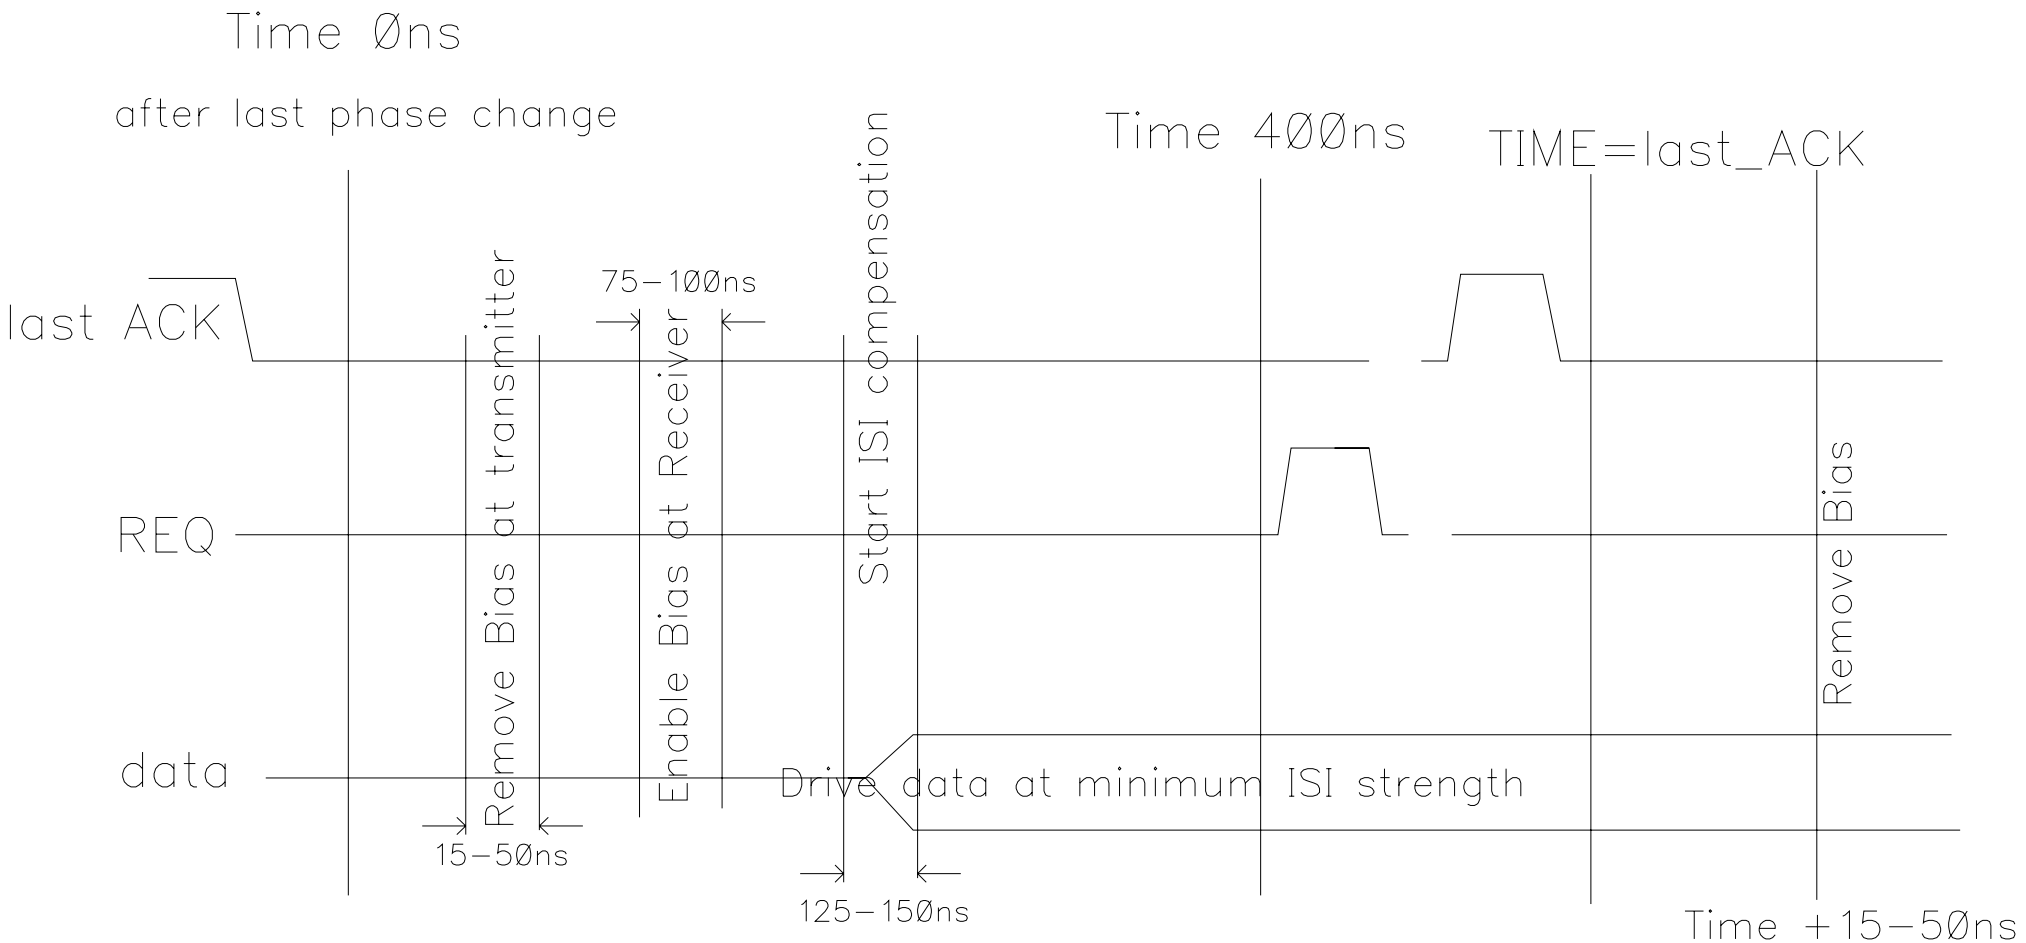

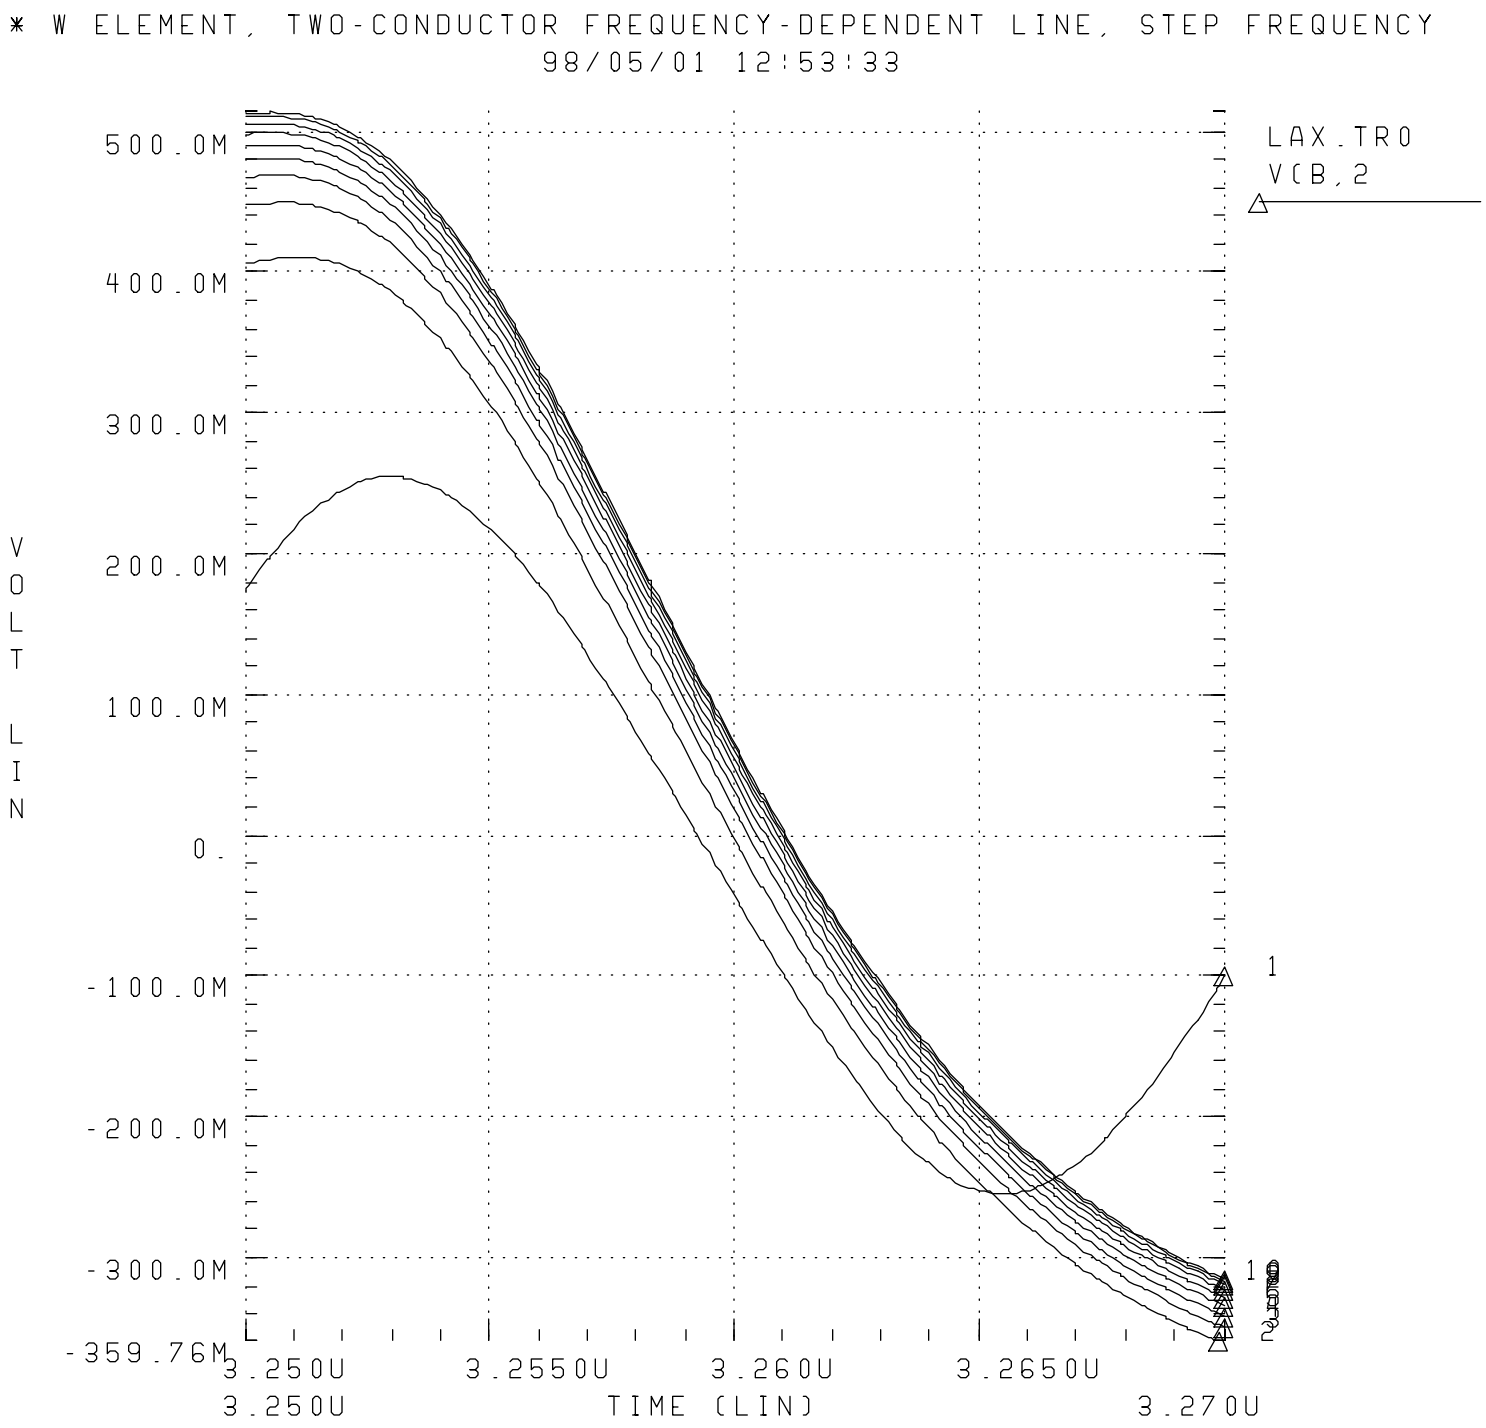
* W ELEMENT, TWO-CONDUCTOR FREQUENCY-DEPENDENT LINE, STEP FREQUENCY
98/05/01 12:53:33
LAX.TR0
V(B,2
500.0M
400.0M
300.0M
200.0M
100.0M
0.
-100.0M
-200.0M
-300.0M
-359.76M
VOLT LIN
3.250U
3.2550U
3.260U
3.2650U
3.250U
TIME (LIN)
3.270U
1
1

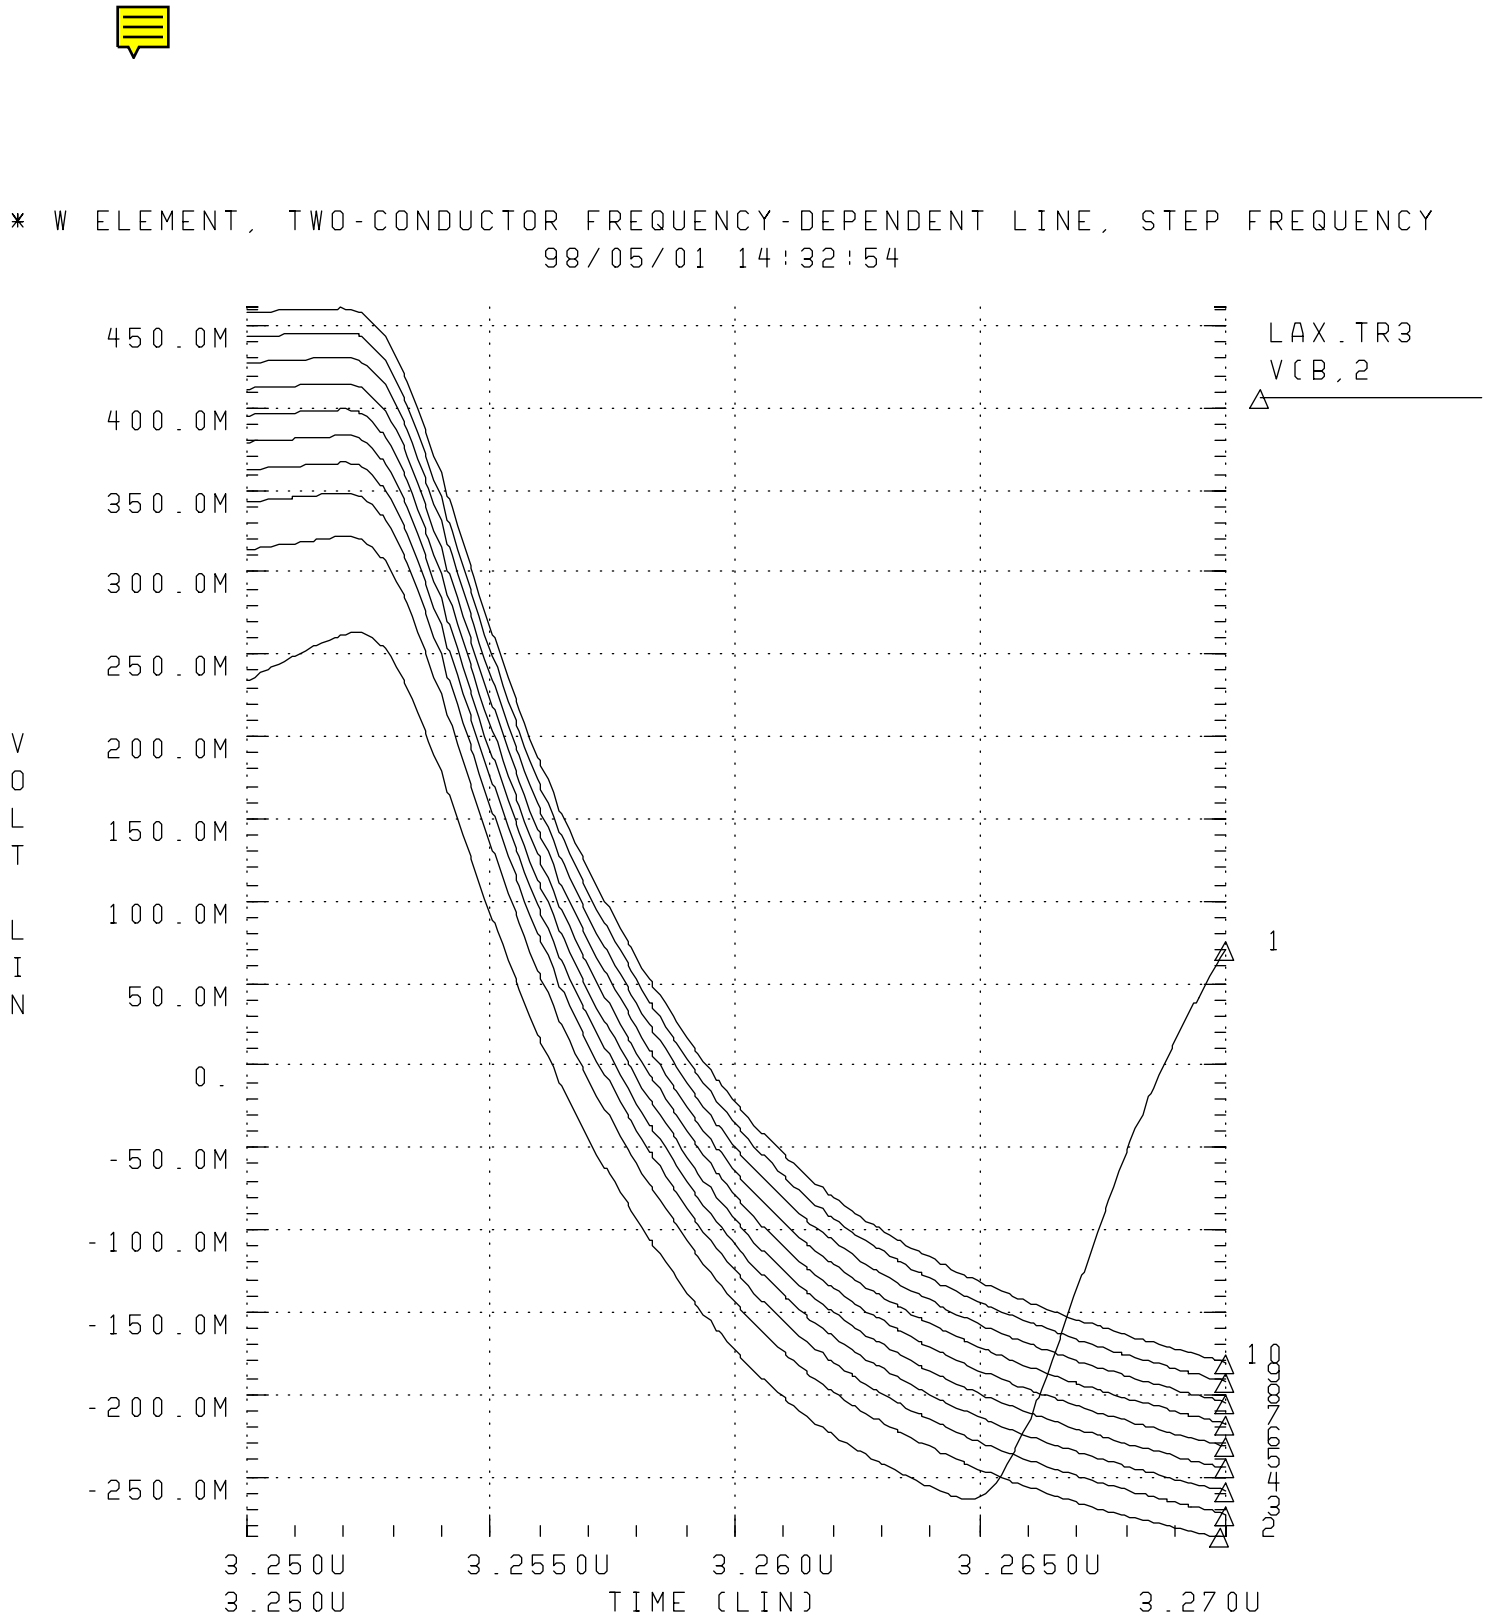
* W ELEMENT, TWO-CONDUCTOR FREQUENCY-DEPENDENT LINE, STEP FREQUENCY
98/05/01 14:32:54
LAX.TR3
V(B,2
450.0M
400.0M
350.0M
300.0M
250.0M
200.0M
150.0M
100.0M
50.0M
0.
-50.0M
-100.0M
-150.0M
-200.0M
-250.0M
VOLT
LIN
1
1
3.250U
3.2550U
3.260U
3.2650U
3.250U
TIME (LIN)
3.270U

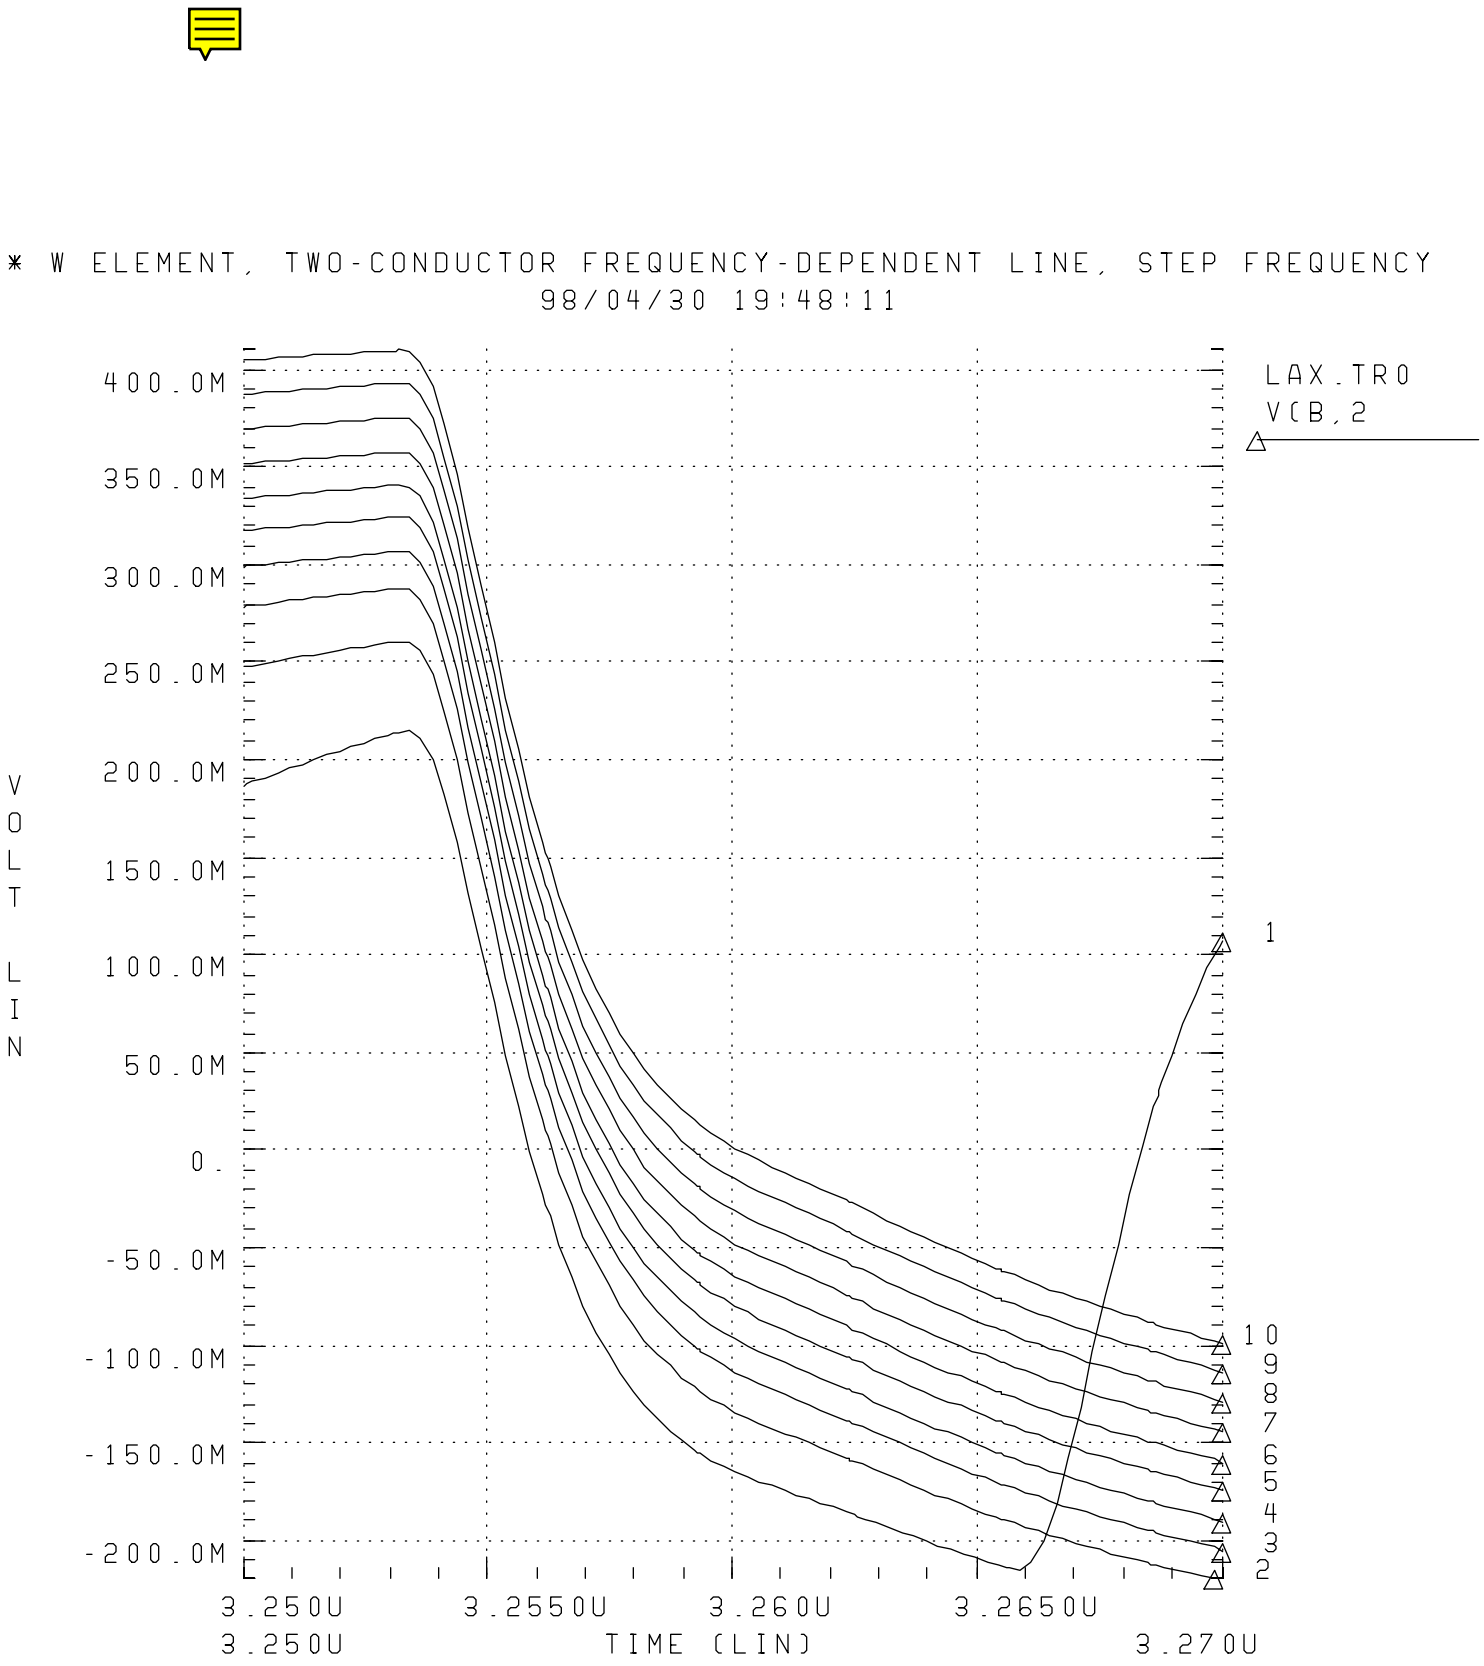
* W ELEMENT, TWO-CONDUCTOR FREQUENCY-DEPENDENT LINE, STEP FREQUENCY
98/04/30 19:48:11
LAX.TRO
V(B,2
400.0M
350.0M
300.0M
250.0M
200.0M
150.0M
100.0M
50.0M
0.
-50.0M
-100.0M
-150.0M
-200.0M
VOLT LIN
1
10
9
8
7
6
5
4
3
2
3.250U
3.2550U
3.260U
3.2650U
3.250U
TIME (LIN)
3.270U

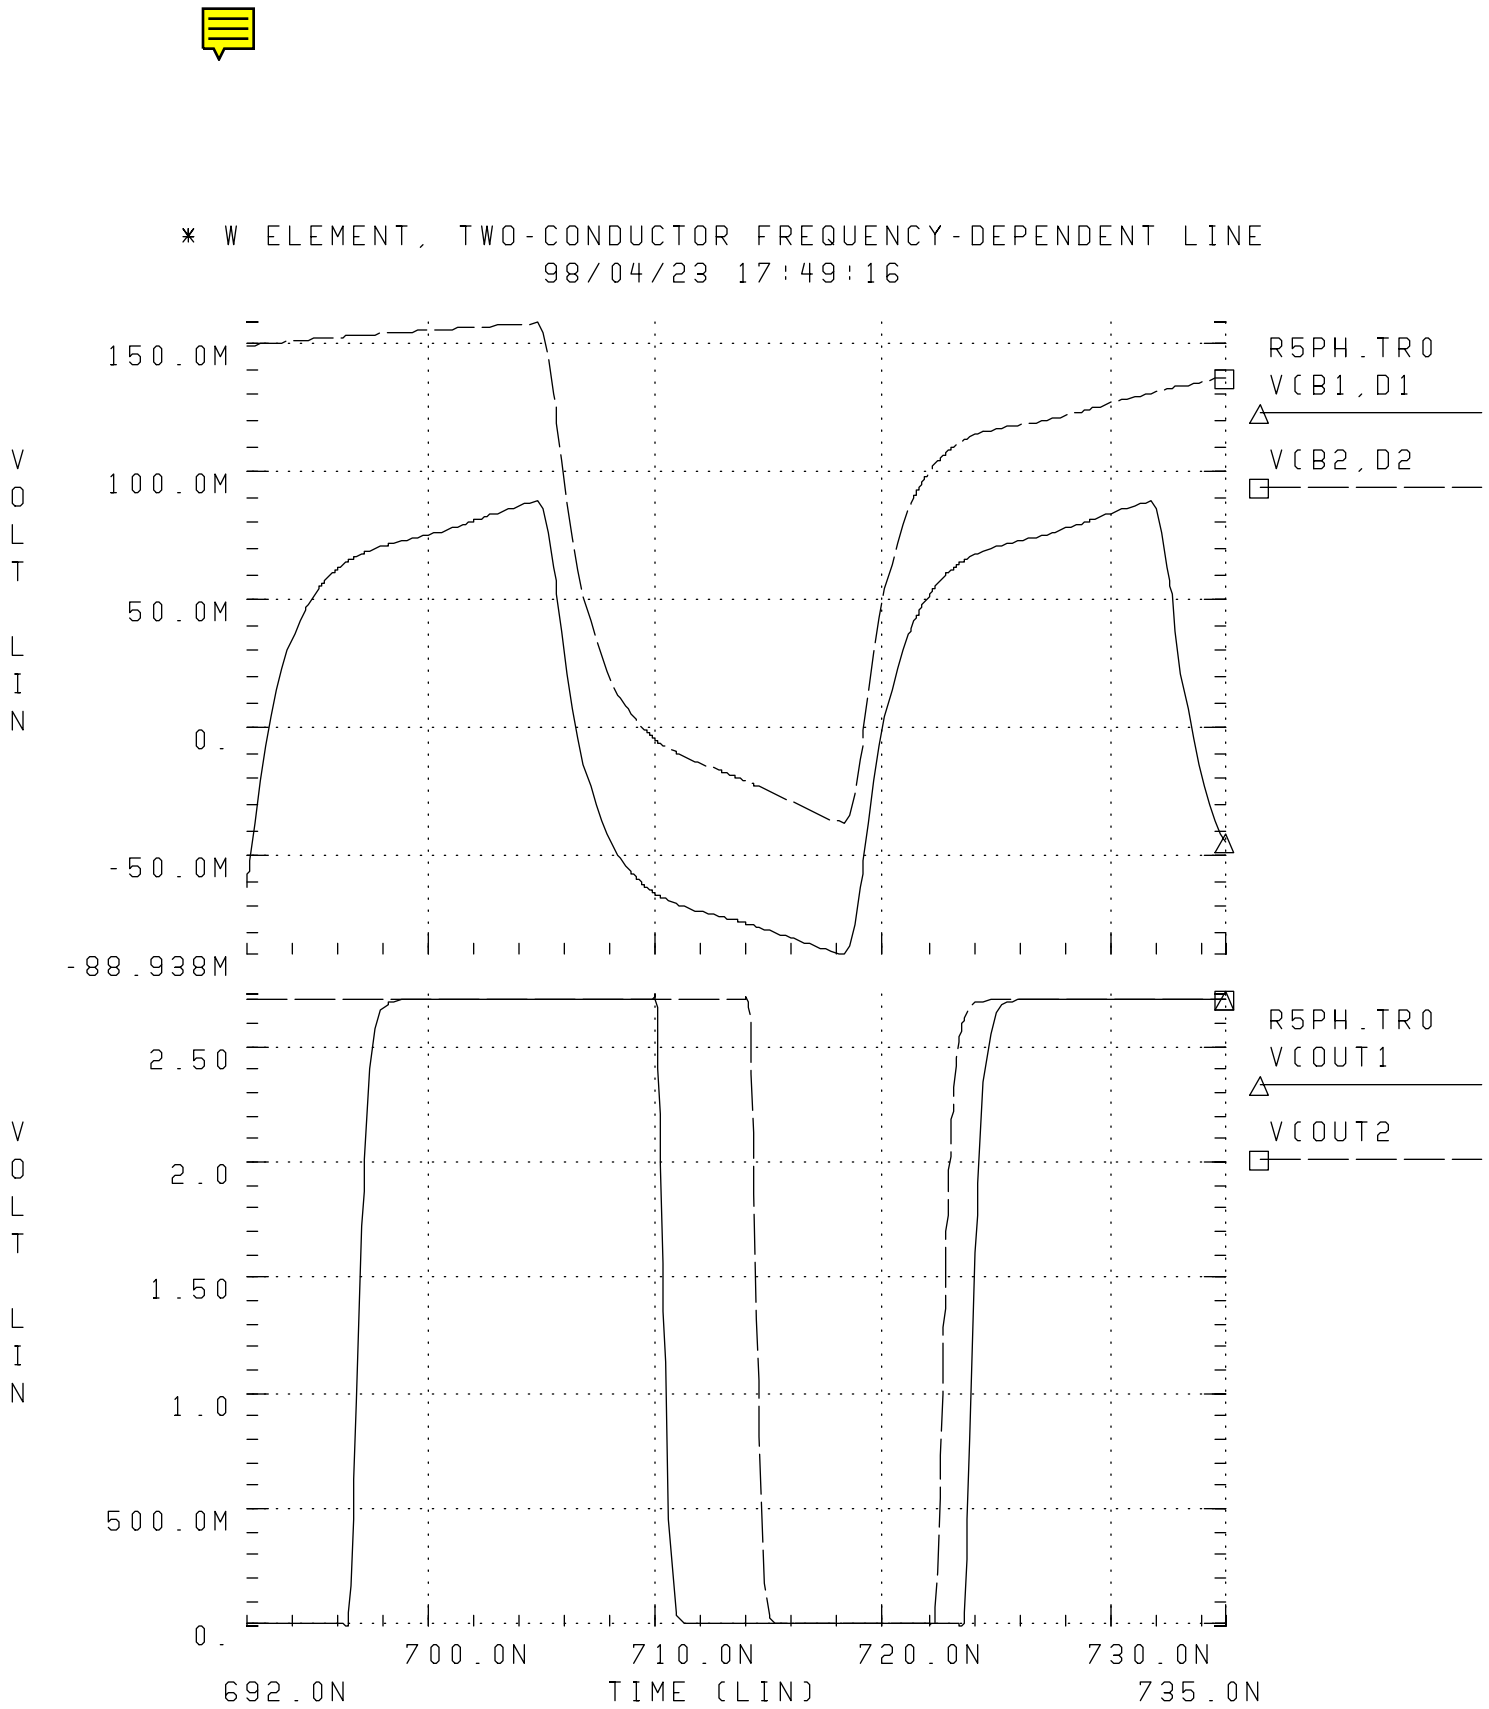
* W ELEMENT, TWO-CONDUCTOR FREQUENCY-DEPENDENT LINE
98/04/23 17:49:16
150.0M
100.0M
50.0M
0.
-50.0M
-88.938M
VOLT LIN
R5PH.TR0
V(B1,D1
V(B2,D2
2.50
2.0
1.50
1.0
500.0M
0.
VOLT LIN
R5PH.TR0
V(OUT1
V(OUT2
692.0N
700.0N
710.0N
720.0N
730.0N
735.0N
TIME (LIN)

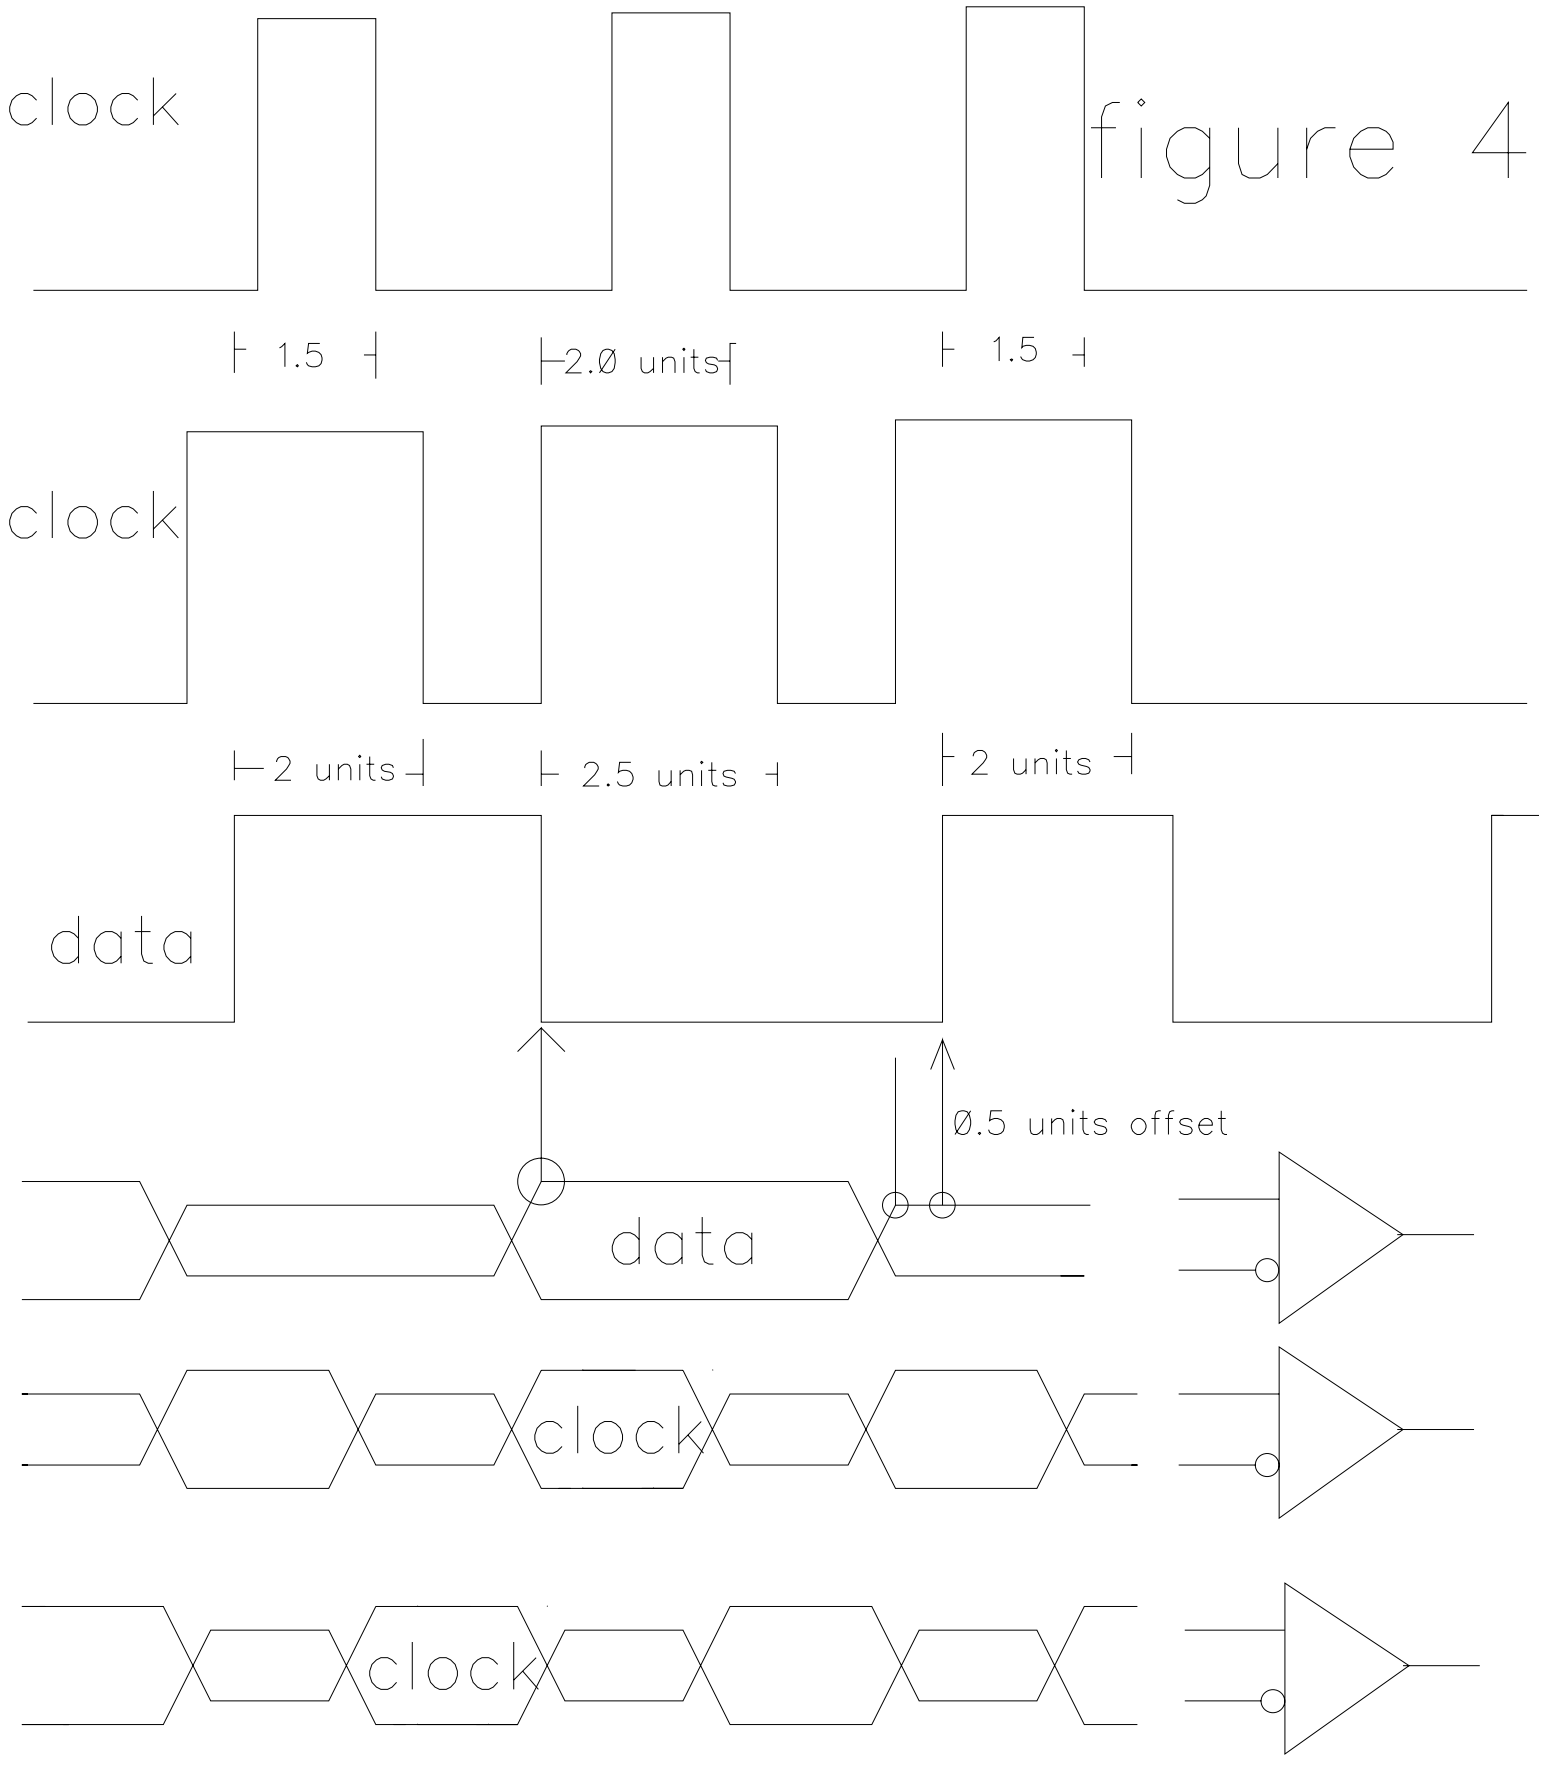
clock
1.5
2.0 units
1.5
clock
2 units
2.5 units
2 units
data
Ø.5 units offset
data
clock
clock

figure 4

figure 5

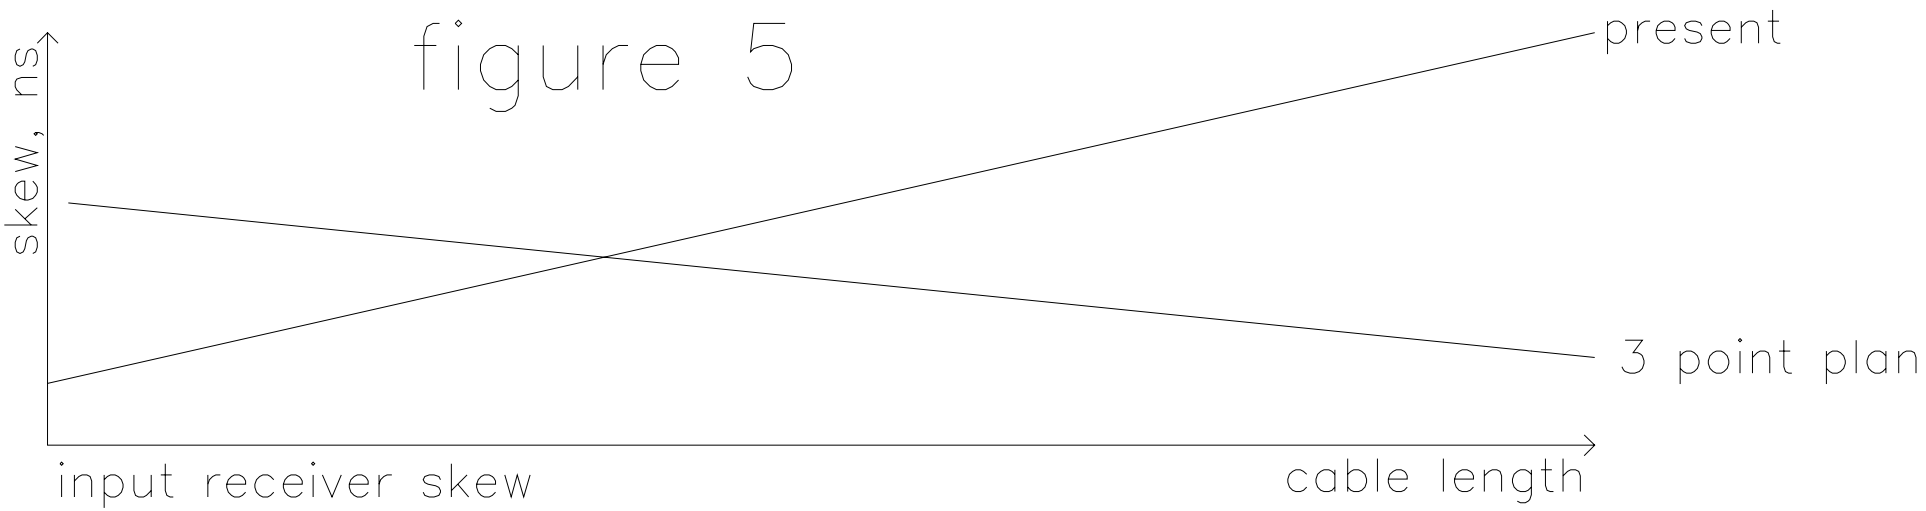
skew, ns
present
3 point plan
input receiver skew
cable length

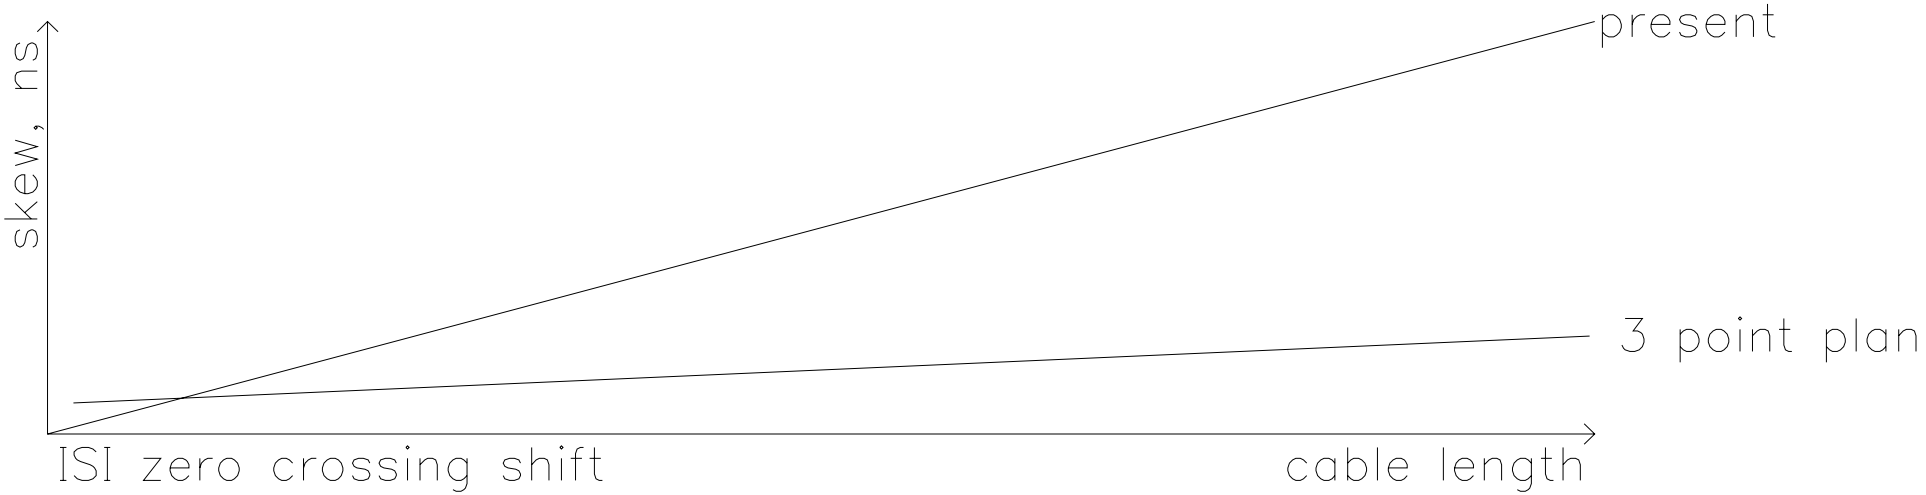
skew, ns
present
3 point plan
ISI zero crossing shift
cable length

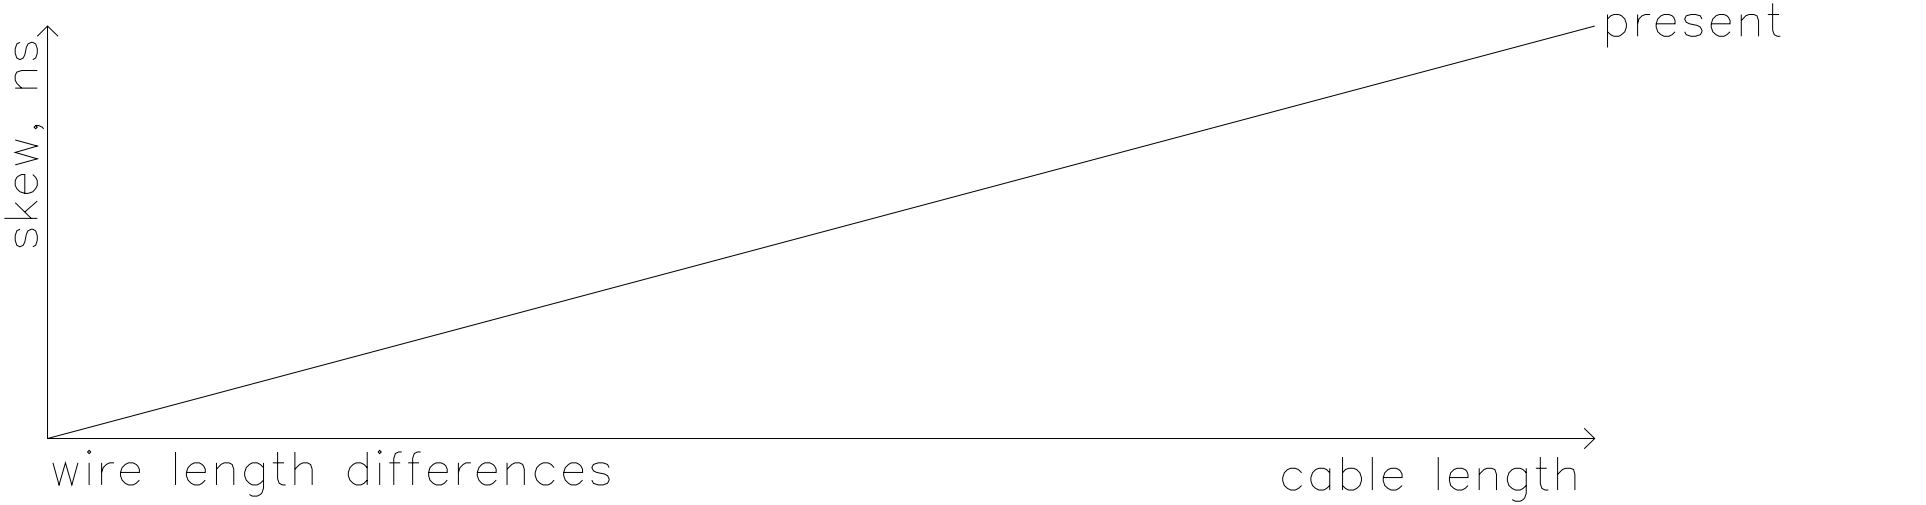
skew, ns
present
wire length differences
cable length

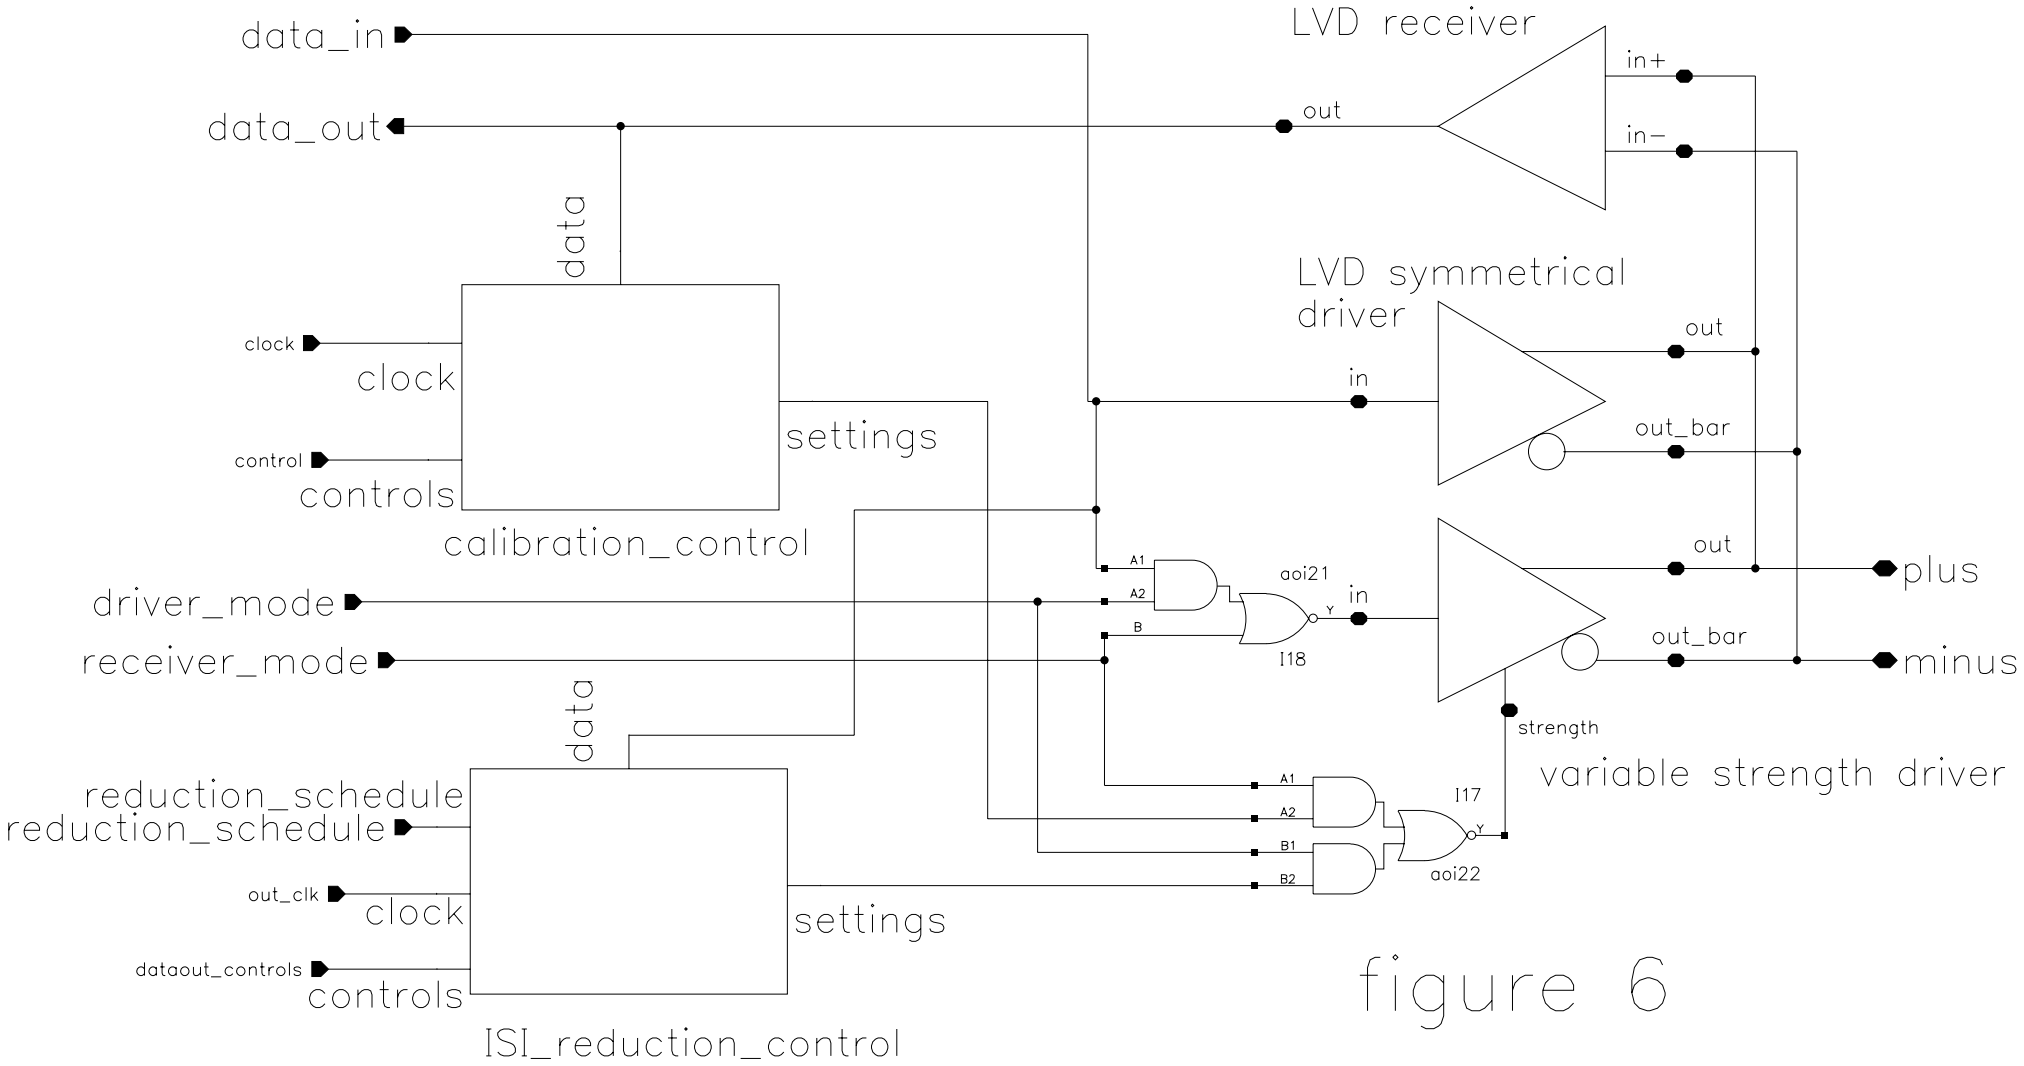
data_in
data_out
data
clock
clock
control
controls
settings
calibration_control
LVD receiver
out
in+
in−
LVD symmetrical driver
in
out
out_bar
driver_mode
receiver_mode
A1
A2
B
aoi21
I18
Y
in
out
out_bar
plus
minus
strength
variable strength driver
data
reduction_schedule
reduction_schedule
out_clk
clock
dataout_controls
controls
settings
ISI_reduction_control
A1
A2
B1
B2
I17
aoi22
Y

figure 6