

Article for Printed Circuit Design

By Lee W. Ritchey, 3Com Corporation

Differential Signaling Doesn't Require Differential Impedance Or, How to Design a Differential Signaling Circuit

That title may seem like a complete contradiction to the "wisdom" written in many design documents describing how to route differential pair signals. It is ---- and that is just what its intent is. Now that your attention has been caught, let's look at differential signaling.

Virtually all design rules in common use require designers to provide a specified differential impedance between pairs of signals involved in differential signaling circuits such as differential ECL or LVDS (Low Voltage Differential Signaling). As a result, much effort is consumed trying to decide on PCB geometries that will provide the desired impedance and an equal amount of effort trying to lay out the traces and build the PCB. On top of this, much time is spent trying to accurately measure the differential impedance.

As it turns out, differential impedance doesn't play a role in this form of signaling and is not necessary. How it was decided to impose this condition on this type of signaling is a mystery. It probably happened because the word differential is in the title. In any case, as will be seen in the following explanation of how this signaling works, other design considerations are more important.

First, it is useful to understand the problem that differential signaling was designed to solve. **Differential signaling is designed to transmit logic signals between two boxes or units that have logic grounds offset from each other by an amount too large for single ended logic signals to function correctly.** The receiver circuit is built with the ability to "float" with the driver circuit. The receive circuit has very high gain, so it can successfully operate on severely attenuated signals. Hence, we can compensate for large amounts of signal attenuation over long paths as well.

There are a couple of other benefits that we get with differential signaling in certain circumstances. They both occur when we have differential signals (read equal and opposite signals) travelling on two wires closely spaced to each other. The first is reduced EMI. This stems from the fact that the magnetic fields produced by the equal and opposite currents cancel each other. This is of very real benefit in the link between laptop computers and their displays. The second is immunity from common mode noise coupled into the pair by outside noise sources. This occurs because any magnetic noise field couples equally into both wires resulting in a common mode signal. Neither of these is of any real benefit when the wires are routed over the planes of a PCB.

How Differential Signaling Works

Figure 1 is a schematic of an ECL differential signaling circuit. A single ended logic signal originating in Box A must be successfully sent to Box B. (Box A and Box B could be two high pin count BGAs, QFPs or PGAs with severe ground offsets or ground and Vcc bounce.) The

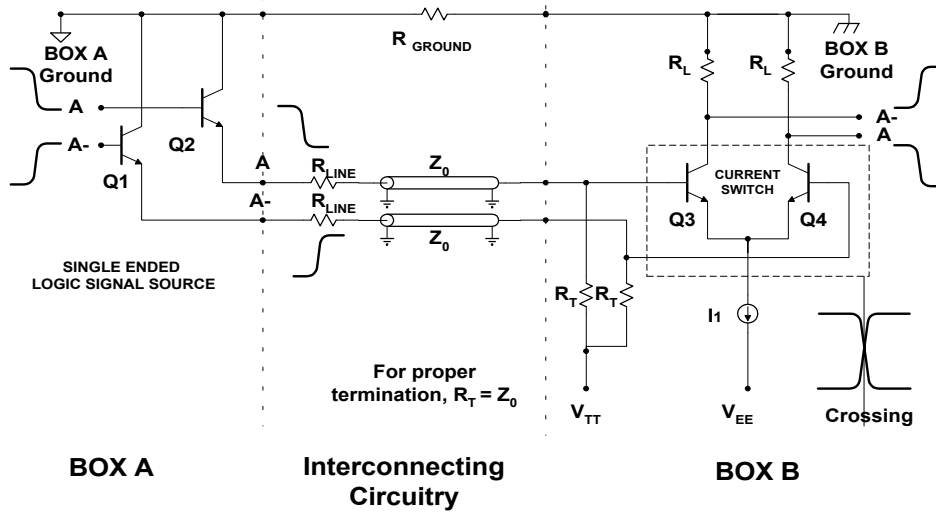


FIGURE 1, TYPICAL ECL DIFFERENTIAL SIGNALING CIRCUIT

path connecting the two logic grounds together has significant resistance impedance or in it, such that it is not possible to hold the two grounds within the DC limits necessary for proper operation of the ECL logic family. In Box A a phase splitter creates two copies of the logic signal, A and A-, each exactly the same size, but 180 degrees out of phase with each other. These two copies are tightly timed to each other. A and A- are sent out of the drive chip on independent emitter followers, Q1 and Q2. These emitter followers are identical to those used to send single ended ECL signals down transmission lines.

The objective is to deliver each of these two signals to the bases of the two transistors, Q3 and Q4 making up the current switch in Box B, with sufficient signal quality for the current switch to accurately measure the “difference voltage” between the two ends of the transmission lines. This is accomplished by sending each signal down its own properly terminated transmission into Box B. The secret to the successful operation lies in the current switch in Box B, not in how the two lines are routed between the two boxes. Again, the objective is to deliver two clean signals to the bases of the current switch. This can be accomplished by routing the signals side by side, on different layers or on the same layer spread apart, so long as the two signals arrive in good condition at about the same time, meaning that each is a well designed transmission line with a proper termination. There is no need for differential impedance to accomplish this.

How the Current Switch Works

The current switch made up of Q3 and Q4 looks to some like a differential amplifier. If the circuit were being driven by an analog signal that was to be amplified, it could be considered as such. However, that is not its function. The current switch is used to switch the current I_1 up through one or the other of the two collector resistors. Depending on which transistor is conducting, the current switch sends a logic one or logic zero into Box B. The path the current takes depends on which transistor base is the more positive of the two. The transistor whose base has the higher voltage of the two will conduct the current while the other transistor will be off. In an ECL current switch, a difference in voltage of only 15-20 millivolts will do this job. The difference in voltages at the emitter followers Q1 and Q2 is nominally 1000 millivolts. From this it can be seen that the signals can suffer severe attenuation and the current switch will still make a successful logic decision. Alternatively, if attenuation is small, the current switch can tolerate very large amounts of noise on the two signals and still make a correct logic decision.

The Current Switch as a Crossing Detector

One way to look at the current switch is as a crossing detector. In fact, that is just what its job is. Deciding whether a logic one or logic zero is being sent turns on which transistor base of the

Q3-Q4 pair is more positive. This pair determines that the logic state has changed from one level to the other by detecting when the voltages on the two bases cross. Indeed, it is preserving this crossing that is key to this form of signaling working successfully. So, what we really have here is a crossing detector tasked with sensing when the voltage levels at the load ends of the two transmission lines cross. This crossing is preserved by making sure that the time of arrival of the two switching levels is within the logic family's timing tolerances and that the noise present on the switching edges is within limits. Note that neither of these requirements involves differential impedance.

Dealing with Ground Offsets

The main function of differential signaling is to compensate for ground offsets between the two ends of the signal path. This is done in the receive end of the signal path. The emitters of the current switch, Q3 and Q4, are fed a current I_1 by a current source. The property of a current source is that the current remains constant for a wider range of terminal voltages. The collectors of transistors are also current sources that will maintain the same current for a wide range of collector to emitter voltages. The magnitude of the current I_1 is chosen such that it produces a voltage drop across either collector resistor, R_L , that represents a logic zero. The difference between this voltage value and the maximum voltage across the current source necessary for it to operate correctly is a "voltage window" that the Q3-Q4 transistor pair is free to move in. Therefore, the Q3-Q4 pair can move up and down as the circuits in Box A shift with ground offsets between Box A and Box B. In ECL circuits, this window can be as large as 2500 millivolts, meaning that the signal path will work properly with ground offsets as large as 2500 millivolts.

Timing Tolerances

The Q3-Q4 transistor pair is a crossing detector. It will accurately determine the time when a logic signal changes from one logic state to another as long as the slopes of the two changing signals are crossing. As can be seen from the diagram, the time interval where this is true is the length of two rise or fall times. With standard 10K ECL this turns out to be approximately 1500 picoseconds. Therefore, if the two signals arrive in time within this time limit, the circuit will accurately determine logic ones and zeros. If this time is turned into distance using the propagation delay of a signal in an FR-4 PCB, the result is about nine inches. This means that the two sides of the differential pair could differ in length by up to nine inches and the circuit would still work correctly.

The significance of the length discussion above is in the affect it has on length matching rules when routing the two halves of a differential pair. Often, routing rules (length matching rules) are imposed on differential pairs that are so tight that it is impossible to route them using layer changing vias or, in most cases to use auto routing. These constraints make routing unnecessarily difficult. Length matching need only be good enough to preserve the timing. For all but the very fastest differential signaling protocols, length matching need not be more precise than 500 mils. This will allow auto routing and layer changing vias on virtually all differential signals.


LVDS Circuits

LVDS, Low Voltage Differential Signaling, circuits are rapidly becoming the circuits of choice for high data rate signaling in a noisy environment. As mentioned earlier, this protocol has been

invaluable in the design of laptop computers and other products that must transfer large amounts of fast data between parts of a system where there is poor grounding. Figure 2 is a typical LVDS circuit showing voltage levels as well as the two transmission lines, each terminated in their characteristic impedance for a total of double Z_o . One might question why the double Z_o resistor isn't called a differential termination. To see why, imagine that each resistor was terminated back to the reference voltage half way between the two logic levels, which is where they should be connected. When this is done, it is easy to determine that the current flowing into that terminal through one resistor is equal and opposite to that flowing out of that terminal through the other resistor, so the net current is zero. Since the net current is zero, it is not necessary to make the connection at all. We can say that each terminating resistor has been tied to a "virtual" ground.

The function of the LVDS circuit is exactly the same as the ECL circuit described above. Therefore, we can use the same layout rules for LVDS as were used for ECL. We will need to determine the system timing tolerance in order to calculate the length matching accuracy limits. In the case of the LVDS family supplied by National Semiconductor, this values is 250 picoseconds, so the length matching would need to be done with an accuracy of 1.5 inches or less.

Logic 1 level is +1.475 Volts.
Logic 0 Level is +0.925 Volts. Z_o



Lay out each line as a stand alone transmission line.
 Match lengths within limits set by LVDS standard.
 Center of terminating resistor goes to a virtual ground.
 Therefore, each line is terminated in Z_o .

Figure 2. Typical LVDS Bus Layout with a Single Driver.

Handling the Transition from Twisted Pairs to PCB Traces

Much confusion centers around how to handle the transition from twisted pair transmission lines to transmission lines on a PCB. Figure 3 shows several types of transmission lines. In all four cases the electromagnetic field travels between the two conductors that make up the transmission line. All four types are characterized by specifying the impedance measured between the two conductors.




Figure 3. Four Typical Transmission Lines

In the case of the **twin lead transmission line**, the electromagnetic field travels between the two wires and the impedance is measured one with respect to the other. This is often referred to as differential impedance, but in truth, it's just the impedance. Whether the wires are twisted or not, this is still a twin lead transmission line. A termination for this type of transmission line is placed across the ends of the conductor pair. The termination could be a transformer as is the case in TV and Ethernet circuits.

In the case of the **coaxial transmission line**, the electromagnetic field travels between the center conductor and the outer shield. The impedance is measured between the center conductor and the outer shield. A termination is placed between the center conductor and the outer shield at the end away from the driver.

In the other two examples, **micro strip line and stripline**, the electromagnetic field travels between a trace and one or more power planes. The impedance is measured between the trace and the power plane.

Where confusion enters the picture is when a twin lead transmission line meets the pair PCB transmission lines. What the designer must do is adjust the impedance of the pair of PCB transmission lines so that added together they are the same impedance as the twin lead. All that is needed is to make sure that each PCB transmission line is half of the impedance of the twin lead to accomplish this. Then, at the ends of each PCB transmission line, terminate them to the power plane with their individual characteristic impedances. No need to get involved with differential impedance.

Conclusions

Differential signaling is a very robust way to transmit logic signals between boxes or sources and loads in an environment where there is substantial noise on the power subsystem. There is a cost penalty associated with this form of signaling. The cost is twice as many wires per data line and special drivers and receivers. However, this cost may be far less than would be incurred by trying to build a power and ground structure that has enough copper in it to maintain the voltage drops in power and ground within limits imposed by single ended logic signaling.

Successful differential signaling does not require differential impedance. What it does require is two signal lines, each properly terminated in its own characteristic impedance at the receiver end. These two lines must be equal in length to within the timing tolerances of the logic family

being used. In virtually all cases, the timing accuracy will allow length differences up to 500 mils or 12.5 millimeters. On PCBs there is no appreciable electrical benefit to routing the two signals side by side. It is convenient to do this, but no harm comes of separating the two wires by any distance needed to move around obstacles such as connector or BGA pinouts. Similarly, no harm comes of allowing the use of vias to change layers, so long as the impedance of the layers being used is the same.