A Practical Guide to Lossy Differential Lines

Wolfgang Maichen and Bo Krsnik
Teradyne, Inc.

The authors can be reached at
Teradyne, Inc.
30701 Agoura Road
Agoura Hills, CA 91301, USA
wolfgang.maichen@teradyne.com, krsnik1@yahoo.com

1 Introduction

In recent years *differential signaling*[1] [1], [2], [1] has made large inroads in high-speed transmission schemes. Estimates are that in a few years almost 100% of all PCB's will have at least some differential signal paths on them. There are several reasons for this: On one hand having a symmetrical pair of lines carrying opposite signals close to each other greatly reduces electromagnetic emissions because the electromagnetic far-fields of the two lines largely cancel. Second, since in such a transmission scheme the total signal current over the two lines of a differential pair is constant (at least as long as no differential skew is present), the current spikes drawn by the drivers are reduced by an order of magnitude, reducing power supply noise and ground bounce (see the section below about power decoupling). Third, differential transmission is much less sensitive to residual ground bounce or external influences than single ended signaling because the influences on the two lines of a pair are of similar size (as long as the two lines are close together) and so largely cancel out since the receiver is only sensitive to their difference.

In contrast to wide-spread belief, there is nothing really special required *per se* for two lines to be "differential" - the only distinctive feature is that the signals these two lines carry are not independent, but are always complementary. Things like "coupling" and "differential impedance" are the result of specific design techniques associated with differential signaling rather than prerequisites for it. This means that differential lines share a lot of issues – for example losses - with single ended lines. This paper intends to show how differential lines can be better understood and modeled with a minimum of effort and with easily available tools. The reader will gain insight into the risks associated with differential signals like delay mismatches (skew), crosstalk, mode conversion, and signal losses. With a good understanding of those effects differential lines lose a large part of their mystery and the risks of using them can be well controlled.

2 Differential signaling

2.1 Differential Line Fundamentals

When two transmission lines are very close to each other, their electric and magnetic fields start to overlap and induce voltages and currents into each other, interfering with the original signals on the lines. In "normal" (single ended) signaling this is observed as capacitive and inductive crosstalk, and it is an unwanted feature. Another way to look at it is that the two lines have some mutual capacitance as well as mutual inductance which, if there is a transition on the other line, adds to or subtracts from the self inductance and the capacitance against the ground plane [1][2][3].

However, if those two lines form a differential pair, then the transitions on them are no longer independent, and the crosstalk has the same effect on both of them for each transition. In this case we talk of "coupling", but it really is just the same physical phenomenon. In which way the signals are influenced depends on the relative polarity of the transitions: If they are of opposite polarity ("odd mode"), the effective capacitance is increased, the effective inductance reduced, and thus the effective "odd mode impedance" Z_{odd} is reduced as well. For same-polarity transitions ("even mode"), the case is exactly the other way around, causing the "even mode impedance" Z_{even} to be higher than the impedance of an isolated line.

[1] When one line is high, the other one is low, and when one line transitions, the other line transitions in the opposite direction.

$$Z_{odd} = \sqrt{\frac{L_u - L_{m,u}}{C_u + C_{m,u}}}, \quad Z_{even} = \sqrt{\frac{L_u + L_{m,u}}{C_u - C_{m,u}}}, \tag{1}$$

where L_u and C_u are the inductance and capacitance of each line of the pair per unit length, and $L_{m,u}$ and $C_{m,u}$ are the mutual inductance and capacitance between them.[2] As long as the electric field is completely contained in a homogeneous dielectric, the propagation speeds v_p (and hence the propagation delays T_{pd}) of both modes are identical (and the same as for a single ended line):

$$v_p = \frac{c_{light}}{\sqrt{e_r}}, \quad T_{pd} = \frac{length \times \sqrt{e_r}}{c_{light}} = length \times \sqrt{C_u \times L_u} = \sqrt{C \times L} \tag{2}$$

Here e_r is the dielectric constant, and C and L are the total line capacitance and inductance, respectively. This is the case e.g. for coaxial cables as well as for striplines (a trace "sandwiched" between two ground planes) in a PCB, but not for microstrip lines (a trace on the surface of the PCB, i.e. with the dielectric and a single ground plane below, but air above).[3] For simplicity we will restrict our analysis in the following sections to the homogeneous case. In this case signal propagation along the differential pair is fully characterized by the two impedances Z_{odd} and Z_{even} plus a single propagation time, while a single ended transmission line would only need a single impedance Z_0:

$$Z_0 = \sqrt{\frac{L}{C}}. \tag{3}$$

These two impedances Z_{even} and Z_{odd} are related to differential impedance Z_{diff} (the total impedance seen by the differential signal, for which the two lines are effectively in series) and common impedance Z_{common} (where the two lines act in parallel) by simple formulas:

$$Z_{diff} = 2 \times Z_{odd}, \; Z_{common} = \frac{Z_{even}}{2}, \; Z_0 = \sqrt{Z_{even} \times Z_{odd}}, \; Z_{even} \le Z_0 \le Z_{odd} \tag{4}$$

The equality between Z_{even}, Z_{odd} and Z_0 happens when the lines are completely uncoupled[4]. The beauty of this concept is that the transmission equations for differential and common signals stay the same as for single ended signals, as long as one replaces the impedance with the proper value for each case!

In real-world applications of differential signaling close line spacing (resulting in considerable coupling) is used for several reasons: First, it reduces emitted radiation—very important for a device in order to be compliant to EMI rules and regulations. Second, it makes differential lines less susceptible to external fields because those influences will cause the same disturbance in both lines of the pair, so they cancel out in the differential receiver. Third, routing both lines close together automatically assures that their propagation times will be well matched, so the differential signal fidelity at the receiver is preserved. Fourth, since crosstalk (coupling) between the two lines is not a concern, routing the lines close together conserves board space, allowing either for smaller boards or less layers, thus decreasing board cost and size. In summary, we see that close coupling is a side effect of other considerations in connection with differential signaling, but it is not an inherent necessity.

So what are the real additional PCB routing requirements, compared to single ended signals [1][2][6][7][8]: A common misconception is that, unlike single ended signals, differential lines don't really need a continuous ground return path because the current flowing into one line already returns through the other line of the pair, so there is no current through the ground plane. This would be true if the signals were really perfectly differential, and - even more important - the traces were so close to each other that the coupling between them would exceed by far the coupling to the ground plane (in a different view, this corresponds to the demand that the return currents below the

---

[2] Differential line theory prefers to think in terms of even and odd mode (instead of differential and common signals), because this results in a more symmetric picture. The connection between even (e) and odd (o) mode signal on one side, and differential (d) and common (c) signal on the other side is straightforward (a and b are the single ended signal levels on the true and the complement line of the differential pair, respectively): $o = (t-c)/2$, $e = (t+c)/2$, while $d = (t-c) = 2o$, $c = (t+c)/2 = e$. Inversely, $a = e + o$, $b = e - o$. For a symmetric differential line (the vast majority of practical cases), those two "modes" (or, alternatively, differential and common signal) have the special property of traveling along the lines without conversion into the respective other mode (while e.g. a single ended signal on such a pair would get converted because of crosstalk to the victim line).

[3] This is also the reason why microstrips exhibit far-end crosstalk, but striplines don't.

[4] Note that Z_0 here is *not* the impedance of an isolated line, rather it is the impedance when one of the two lines of the pair is driven by a signal and the other one kept silent. Like Z_{odd} its value decreases with increasing coupling between the lines, but much less so than Z_{odd}.

traces completely overlap and thus cancel, leaving current only in the traces but not in the ground plane(s)). But in reality, due to rise/fall time mismatches, skew caused by path length differences or driver mismatches, etc. there is always some amount of common mode signal present that has to return to its source though the ground plane. And even more important, the trace-to-trace coupling in real designs is always much weaker than the coupling to a massive ground plane [2]. For edge-coupled striplines with the closest achievable spacing, where the trace distance of two 50 O lines is equal to the trace width, return current overlap (and cancellation) in the ground plane is only about 10% [2]. So there is no difference to single ended signals here - ground planes are indispensable. (One notable exception are *twisted-pair lines*[5] [1][2].)

Still, for differential signal integrity purposes it is good practice to keep the two lines of a pair similar in performance (losses). Routing them similarly shaped and on the same PCB layer(s) (to avoid variations caused by dielectric material tolerances) helps achieve this goal, but does not put very stringent requirements on the matching tolerances.

2.2 Skew Within a Differential Pair – Mode Conversion

A well-designed differential channel consists of two well-matched lines. Any mismatch between those two lines of the pair will result in performance degradation. Common causes for mismatches are skew, i.e. a difference in the propagation delays of the two lines, and differences in parasitics (which constitute low-pass filters with delays on the order of their time constant, and distort the signal in proportion to their size), or in short, any asymmetry between the two lines.

Such asymmetries have two important effects: They lead to so-called mode conversion, and they reduce the differential path bandwidth. Mode conversion, despite its impressive sounding name, is actually very easy to understand. Figure 1 gives an example where two initially perfectly aligned signal edges of opposite polarity (i.e. a pure differential signal without any common component) travel down a differential path that has a slight delay mismatch between the lines. At the end we see that the sum of the two skewed signals now contains a common signal spike (which will lead to excessive EMI radiation). The same process will also transform common signals from the input (or from external interferences) into differential signals at the output, thus reducing the common noise immunity (common mode rejection ratio) of the signal path.

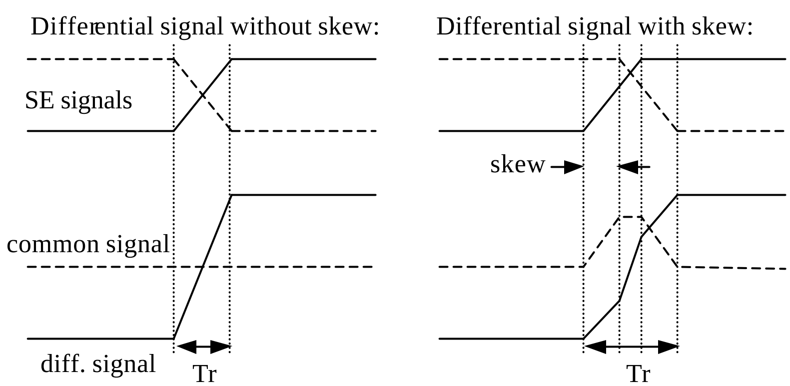

Figure 1: Skew within a differential pair creates common mode spikes and increases the differential signal's rise time (i.e. reduces the effective differential bandwidth).

In addition, skew between the two signals of a differential pair causes a rise time increase of the resulting differential signal (see also Figure 1). The interesting conclusion from that is that for differential signals, intra-pair skew is a source of bandwidth limitation compared to single ended signals, even when there are no "real" losses (see

[5] In twisted pair lines, where two wires are snugly wound around each other, the spacing between these two signal wires is usually small compared to the spacing to the shield, and thus the two return currents in the shield nearly cancel – in other words, the shield is hardly carrying any current, and virtually all the current is flowing in one signal line and returning through the other. In this case one can even remove the shield without affecting the line impedance much (but immunity against external fields will be somewhat reduced).

later) present. To get an approximation for the effective limitation one can use the well-known relationship between rise time Tr and -3dB bandwidth BW, BW ≈ 0.33/Tr, using the total skew as a rough measure for the rise time. Note that this rise time / bandwidth does not behave exactly like the bandwidth limitation caused by e.g. a parasitic in the path. While the latter adds as RMS to the incoming rise time, the former adds up linearly, and in addition the rise time in the formula is actually the 10/90 rise time. As an example, a skew of 100 ps would give an estimate of 3.3 GHz of bandwidth, while measurement of such a structure (two cables of unequal length) with a Vector Network Analyzer gives a -3dB bandwidth of 2.34 GHz. Still, the order of magnitude is correctly predicted.

In order to avoid excessive radiation and to preserve signal integrity, the requirement is that the total differential skew does not exceed a small fraction of the signal rise time. This puts high demands on the routing of the PCB: For example, if the signal has 100 ps rise time (this corresponds to a propagation of about 2 cm using a low-loss dielectric), the matching should be better than 10 ps or about 2 mm. While this may sound easily feasible, keep in mind that the exact delays of discontinuities like vias or bends tend to be very difficult to calculate and have considerable manufacturing variations.

2.3 Crosstalk between Adjacent Lines

One of the reasons stated very often for why differential lines are useful is crosstalk reduction. While it is true that going differential reduces crosstalk from *external* sources (i.e. sources whose distance is large compared to the spacing between the two lines of the pair), this is not necessarily true for crosstalk between traces on a PCB, e.g. in a wide bus. The crosstalk from a single ended stripline trace to another[6] is given by the following approximation formula (using expressions for C_u, L_u, $C_{m,u}$, and $L_{m,u}$ from [9]):

$$NEXT = \frac{1}{4} \times \left(\frac{C_{m,u}}{C_u} + \frac{L_{m,u}}{L_u} \right) \approx \frac{1}{4p} \times \frac{H \times (H+W)}{(S+W)^2} \times 100\,\% \,. \tag{5}$$

H is the spacing between signal trace and ground plane, W is the width of the trace, and S is the edge-to edge spacing between the two traces. To use this formula to calculate the crosstalk between two differential pairs we simply have to calculate the crosstalk of each line of the aggressor pair into each line of the victim pair and then add those four partial results (taking into account that some of the signals have inverse polarity).

Figure 2 plots the calculated crosstalk for different edge-to-edge spacings for three different cases: First, two closely coupled differential lines where the spacing within the pair is equal to the line width (which is usually the minimum achievable value). Second, two single ended lines with a spacing equal to the differential pair-to-pair spacing of the first case. And third, two single ended lines that use the same routing area as the differential case.[7] One can argue that it makes most sense to compare cases (1) and (3) since between them everything stays the same except for differential vs. single ended signaling, while case (2) would take much less routing space than the differential case, which in turn – not all that surprisingly – will lead to higher crosstalk.

The surprising result from Figure 2 is that for close spacing (D smaller than 2.5 line widths) the differential bus actually behaves *worse* than the widely spaced single ended bus taking up the same routing area (width). The reason is that according to above formula the crosstalk falls off approximately with the square of the distance, so the crosstalk between the two inner traces (one aggressor and one victim) completely dominates the result. The conclusion is that differential signaling does not significantly relieve us from following roughly the same spacing requirement as for single ended lines: E.g. to keep the crosstalk below 1%[8] the edge-to-edge distance between differential pairs should be at least twice the line width (compared to 3 times for single ended lines). With the additional routing space required for the second line of the differential pair this actually results in a *larger* required routing space than for a single ended line.

[6] Striplines exhibit only near-end crosstalk (NEXT), but no far-end crosstalk (FEXT) [2].

[7] We should note that in order to make Z_0 of the single ended lines equal to Z_{odd} of the differential pair the line width of the single ended lines would have to be slightly larger. However, since the formula is mainly dependent on the sum of spacing and width this would only mean a very minor correction to the overall result.

[8] This is the crosstalk from one line to the next. In a bus the total crosstalk on the victim can be more than twice that high since it has *two* closest neighbors and additional ones farther away (the two closest neighbors contribute roughly 95% of the total crosstalk for line spacing equal to line width [2], as is the case in our example).

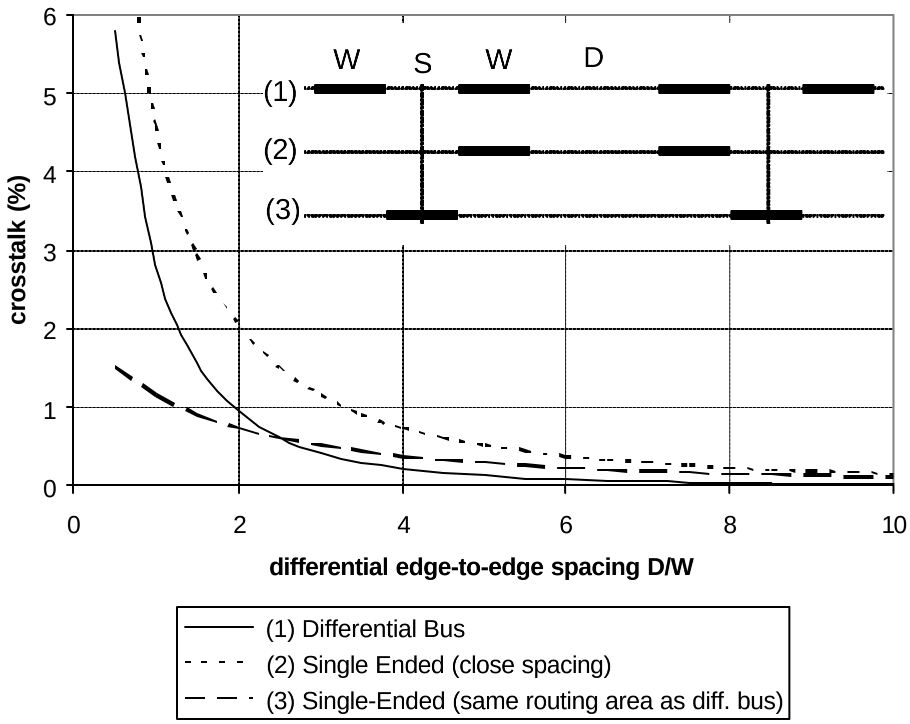

Figure 2: Crosstalk comparison for differential lines vs. single ended lines

3 Lossy Differential Lines

3.1 General Considerations

As mentioned in the previous section, as long as we only need to consider differential signals alone, i.e. neglect common signal components, there isn't much difference between a single ended line (characterized by Z_0 and T_{pd}, or alternatively by L and C) and a differential pair (characterized by Z_{diff} and T_{pd}, or alternatively by $L_{diff}=L-L_m$ and $C_{diff}=C+C_m$). To apply the theory below to single ended lines instead, we would only need to substitute $Z_{diff} \to Z_0$, $L_{diff} \to L$, and $C_{diff} \to C$.

The general model of a lossy differential transmission line, carrying only differential signal, is shown in Figure 3: Its components are the series inductance L_{diff}, shunt capacitance C_{diff}, series resistance R, and shunt conductance G.[9] For an ideal, loss-less transmission line, R and G are zero. Ohmic resistance and skin effect loss (which is nothing else than ohmic resistance aggravated by inhomogeneous current distribution due to skin effect) are modeled by a frequency dependent $R(f)$, while dielectric losses are represented by $G(f)$, again frequency dependent.[10] This frequency dependency is what causes all the trouble in modeling the signal propagation. The general expression for the characteristic impedance Z_{diff} of this line is [1][4]:

$$Z_{diff} = \sqrt{\frac{R + jwL_{diff}}{G + jwC_{diff}}}, \quad w = 2pf \cdot \tag{6}$$

For sufficiently high frequencies and/or relatively small losses ($R << wL_{diff}$ and $G << wC_{diff}$) this can be approximated as

[9] For a homogeneous transmission line all those parameters are distributed evenly along the length of the line.
[10] Note that R is usually given as a *resistance*, while G is seen as an *admittance* (i.e. inverse of a resistance).

$$Z_{diff} = \sqrt{\frac{L_{diff}}{C_{diff}}}, \tag{7}$$

which is the same as for a loss-less transmission line. Practically usable digital interfaces all fall into this range. This is an important piece of information because it tells us that – apart from some signal attenuation – the signal propagation is the same as for loss-less lines, e.g. Z_{diff} is largely independent of frequency, and the following expressions hold [4]:

$$\begin{aligned} L_{diff} &= T_{pd} \times Z_{diff} \\ C_{diff} &= T_{pd} / Z_{diff} \end{aligned} \tag{8}$$

This is the first set of equations that will allow us to determine model parameters (L and C in this case) based on measurements. The general propagation constant of a lossy line is [4]:

$$g = a + jb = \sqrt{(R + jwL_{diff}) \times (G + jwC_{diff})}, \tag{9}$$

where a (the real-valued part) describes the signal attenuation, and jb (the imaginary part) describes the wave propagation along the line. Under the same low-loss assumptions as before ($R << wL_{diff}$ and $G << wC_{diff}$) this approximates to [4]:

$$\begin{aligned} a &\approx \frac{1}{2} \times \left(\frac{R}{Z_{diff}} + GZ_{diff} \right) \\ b &\approx w \times \sqrt{L_{diff} C_{diff}} = w \times T_{pd} \end{aligned} \tag{10}$$

For a loss-less line a is of course zero. An important conclusion from a in formula (5) is that the signal gain[11] through a lossy line is

$$gain = \frac{V_{out}}{V_{in}} = e^{-\frac{R}{2Z_{diff}}} \times e^{-\frac{GZ_{diff}}{2}}. \tag{11}$$

V_{in} and V_{out} are the signal amplitudes going into and coming out of the line, respectively. This is the third formula we need to build our model, since it relates the measured loss through a cable (or trace) to our loss parameters R and G. Since gain (or loss) can span a wide range of magnitudes, it is usually stated in the logarithmic dB scale, thus

$$\begin{aligned} gain_{dB} &= 20 \times \log(gain) = -20 \times \log(e) \times \left(\frac{R}{2Z_{diff}} + \frac{GZ_{diff}}{2} \right). \\ loss_{dB} &= -gain_{dB} \end{aligned} \tag{12}$$

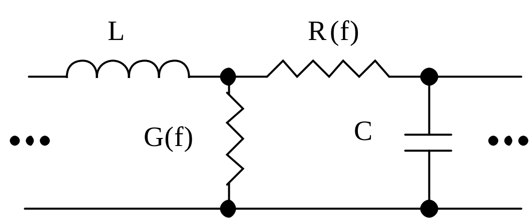

Figure 3: General model element for a lossy transmission line

[11] Although transmission theory usually talks in terms of gain, in our case *gain < 1*, i.e. it really is loss.

3.2 Loss Models

The only piece of theory still missing is a set of models for the behavior of the different loss contributors versus frequency. The easiest is ohmic (DC) loss, since it does not depend on frequency. It forms one part of the series resistance R, and we denote it with R_{DC}.

The second part of the series resistance is caused by skin effect. For a perfect coaxial cable the skin resistance is proportional to the square root of the frequency and one can even derive an analytical expression for the proportionality factor [1][4]. For arbitrarily shaped transmission lines (e.g. striplines in a PCB) where the field distribution is more complicated this is no longer possible, but the square-root-behavior remains at least approximately valid. Thus we can model the series resistance due to skin effect as

$$R_{skin} = k_{skin} \times \sqrt{f} \tag{13}$$

For our purposes the factor k_{skin} will simply be a fit parameter, to be adjusted to match the measured behavior. The skin resistance adds up with the DC resistance to the total resistance[12] to give

$$R = R_{DC} + R_{skin}. \tag{14}$$

Finally dielectric losses – represented by the shunt conductance G – increase approximately linearly with frequency[13], i.e.

$$G = 2\boldsymbol{p} \times f \times C_{diff} \times \tan\boldsymbol{d} = k_{diel} \times f = k_{diel} \times \left(\sqrt{f}\right)^2. \tag{15}$$

We'll se in a second what the strange square-root formulation is good for. Just as for the skin effect, since we usually don't know the effective loss tangent, for us the factor k_{diel} is a fit parameter to be adjusted to best fit the measurement.

We can now insert those loss models into the general formula (7) for the total line losses:

$$gain_{dB}(f) = -20 \times \log(e) \times \left(\frac{R_{DC}}{2Z_{diff}} + \frac{k_{skin}}{2Z_{diff}} \times \sqrt{f} + \frac{Z_{diff}\,k_{diel}}{2} \times \left(\sqrt{f}\right)^2 \right). \tag{16}$$

If we look closely we see that – if we plot the measured gain in dB versus the square root of the frequency (instead of the frequency itself) this reduces to a simple parabola:

$$\begin{aligned} & gain_{dB}(x) = a + bx + dx^2 \\ & x = \sqrt{f} \\ & a = -20 \times \log(e) \times \frac{1}{2Z_{diff}} \times R_{DC} \\ & b = -20 \times \log(e) \times \frac{1}{2Z_{diff}} \times k_{skin} \\ & d = -20 \times \log(e) \times \frac{Z_{diff}}{2} \times k_{diel} \end{aligned} \tag{17}$$

In other words, once we have measured the loss (in dB) over a range of frequencies, the actual data fitting becomes almost trivial. One possibility is to use Excel for this task, as almost every engineer has it readily available, and it can perform polynomial fits. But any other fitting software will do the task as well. From the fit parameters a, b, and d it is then easy to calculate the loss parameters R_{DC}, k_{skin}, and k_{diel}.

---

12 This is only a rather crude approximation, but it greatly simplifies fitting the loss behavior to the measured data. Reference [2] mentions that a better fit to the actual resistance trend can be achieved by adding R_{DC} and R_{skin} as RMS (root-mean-square), i.e. R=sqrt(R_{DC}^2+ R_{skin}^2)).

13 This is really just an approximation because for any real-world dielectric material the loss tangent tan δ varies somewhat with frequency, although the variation is small compared to the frequency variation.

3.3 Measurement and Model-Building

For the practical measurement, a Vector Network Analyzer (VNA) is the tool of choice.[14] The reasons are two-fold: First, as we have seen losses have an easy-to-describe behavior in the frequency domain, in contrast to rather difficult-to-interpret time domain (e.g. step response) behavior. Second, they have unparalleled amplitude (and therefore loss) measurement resolution, so one can cover the whole range from very small to very large losses.[15] All this does not mean the resulting Spice model is restricted to frequency domain simulations (AC sweeps), rather it will do just as well in a time domain simulation (transient response). If we have a differential line, we will need a four-port VNA that can perform differential through-measurements.

We need to set the frequency sweep to logarithmic and measure the absolute value[16] of the differential transmission coefficient (S_{dd21}) – plotted on a dB scale this is then identical to our $gain_{dB}$ from before. This data is transferred to a PC equipped with Excel (or some other plot program), the curve is re-plotted vs. the square root of the frequency and fitted with formula (12) to yield R_{DC}, k_{skin}, and k_{diel}.[17]

As the next step we change the setup of the VNA to display the *group delay* of S_{dd21}. The group delay is defined as the derivative of the phase j over the frequency f, i.e.

$$T_g = \frac{dj}{df} \,. \tag{18}$$

For dispersion-less paths the group delay has a very simple meaning – in those cases it is identical to the propagation delay T_{pd} of the path. Fortunately this is normally a good approximation for our situation.[18] If the VNA does not have the option to display the group delay directly, we can instead have it display the phase and then calculate it from the slope of the phase curve.[19] Usually the group delay display only yields useful readings in the high-frequency region (above 100 MHz) because for lower frequencies the phase change is very small (since the wavelength is much longer than the propagation time) and the result thus heavily impacted by measurement noise and other inaccuracies. Since the differential path impedance Z_{diff} is usually known (in digital application it is almost without exception 100 Ω), from the measured propagation delay we can calculate the total differential line capacitance C_{diff} and the total line inductance L_{diff} using formula (3). Otherwise we can use differential time domain reflectometry (TDR) to measure Z_{diff}.

A few more measurement hints may be useful. First, we should always start out with the longest line available to minimize the impact of fringe effects – in other words, path loss must dominate effects from impedance mismatches at the connection points on both ends (e.g. SMA connectors, probes, etc.). Otherwise characterization and mathematical de-embedding of those connectors would be necessary, which goes far beyond the scope of this article (this is one area where professional modeling tools come into play). Second, we want to use the widest frequency range available – as low and as high as the VNA can go – since the wider the range, the more reliable the curve fit will be (it is always more accurate and reliable to interpolate than to extrapolate[20]). Ideally we should go to at least twice the highest frequency present in our signal (i.e. the signal bandwidth). For a digital signal, the approximate bandwidth (*BW*) is *not* given by the data rate or the clock frequency, but by the rise time $T_{R,10/90}$ of the signal:

---

[14] A scalar network analyzer will do just fine for the loss measurement, but we will need to determine the propagation delay T_{pd} with some alternative method (e.g. time domain reflectometry (TDR)). If we don't have any network analyzer available, then one can measure the attenuation using a sine-wave source (an RF generator) and an oscilloscope, though it is tedious to acquire a sufficient number of points, and small losses will be difficult to measure accurately.

[15] A VNA can easily reach a signal-to-noise ratio of 100 dB (a factor of 10^5), while an oscilloscope will be hard-pressed to reach even 60 dB (a factor of 10^3).

[16] Note that S_{dd21} is a complex value (it has both magnitude and phase).

[17] The choice of frequency sweep (linear or logarithmic) has a slight influence on the fit result: A linear sweep would over-emphasize the high frequency range, while a logarithmic sweep distributes the measured data points evenly over the whole range.

[18] The condition for constant group delay is that the dielectric constant e_r does not vary with frequency. This can only be exactly true for loss-less media; whenever there are losses, the dielectric constant changes over frequency in proportion to those losses.

[19] One potential trap here is that the VNA calculates the group delay numerically based on the phase difference between adjacent data points. If the frequency spacing between is too large this can give erroneous results because the phase wraps around every 360 degrees. To avoid this, we need to assure that the VNA sweep uses enough (i.e. sufficiently closely spaced) points. A simple test is to double the number of points and verify that the displayed group delay curve does not change.

[20] Although one of the benefits of the presented fitting method is that it allows extrapolating losses beyond the measured range if necessary with good confidence.

$$BW \approx \frac{0.33}{T_{R,10/90}} \tag{19}$$

As far as simulation is concerned, a popular tool is Spice. Among the many Spice variants available we chose Orcad's PSpice because it offers a lossy transmission line model as a built-in library component (TLOSSY), and a free demo version (restricted in the number of components in a design, but otherwise virtually fully functional) is freely available for download [5]. The first issue one encounters is that PSpice has only models for *either* single ended lossy lines *or* coupled (differential) lines, but no coupled lossy lines. But since we only want to model the differential signal propagation (odd mode), we can simply use the single ended lossy line and make its impedance equal to Z_{diff}. The second issue is that PSpice does not offer the frequency as a parameter for the lossy line, but provides the Laplace parameter $s = 2pjf$ instead. This is quickly overcome by using the substitution

$$f = abs\left(\frac{s}{2p}\right). \tag{20}$$

In addition to the parameters C, L, R, and G, the TLOSSY model in PSpice offers a fifth parameter called LEN. Since our parameters are referred to the total line, we can simply set LEN to 1.[21] Finally, if we want to use formulas for parameters in PSpice, they need to be enclosed in curly braces (see the example below).

As a final consideration, what can we do if we do not have access to the far end of the line, e.g. because it ends in a socket or in a needle probe head? This will preclude transmission measurements. The solution is to measure the reflected signal (S_{dd11}) instead of the transmitted (S_{dd21}), since this requires only a single connection to one end of the line. For this to work the other (far) end must remain open (unterminated). The signal then transverses the line, gets fully reflected at the other end, and returns to the source, so it effectively traverses the line *twice*. Data collection and fitting is done as before, the only difference is that we must divide all the measurements (T_{pd} and dB-loss) by two to represent a single traversal.[22]

3.4 An Example

In the following we will show the theory applied to a practical example. The object to be modeled was a pair of coaxial cables with an SMA connector on each end.[23] The intended signal bandwidth was around 2 GHz. The VNA could cover the frequency range between 300 kHz and 8.5 GHz. Care must be taken with calibration since even slight errors affect the accuracy of the measured parameters, especially in the low-frequency region where the losses are small.

Figure 4 shows the measured loss curve as well as the model simulation results. The match is excellent over the whole range, especially considering that the model consists of just a single circuit element in PSpice. The slight "wiggle" at high frequencies is due to reflections at the connector discontinuities, and is obviously not included in the model. The DC resistance in this particular case was negligible and was omitted (by telling Excel to set the y-intercept to zero). The measured and fitted parameters are:

$$\begin{aligned}
&Z_{diff} = 100\ \Omega \text{ (known, and verified with TDR)}\\
&T_{pd} = 2.80 \text{ ns (from group delay, and also from TDR)}\\
&a \equiv 0 \text{ (DC losses negligible)}\\
&b = 2.826 \cdot 10^{-5} \text{ (from curve fit)}\\
&d = 8.148 \cdot 10^{-11} \text{ (from curve fit)}
\end{aligned} \tag{21}$$

From this we can calculate the model parameters for the lossy cable (using formula (3), (12), and (15)):

[21] After the line is modeled, this parameter is a nice way to scale the model up or down for a different line length, e.g. if we have measured a long line to minimize connector effects, but want to built a model for a shorter section of the same type of line.
[22] We neglect here the effect of the parasitic fringe capacitance that can be different when the line is open compared to when the line is terminated at the far end.
[23] The quality of the connection was verified using time domain reflectometry, but it is also visible in the fact that there is virtually no "wiggle" (caused by reflections between the connectors) in the measured S_{dd21} curve at high frequencies.

$$
\begin{aligned}
&LEN = 1 \\
&C = 28.0\,\text{pF} \\
&L = 280\,\text{nH} \\
&R = \left\{6.508 \cdot 10^{-4} \times \sqrt{abs(s/2p)}\right\} \\
&L = \left\{7.504 \cdot 10^{-13} \times abs(s/2p)\right\}
\end{aligned} \tag{22}
$$

To match the VNA measurements, the lossy line in the Spice model must be terminated with a matched 100 Ω differential termination, and the AC sine wave source has to have 100 Ω differential impedance as well. Even though the example plot is a frequency domain sweep, the PSpice model could just as well be employed for time-domain (transient) simulations, without any changes to the model itself.

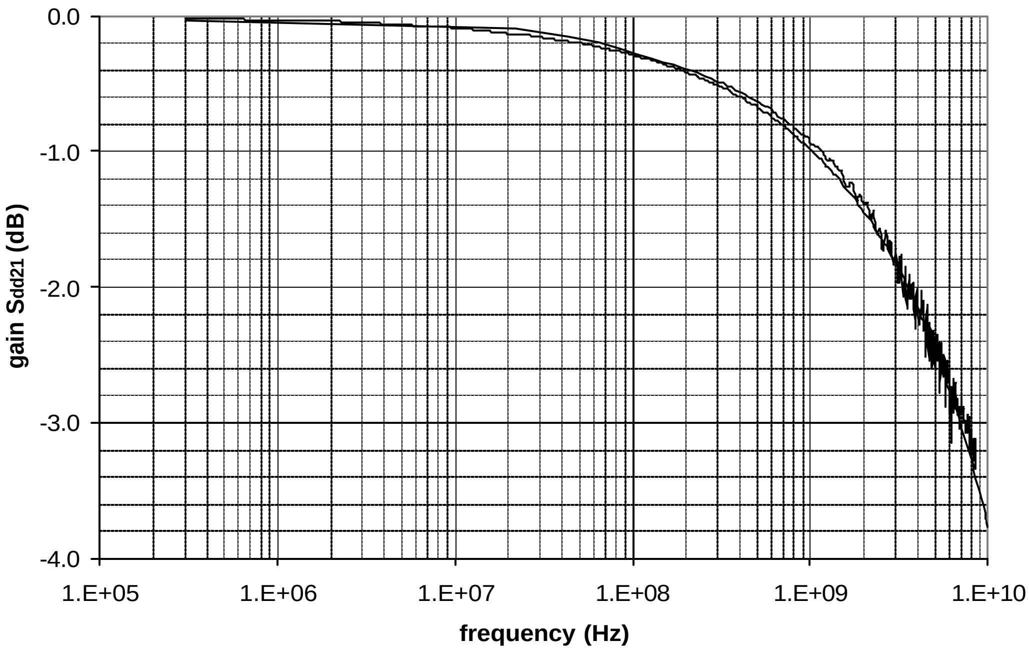

Figure 4: Measured and simulated loss curves.

4 Bandwidth and Data Dependent Jitter (ISI)

The final question that remains is – why is maximizing the path bandwidth so important? It is clear that when the bandwidth is *much* too low for our particular application, the signal rise time would become longer than a bit interval, the signal would not cross the threshold within the bit interval, and therefore the receiver would not switch within the allotted time frame. This limit is certainly a concern in actual applications where we try to squeeze out the last bit of performance (speed) from a given backplane or cable, but less so in test where we can afford to spend more time and effort to make the system perform well – after all, in semiconductor test we want to measure the behavior of the device under test, not of our connection path.

But even below the absolute rise time limit bandwidth limitation have an influence on our test results. We can sum up this influence with the following statement: *Whenever there are bandwidth limitations in the path, the signal will exhibit pattern dependent timing errors (also called pattern dependent jitter, inter-symbol interference, pulse pulling, or pulse width effects).* Since in reality every signal path has *some* bandwidth limitation (even though it may be high), the question is less *if* we will get pattern dependent jitter, but rather *how much*. Figure 5 shows the basic mechanism that creates this jitter: If the signal has been high for a long time, it has settled out to the "true" high level, and a transition to low crosses the threshold at a certain time. But if the signal has been low and just briefly was driven high (in the case displayed in Figure 5 for a single bit interval) before going low again, the signal hasn't yet fully settled to the high level and thus goes into the low transition with a "head start". As a result it will reach the threshold level earlier. Put together, the exact timing of the threshold crossing becomes dependent on the "data history" of the path, i.e. preceding data bits affect the timing of the current one.

For effects that result in a monotonously increasing transmission loss with frequency, the two most extreme cases are precisely what Figure 5 indicates – a minimum-width pulse vs. a transition after staying a long time at the same level. This is the case for skin effect and dielectric losses, rise time degradation through parasitic elements, rise time limitations of the signal source itself, and for all practical purposes also for intra-pair skew, so Figure 5 is the perfect test case for all those effects. Reflections (caused by parasitics or by impedance mismatches) don't have such a simple behavior and thus need more elaborate analysis to find the maximum possible timing error.

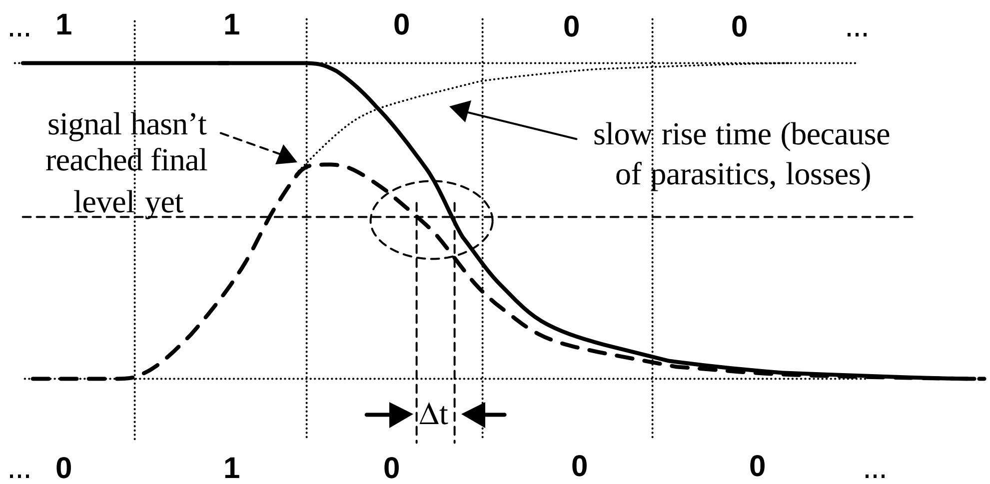

Figure 5: Rise time and settling limitations (or, equivalently, bandwidth limitations) always lead to pattern dependent timing errors Dt (jitter, ISI): The timing of the same transition (1 to 0 in this example) becomes dependent on the "data history" (the preceding data bits) before this transition (here, …1111 vs. …0001).

5 Conclusion

We have given an introduction into the issues of interest for practical differential signaling. Emphasis was on the close connection between single ended and differential lines – while they differ in some respect, they have much more in common. We have shown how crosstalk and losses affect differential lines just as single ended ones. In addition, the effect of intra-pair skew turns out to be another bandwidth-limiting factor.

6 References

[1] Howard Johnson, Martin Graham, "High-Speed Signal Propagation", Prentice Hall, 2003.

[2] Eric Bogatin, "Signal Integrity - Simplified", Prentice Hall, 2004.

[3] Stephen H. Hall, Garrett W. Hall, James A. McCall, "High-Speed Digital System Design", John Wiley & Sons, Inc., 2000.

[4] David M. Pozar, "Microwave Engineering" (Second Edition), John Wiley & Sons, Inc., 1998.

[5] Orcad PSpice Version 10.0 (Free Demo Version), download from www.orcad.com.

[6] Mark I. Montrose, "Printed Circuit Board Design Techniques for EMC Compliance", Wiley-IEEE Press, 2000.

[7] Douglas Brooks, "Differential Trace Design Rules: Truth vs. Fiction", UltraCAD Design, Inc., download from www.ultracad.com.

[8] Lee W. Ritchey, "Differential Signaling Doesn't Require Differential Impedance", Printed Circuit Design Magazine, March 1999.

[9] Charles S. Walker, "Capacitance, Inductance and Crosstalk Analysis", Artech House, 1990