



# Driving Copper Cables with HOTLink™

## Overview

The HOTLink™ family of data communications products are designed to support communication over both optical and copper cables. Each media type has specific cost, bandwidth, emissions, and distance criteria. This application note covers the methodology and evaluation of various forms of attachment to copper media. It is expected to be used in conjunction with a companion application note titled, "HOTLink Design Considerations."

## Primary Topics

The primary topics covered in this application note are:

- Transmission lines
- Copper cable types
- Direct coupling
- Capacitive coupling
- Transformer coupling
- Quantitative Interface Comparison

## Introduction

The electromagnetic spectrum covers all wavelengths from near zero through infinity. This includes all radio, microwave, light, x-ray, and cosmic-ray wavelengths. Table 1 lists the classifications of those frequencies and wavelengths usually associated with data communications over copper media. Communication links based on HOTLink products utilize frequencies in the HF, VHF, and UHF bands.

Copper cables (or circuit board traces) are used to move electromagnetic energy from one place to another. With slow signal-switching speeds (and short interconnect distances), a signal placed on one end of a cable eventually shows up at the other end. Systems of this type are seen and used in homes and offices every time a light switch is opened or closed. Here the primary concern is delivering energy to a load.

In high-speed communications systems, many other concerns exist. Not only must energy be delivered to the communications link receiver, but the signal delivered must arrive with minimal distortion. Delivery of electromagnetic energy with minimal (or controlled) distortion requires the proper use of transmission lines.

## Transmission Lines

In the most general sense, a transmission line is any closed system for directing electromagnetic energy. (While antennas may also direct electromagnetic energy, they are not part of a closed system and are thus not considered transmission lines.) Any transmission line meets the following three criteria:

- Has a system of material boundaries
- Has a start and end point
- Capable of directing electromagnetic energy

Electromagnetic energy moves along a transmission line as an electromagnetic wave, composed of electric and magnetic fields. These waves and fields travel (or propagate) down a transmission line at a finite rate, determined primarily by the dielectric in the transmission line.

Transmission lines generally fall into two different types, based on the orientation of the electromagnetic fields as they propagate down the transmission line. All dual-conductor transmission lines (coaxial, twisted-pair, twin-axial, microstrip, stripline, etc.) propagate their electromagnetic energy with both the electric and the magnetic fields oriented perpendicular to the direction of propagation. This is known as Transverse Electric Magnetic (TEM) mode. Figure 1 shows a graphic representation of these fields within a coaxial cable.

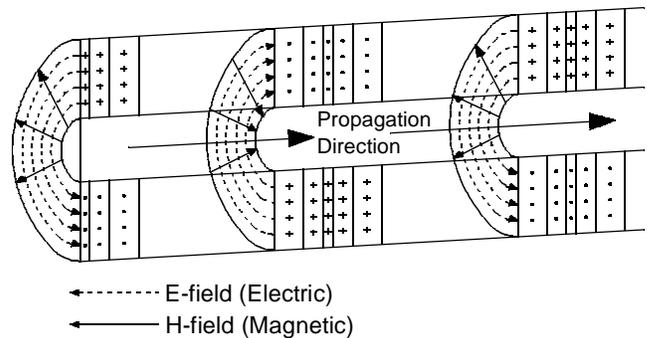


Figure 1. Electric and Magnetic Fields for TEM Mode in a Coaxial Transmission Line

Single-conductor transmission lines (also known as waveguides) propagate their energy in multiple modes known as either TE (Transverse Electric field) or TM (Transverse Magnetic field). In these modes, one or the other of the fields is oriented parallel to the direction of propagation.

Both TEM and TE/TM transmission lines have cutoff frequencies—points in the electromagnetic spectrum where the transmission modes change. For TEM transmission lines the cutoff frequency determines the upper frequency limit for TEM propagation. Signal components higher than the cutoff frequency must propagate in TE/TM modes.

For TE/TM (waveguide-type) transmission lines, the cutoff frequency determines the frequency below which energy cannot propagate. This cutoff frequency is determined by the physical dimensions of the waveguide, and is calculated using Equation 1.

$$f_{(c)} = \frac{300,000 \text{ km}}{2 \times \text{Wall\_Width}} \tag{Eq. 1}$$

Applying this equation to the data rates used with HOTLink shows that such a structure would be very impractical. It

**Table 1. Electromagnetic Band Classifications**

Band	Band Name	Frequency Range	Wavelength Range	Common Uses
ELF	Extremely Low Frequency	30 Hz–300 Hz	10 Mm–1 Mm	Commercial AC Power Distribution
VF	Voice Frequency	300 Hz–3 kHz	1 Mm–100 km	Analog Telecommunications
VLF	Very Low Frequency	3 kHz–30 kHz	100 km–10 km	Voice and Music Reproduction, Submarine Communications, Sonar
LF	Low Frequency	30 kHz–300 kHz	10 km–1 km	Commercial AM Radio, Shallow-to-Medium Depth Sounders
MF	Medium Frequency	300 kHz–3 MHz	1 km–100 m	Commercial SW Radio, Amateur Radio, Marine Radiotelephone
HF	High Frequency	3 MHz–30 MHz	100 m–10 m	Commercial SW Radio, Amateur Radio, Citizen Band Radio
VHF	Very High Frequency	30 MHz–300 MHz	10 m–1 m	VHF Television Broadcast (Channels 2–13), FM Radio, Amateur Radio, Cordless Telephones
UHF	Ultra High Frequency	300 MHz–3 GHz	1 m–10 cm	UHF Television (Channels 14–83), Microwave Ovens, Aeronautical Radionavigation
SHF	Super High Frequency	3 GHz–30 GHz	10 cm–1 cm	Microwave Communications, Marine Radar, Aircraft Tracking and Radar
EHF	Extremely High Frequency	30 GHz–300 GHz	1 cm–1 mm	Space Communications, Radio Astronomy

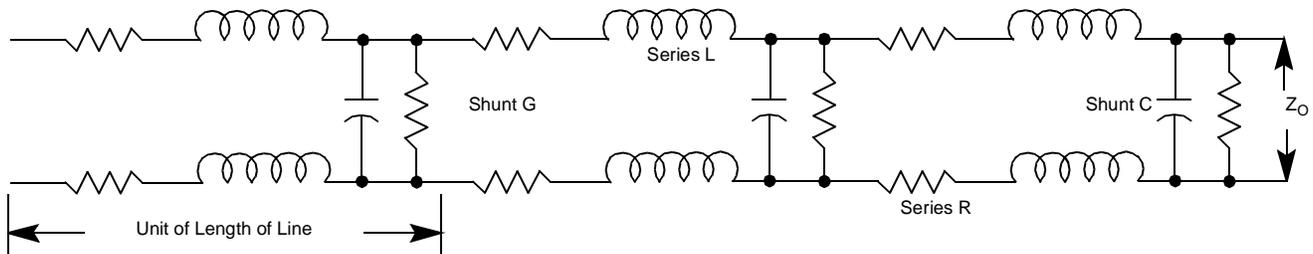
would require a cross-sectional width of nearly 4 meters to propagate the low-frequency signal components (40 MHz) of even the highest operating data-rate (400 Mbps) of HOTLink. Because of this restriction (and others) all remaining discussion will only deal with TEM-type transmission lines.

**TEM Transmission Line Characteristics**

The transmission line has numerous distributed parameters that determine its operation and characteristics. These distributed parameters include the series inductance (L) of the conductors in the transmission line, the shunt capacitance (C) between the conductors, the series resistance (R) of the conductors, and the shunt conductance (G) between the conduc-

tors. Because these properties remain constant per unit length of the transmission line, they are referred to as distributed properties. These parameters are functions of the diameter and spacing of the conductors and the dielectric constant of the spacer used between them. A schematic equivalent of these elements in a balanced (two-wire) transmission line is shown in *Figure 2*.

Transmission lines are usually characterized by two parameters: characteristic impedance ( $Z_0$ ) and velocity of propagation ( $V_p$ ). Proper determination of these values is imperative to allow the transmission line to be used correctly.



**Figure 2. Equivalent Circuit of a Transmission Line**

### Characteristic Impedance

The characteristic impedance identifies the impedance seen by a source when driving a transmission line terminated at the load-end in a pure-resistance equal to the characteristic impedance. While this appears to be a circular definition, it is valid. If the load end of the transmission line is terminated in an impedance other than the characteristic impedance of the line, the source end of the line will see an impedance different than either that of the load or the characteristic impedance of the line. Because this characteristic impedance is generally unaffected by frequency, a transmission line terminated in its characteristic impedance has the same load characteristic of a fixed resistor.

In most transmission lines the series-R and shunt-C values are usually very small and have minimal effect on the impedance of the line. This means that the characteristic impedance is determined almost entirely by the series-L and shunt-C shown in *Figure 2*. This relationship is shown in Equation 2.

$$Z_0 = \sqrt{\frac{L}{C}} \quad \text{Eq. 2}$$

### Velocity of Propagation

In space, an electromagnetic wave travels at nearly 300,000,000 meters per second (speed of light). Moving this same signal through a transmission line with a vacuum for the dielectric separator between the conductors allows the wave to propagate at or near this same rate.

Real transmission lines are seldom found with a vacuum dielectric. Instead, various non-conductive materials are used to maintain the spacing between the two conductors of the transmission line. These separators all have different dielectric constants, and all of them slow down the propagation of the signal. The rate the signal propagates, relative to the speed of light, is known as the Velocity of Propagation ( $V_P$ ) and is usually expressed as a percentage (sometimes expressed as a propagation delay in time per unit distance). This velocity difference may be calculated using Equation 3, where  $\epsilon_r$  is the relative dielectric constant of the transmission line.

$$V_P = \frac{1}{\sqrt{\epsilon_r}} \quad \text{Eq. 3}$$

For this calculation to work, the entire electromagnetic field must propagate in the dielectric. Many transmission lines are structured such that some of the field propagates in the dielectric, while other parts propagate in the surrounding air. For transmission lines of this type the equation must be modified to account for the mixed dielectrics.

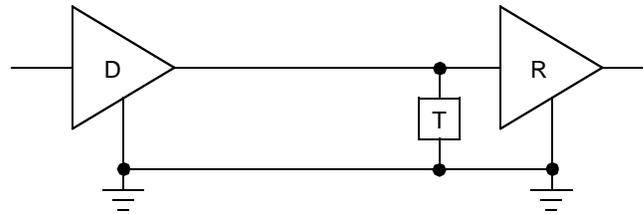
### TEM Transmission Lines

TEM transmission lines may be grouped in any number of different ways: by length, by construction, by dielectric, by usage, etc. For operation with HOTLink they are generally split into two categories: unbalanced (single-ended) and balanced (differential) transmission lines.

#### Unbalanced Transmission Line

*Figure 3* shows a driver/receiver combination used in an unbalanced transmission line. In this configuration, a single

driver sources and sinks current into the transmission line with the return path provided by a common ground.



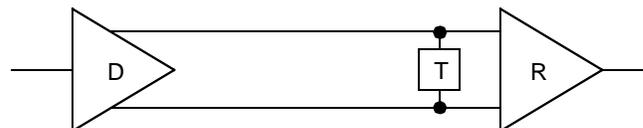
**Figure 3. Unbalanced Transmission Line**

In this configuration, other communications paths can share the common ground. This allows for fewer wires in a cable, and fewer contacts in a connector. The main problems suffered by this type of transmission line are susceptibility to external noise, crosstalk, ground potential differences, and limited noise margin.

In an unbalanced transmission line, the electromagnetic field necessary for signal propagation exists between the driven line and the ground path. The receiver operates by comparing the amplitude of the received signal relative to ground or some other reference.

#### Balanced Transmission Line

*Figure 4* shows a driver and receiver configured for use with a balanced transmission line. In this configuration, two drivers source and sink complementary signals into the two wires of the transmission line. These signals need to be matched in amplitude, and must be 180° out of phase with each other for the transmission line to work properly.



**Figure 4. Balanced Transmission Line**

In this configuration, a common ground is not always necessary. Since there is no ground requirement, the sensitivity to ground potential differences is greatly reduced. All that is required is that the signals remain within the input (common-mode) range of the receiver.

Susceptibility to crosstalk is also greatly reduced. The construction of a balanced transmission line requires that the two conductors be in close proximity to each other (without an intervening ground or power plane). This means that any transients induced in one conductor of a balanced transmission line will have the same (or nearly the same) transient (with the same magnitude and phase) induced in the other conductor. This crosstalk is, in effect, a form of common mode noise that (within limits) is rejected by the differential receiver.

In a balanced transmission line, the electric and magnetic fields exist *between* the two driven lines—there is no dynamic current flow in any present ground path. These fields are shown in *Figure 5*. The receiver is implemented as a differential amplifier that operates by comparing the amplitude *difference* between the two received signals.

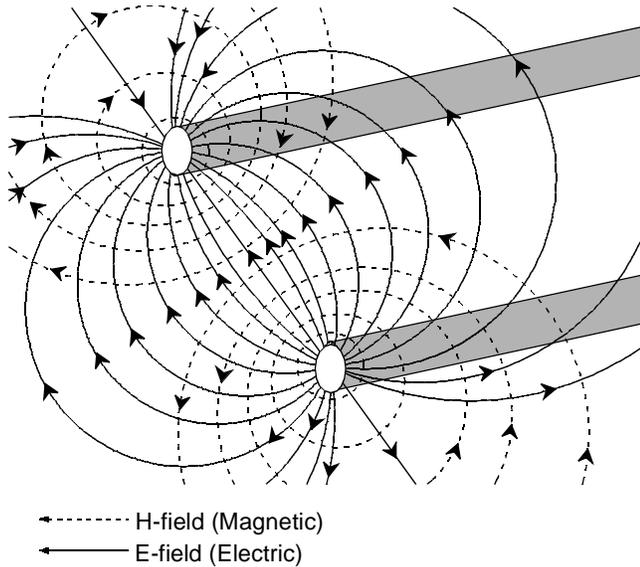


Figure 5. Electric and Magnetic Fields in a Balanced Transmission Line

### HOTLink Usage of Transmission Lines

When driving transmission lines with HOTLink, the first selection criteria is usually how far the signals must travel. For very short interconnects, the transmission line is often created using circuit board constructs that allow the high-speed signals to be routed across a card or backplane. For distances greater than a meter, cables of various configurations are generally used instead.

### Circuit Board Transmission Lines

Figure 6 shows the cross-sectional construction of the two primary types of circuit-board-based transmission lines. While other configurations are possible, the stripline and microstrip constructions follow standard circuit board manufacturing flows, and thus see the largest industry usage.

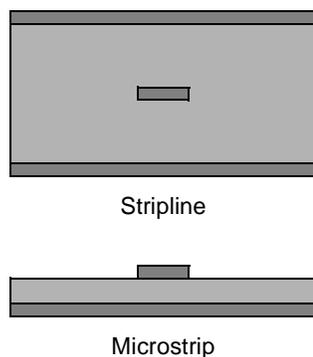


Figure 6. Circuit Board Transmission Lines

These types of transmission lines are used to route high-speed signals from a few centimeters to around a meter of circuit board. They are often routed through connectors as well as backplanes. Because of the relatively short distances

used with these types of transmission lines, they are usually considered to be lossless.

### Microstrip Transmission Line

Microstrip transmission lines are characterized by having a single strip-conductor spaced above a ground plane by a dielectric. This dielectric is usually the same material used for the remainder of the circuit board.

The key to using such a construct as a transmission line is stability of dimensions. Three dimensions determine the characteristic impedance ( $Z_0$ ) of the transmission line as shown in Figure 7: the width of the trace, the thickness of the trace, and the height of the dielectric.

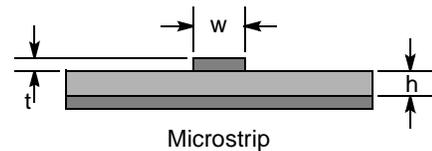


Figure 7. Microstrip Dimensions

With standard circuit boards the thickness of the trace is determined by the weight of copper specified for that specific (strip) layer. Standard thicknesses are usually specified in ounces; i.e., 1-ounce copper yields a trace 0.0356 mm (0.0014") thick. The width of the trace is specified in the artwork used to generate the circuit card, while the height of the trace from the ground plane is determined by the thickness of the laminate specified for the board construction.

A close approximation of the characteristic impedance of a microstrip transmission line may be calculated using Equation 4, where  $\epsilon_r$  is the relative dielectric constant of the board and  $w$ ,  $h$ , and  $t$  are the dimensions shown in Figure 7.

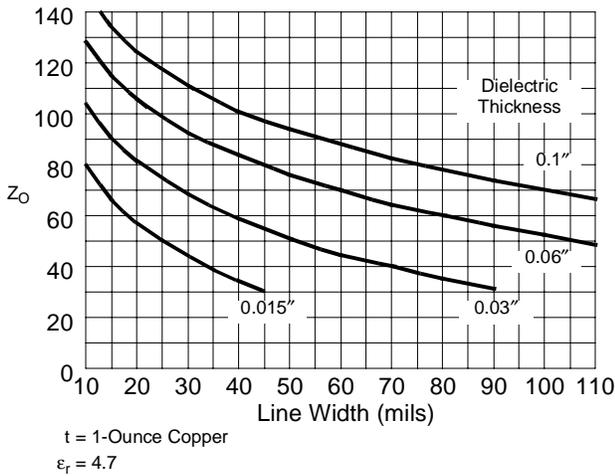
$$Z_0 = \frac{87}{\sqrt{\epsilon_r + 1.41}} \ln\left(\frac{5.98h}{0.8w + t}\right) \quad \text{Eq. 4}$$

This equation is an approximation and is not accurate for all ratios of width-to-height-to-thickness. Per experimental observation it does remain accurate ( $\pm 5\%$ ) for width-to-height ratios between 0.1 and 3.0 if the dielectric constant remains in the 1–15 range (Reference 2).

The transfer function for  $Z_0$  versus trace width for a microstrip transmission line is shown in Figure 8. All curves are based on standard FR4/G10-type laminate with 1-ounce copper. Varying the copper thickness has the least effect on the trace impedance. Going to 2-ounce copper will lower the trace impedance from 1–5%, while changing to 0.5-ounce copper will raise the impedance a similar amount.

Because of the variation in trace widths caused by etching, it is not advisable to use line widths under 10 mils for controlled impedance transmission lines. As the trace widths get smaller, the variation in line width has a much larger impact on trace impedance.

In a transmission line of this type some of the electromagnetic field propagates in the air above the strip conductor, while the remainder propagates through the circuit board dielectric. Because of this mixed medium, the  $V_P$  calculation for a microstrip transmission line (shown here in Equation 5) is different from that in Equation 3 (Reference 2).



**Figure 8. Calculated Impedance vs. Trace Width for Microstrip Transmission Lines**

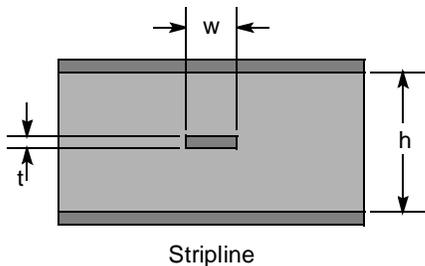
$$V_P = \frac{1}{\sqrt{0.475\epsilon_r + 0.67}} \quad \text{Eq. 5}$$

*Stripline Transmission Line*

Stripline transmission lines are characterized by having a single strip-conductor spaced between two ground planes by a dielectric. This dielectric is usually the same material used for the remainder of the circuit board.

Just as with a microstrip line, the key to using a stripline construct as a transmission line is stability of dimensions. Three dimensions determine the characteristic impedance ( $Z_0$ ) of a stripline transmission line as shown in *Figure 9*: the width of the trace, the thickness of the trace, and the height of the dielectric.

A close approximation of the characteristic impedance of a stripline transmission line may be calculated using Equation 6, where  $\epsilon_r$  is the relative dielectric constant of the board and  $w$ ,  $h$ , and  $t$  are the dimensions shown in *Figure 9*.



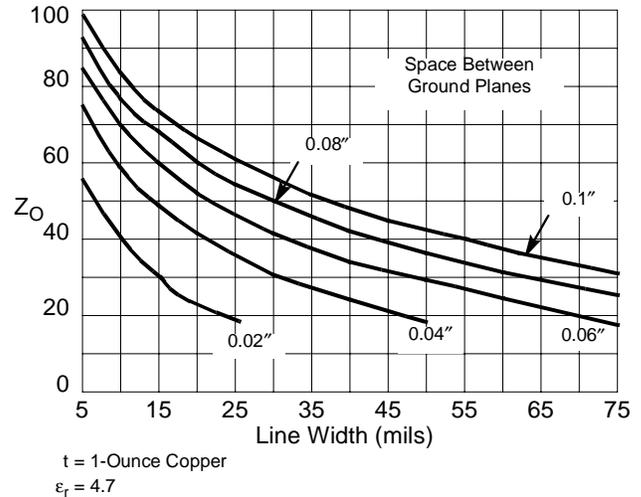
**Figure 9. Stripline Dimensions**

$$Z_0 = \frac{60}{\epsilon_r} \ln \left[ \frac{4h}{0.67\pi w \left( 0.8 + \frac{t}{w} \right)} \right] \quad \text{Eq. 6}$$

This equation is also an approximation and is not accurate for all ratios of width-to-height-to-thickness. Per experimental

observation it does remain accurate ( $\pm 5\%$ ) when  $w/(h-t) < 0.35$  and  $t/h < 0.25$  if the dielectric constant remains in the 1–15 range (Reference 2).

The transfer function for  $Z_0$  versus trace width for a stripline transmission line is shown in *Figure 10*. All curves are based on standard FR4/G10-type laminate with 1-ounce copper. Varying the thickness of the buried copper trace has the least effect on the trace impedance. Going to 2-ounce copper will lower the trace impedance from 3–8%, while changing to 0.5-ounce copper will raise the impedance from 1–5%.



**Figure 10. Calculated Impedance vs. Trace Width for Stripline Transmission Lines**

Unlike a microstrip transmission line, where part of the electromagnetic field propagates in air, in a stripline transmission line the field is bounded by the ground planes and must remain within the circuit board dielectric. This means that the  $V_P$  for a stripline transmission line is determined only by the dielectric constant and thus follows the calculation in Equation 3.

*Other Circuit Board Concerns*

When building microstrip transmission lines, interaction with the circuit board manufacturer is a must. To insure a constant dielectric thickness, the user should verify with their board manufacturer that a double-sided laminate is used (versus the B-stage or pre-preg layers) for this part of the circuit board. These “fill” layers in multilayer circuit boards cannot maintain the same dimensional stability between the strip-trace and the ground (or power) plane.

Verification of the relative dielectric constant should also be done. While often approximated at 4.7 for G10/FR4 substrates, this value can range from 4 to 6.

Unlike a microstrip transmission line, which may be forced onto a two-layer laminate for its construction, a stripline transmission line must, by its very nature, be composed of two separate circuit boards that are then laminated together in a multilayer assembly. This makes it much more difficult to control the dielectric height specification. The multilayer construction also raises the assembled board cost. This often limits the use of stripline transmission lines to those areas of

a design that require the additional shielding provided by the embedded strip construction.

Care must also be exercised in the placement position of the strip conductors relative to any significant discontinuities in the ground planes. As a general rule the strip should remain at least  $5(w+h)$  away from the discontinuity for microstrip, and  $5(w+ h/2)$  for stripline; e.g., do not route these types of transmission lines along the edges of cards.

Table 2 lists the relative dielectric constants for a number of common circuit board substrates. This dielectric constant alone determines the  $V_P$  (and the propagation delay) for the transmission line.

While it is theoretically possible to create a balanced transmission line on a circuit board, such construction is both difficult and costly due to the geometries involved. Almost all circuit-board-based transmission lines are unbalanced; i.e., transmitted as a signal relative to ground.

**Table 2. Properties of Circuit Board Substrates**

Material	Dielectric Constant	Prop Delay (ps/cm)	
		Microstrip	Stripline
G10/FR4	4.7	56.8	72.3
Mylar	5	58.2	74.5
Alumina	9.9	77.2	105
Teflon	2.1	43.0	48.3

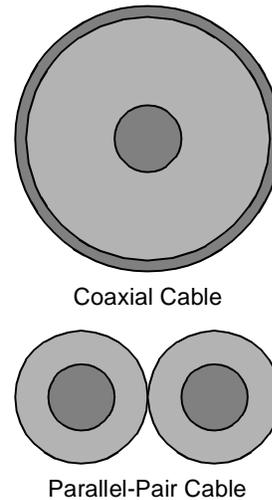
For those cases when the added noise immunity or other signal characteristics of a balanced transmission line are desired, the circuit may employ two unbalanced transmission lines that are then examined by the receiver differentially.

If matched delays are necessary in a system (for clock traces, pseudo-differential signals, etc.), do not attempt to route some of the signals as stripline and others as microstrip. The  $V_P$  calculations for each of these transmission lines are approximations based on specific dimensions that can vary significantly over manufacturing runs. By selecting either strip-line or microstrip for *both* transmission lines, the manufacturing variations present should affect both transmission lines in similar amounts and thus have minimal effect.

### Copper Cable Transmission Lines

Copper cables are generally used either for difficult signal routing (may even be used on a circuit board) or when long distances are involved. They have the advantage of being available in many configurations, with tightly controlled impedances, and allow communications at high bit-rates over hundreds of meters.

All copper cables fall into two categories: coaxial and parallel pair. The principal selection criteria between these two cable types is if the signal is to be transmitted unbalanced (coaxial) or balanced (parallel pair). The cross sectional construction of these two cable types is shown in Figure 11.



**Figure 11. Copper Cable Cross-Sectional Constructions**

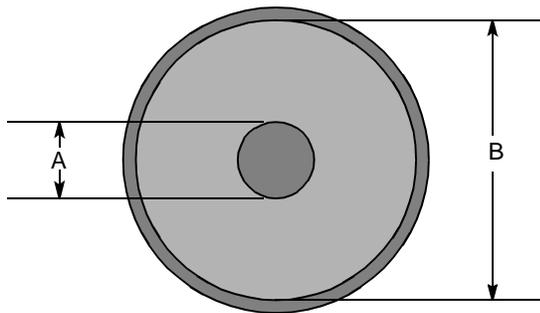
### Coaxial Cables

Coaxial cables are composed of two concentric conductors, maintained at a fixed spacing by a dielectric separator. This type of cable may only be driven in an unbalanced or single-ended form. They are available in flexible, semi-rigid, and rigid configurations, in diameters from 0.25 mm up to around 10 cm. Commercial cables are available in many impedances ranging from  $32\Omega$  to  $125\Omega$ , with the primary standards being  $50\Omega$ ,  $75\Omega$ , and  $93\Omega$ .

The  $50\Omega$  standard was developed for the military services in the early 1900s for use in radio broadcast. They needed a cable to feed vertical ground plane omnidirectional antennas which, by construction, had a  $50\Omega$  impedance. The  $75\Omega$  standard was adopted by the video and telecom industries because this impedance is the most efficient (considering only the voltages, currents, and powers to be driven) for transmission. The  $93\Omega$  standard was developed for the instrumentation industry to address their need for a low capacitance cable. Many other construction variants exist (triax, quadax, etc.) that differ primarily in the number and usage of the outer cable shield (Reference 3).

The characteristic impedance of a coaxial cable is determined by the ratio of its inner conductor to outer conductor diameters as shown in Figure 12, and the dielectric constant of the spacer material. This relationship is shown in Equation 7 (Reference 4).

$$Z_0 = \frac{138}{\sqrt{\epsilon_r}} \log_{10} \frac{B}{A} \tag{Eq. 7}$$



Coaxial Cable

Figure 12. Coaxial Cable Critical Dimensions

The entire electromagnetic field in a coaxial cable propagates through the dielectric (see Figure 1). This means that the  $V_P$  for a coaxial transmission line is determined only by the dielectric constant and thus follows the calculation in Equation 3. A comparison of the propagation velocities of common coaxial cable dielectrics is given in Table 3 (Reference 5).

Table 3. Propagation Velocity of Dielectrics

Insulation Type	$\epsilon_r$	$V_P$ Coaxial (%)	Prop Delay (ns/m)
PVC (Standard)	4–6	50–41	6.7–8.2
PVC (Premium)	3–5	58–45	5.8–7.5
Polyethylene	2.27	66	5.02
Polypropylene	2.24	67	4.99
Cellular Polyethylene	1.5	82	4.08
FR Polyethylene	2.5	63	5.27
FEP/TFE Teflon	2.1	69	4.83
Cellular FEP	1.4	85	3.94

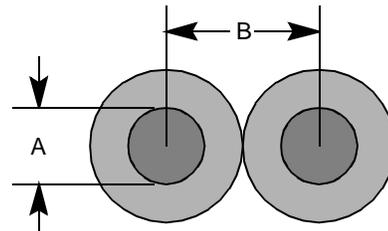
*Parallel-Pair Cables*

Parallel-pair cables are formed from two conductors, each having the same diameter, maintained a fixed distance apart from each other. This distance separation is usually maintained by the insulation around the individual conductors, but other types of spacers are also used.

While individual coaxial cables may only be driven in a single-ended (unbalanced) connection, parallel-pair cables may be driven either single-ended or differentially. What surprises many people is that the characteristic impedance for the cable is *different* depending on how the line is driven.

Equation 8 (along with the dimensions shown in Figure 13) is the standard equation used to calculate the  $Z_0$  for a parallel-pair transmission line. What is not usually identified is that this equation is only valid for differentially driven cables. When the exact same cable is driven single-ended (i.e., one line of the pair is a signal ground), the cable impedance is about 25%–35% lower (Reference 6).

$$Z_0 = \frac{276}{\sqrt{\epsilon_r}} \log_{10} \frac{B}{A} \quad \text{Eq. 8}$$



Parallel-Pair Cable

Figure 13. Parallel-Pair Cable Critical Dimensions

Equation 8 also makes the assumption that the entire electromagnetic field propagates through the dielectric. Except for those transmission lines that are either air dielectric (open wire) or a specialized construction, the propagation is actually split across multiple dielectric types and Equation 8 is generally not as accurate.

The  $V_P$  of a parallel-pair cable is also usually calculated using Equation 3, however the accuracy of this equation (because of the mixed dielectric) varies depending on cable construction. It is usually slightly faster than the calculation, which assumes only the physical (non-air) dielectric.

In theory, in a balanced transmission line the electromagnetic fields created around the two parallel conductors are equal in magnitude, but opposite in phase. The total field around such a transmission line has a net field-strength of zero; i.e., the fields cancel each other out and no energy is radiated. In actuality the two fields do not quite cancel. To do so would require both conductors to occupy the same physical space. To keep radiation to a minimum, the distance between conductors should be kept to no more than 1% of the signal wavelength (Reference 4).

Current balance is also important to minimize radiation. Because the fields generated are based on the currents present in the two conductors, any difference in the magnitude or phase of the driven signals will generate a different electromagnetic field. This difference, because it is not canceled out by the opposing field on the other conductor, radiates energy. This mismatch can be a significant contributor to EMI in a system.

Care must also be exercised in routing the conductors of balanced transmission lines to make sure that adjacent objects do not induce an unbalance into the system. If one of the two conductors is routed close to a ground or other conductors, the shunt capacitance can unbalance the line currents and increase radiation.

Two primary techniques are available to help reduce the interference affects of parallel-pair transmission lines, both from a radiation and from a susceptibility standpoint. The first of these is to twist the two conductors together at a controlled number of twists per unit length. In such a construction, the conductors must radially remain at the same centerline spacing throughout the twists to maintain the transmission line characteristic impedance. Average twist densities are from 1 to 0.1 twists per centimeter.

Twisting the lines together allows magnetic field cancellation and minimizes the affects of other nearby conductors. While the shunt capacitance still exists, it is now applied in nearly

equal amounts to both conductors, maintaining the field balance.

This same twisting also improves immunity to crosstalk in a system. With a true parallel-wire system, the currents induced by the fields present around an adjacent conductor are not always of the exact same magnitude on both conductors of a parallel-pair (due to the physical spacing between the conductors). The twists in a twisted-pair cable tend to bring both conductors into the same proximity of the noise generating conductor. This not only maintains the field balance in the cable, but also keeps the noise pickup truly common-mode, which can then be canceled by the receiver differential amplifier.

Twisted-pair cables also offer significant immunity to external e-fields (electric) and h-fields (magnetic). Because the signal wavelength is significantly longer than the twist-length on the cable, an external electromagnetic field's influence is spread across each propagating wave in multiple twists of the cable, each of which presents an opposite field intensity. These opposing fields tend to cancel out the affect of the external field.

The other method used to limit interference on parallel-pair conductors is shielding. A shield is an additional conductor surrounding both signal conductors in the parallel-pair. The purpose of this shield is two-fold: to constrain the electromagnetic fields generated by the transmission line, and to isolate external fields from this same transmission line.

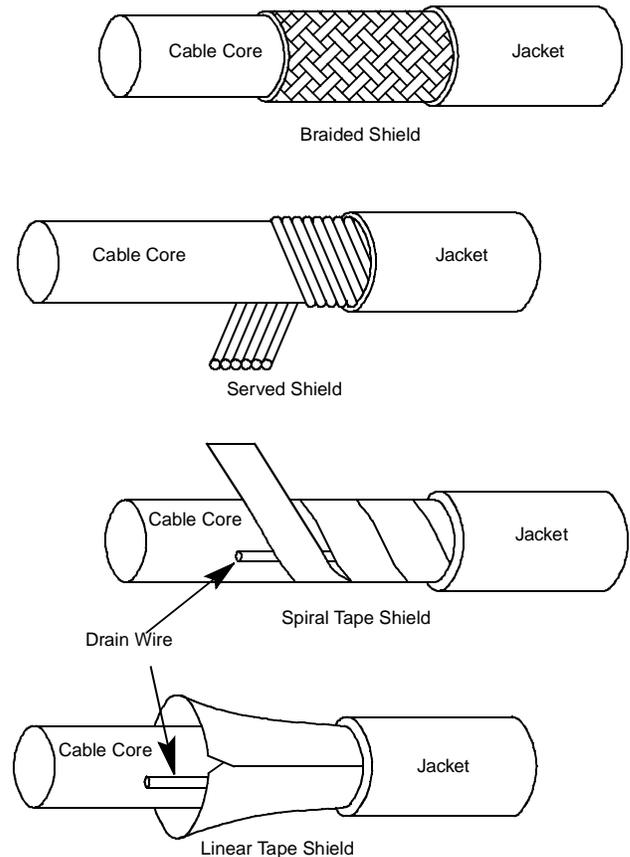
### Shields

Shields are used to keep what is outside out and what is inside in. How effective they are depends on their construction and how they are used in the system. *Figure 14* shows the construction of a number of different types of cable shields. Shields of these types operate as an electrostatic or Faraday shield. This means that they can block e-fields (electric) but offer only minimal protection from external h-fields (magnetic).

In *Figure 14* the part identified as the cable core could be any of the previously described cable types. In the case of coaxial cables the core, in its simplest form, would consist of a single conductor surrounded by its dielectric spacer, with the shield being the ground return conductor of the transmission line. Other constructions of transmission line cables can actually have multiple shields. In these configurations the cables are usually identified by the names triax (a center conductor, its ground, and an overall isolated shield) and quadax (a shielded parallel or twisted-pair cable with an overall isolated shield).

A perfect shield would be a seamless metallic tube running the length of the transmission line. Construction of this type is actually used for some forms of coaxial cable known as *hardline*.

For flexible cables, a compromise must be made. This compromise trades off shielding effectiveness for cable flexibility. Now instead of the shield being completely seamless, it has multiple seams that allow the cable to bend. These shields are made of either braided or spirally wrapped (served) layers of fine-gauge copper (sometimes aluminum if used as a secondary shield) wire, or spiral or linear-wrapped metallic tape.



**Figure 14. Cable Shield Constructions**

Braided shields consist of multiple groups of 34 to 40 AWG copper wire, braided together in a circular fashion around the core section of the cable. These strands may be bare copper but are often tin or silver plated. Shields of this type are rated in terms of braid coverage; i.e., how close to a seamless tube they get. Because of the high-frequencies present in a HOTLink-based serial connection, shield coverage should be a minimum of 85%. As a rule of thumb, if any dielectric is visible through the braid, there is insufficient coverage.

Served shields consist of the same fine-gauge copper wire wrapped in a continuous spiral around the cable core for the length of the cable. These strands may be tin plated, but are generally not silver plated. Cables of this construction should never be used for frequencies above 10 MHz because the spiral-wrap construction contains many long spiral gaps (especially near cable bends) that will leak EMI.

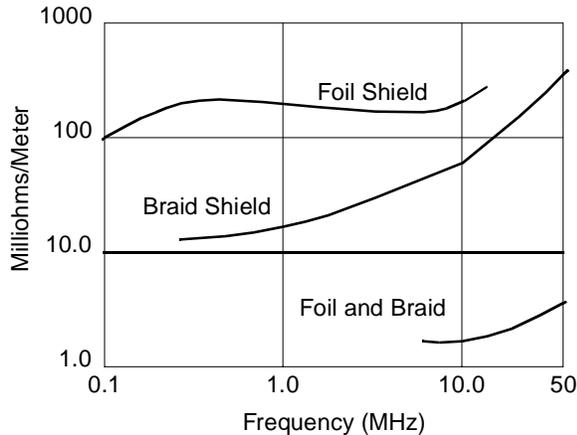
Metallic-tape shields are often used for high-frequency signaling because of the high degree of shielding coverage they provide. The metallic tape is made from either thin aluminum foil, or a plastic strip that is coated with aluminum on one or both sides. To allow termination of the shield at either end of the cable, and to make sure that each wrap of the shield tape is shorted together, these cables usually include an uninsulated drain wire that is in direct contact with the tape shield for its entire length.

Shields are often combined for even better shielding. Often a tape-shield will be covered by a braided shield. In this configuration the drain wire is eliminated because the braided shield performs the same function.

### Shield Transfer Impedance

One of the best ways to judge a shield's effectiveness is by its transfer impedance. This is a specification that relates how currents on one surface of a shield generate a voltage drop on the *other* surface of the shield. It is usually specified in  $m\Omega$ /meter of cable. The effectiveness of any shield is directly proportional to its transfer impedance. As the term impedance implies, this is a frequency sensitive parameter.

Because of their high DC resistance, aluminum-based tape shields do not fare very well in this measurement. Braided and served shields do much better due to their low-resistance copper construction. The best results are achieved by the combination of tape and a braided or served shield. *Figure 15* shows how shield construction effects transfer impedance.



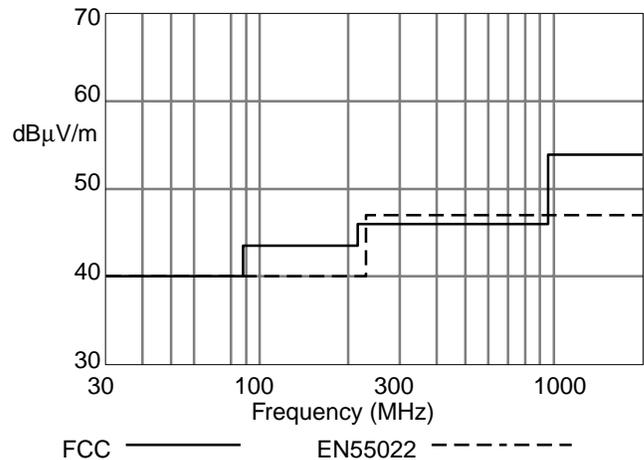
**Figure 15. Shield Transfer Impedance**

### Electromagnetic Compatibility

Shields are also necessary in many systems to allow equipment to meet various national and international electromagnetic compatibility (EMC) requirements. EMC deals with how much electromagnetic energy a piece of equipment is allowed to radiate, as well as how much external energy it must tolerate. Specific limits for both of these are set by a number of different international governing bodies. In the United States the limits for compatibility are set by the Federal Communications Commission (FCC) in Part 15 of their regulations. In Europe, the Common Market countries are now governed by a single EMC Directive in standards EN55022, EN55014, and EN60555-2, developed by CENELEC (Committee for European Electrotechnical Standardization). These standards deal with any digital equipment operating with any clocks or switching present at greater than a 9-kHz rate, and cover all frequencies up to 40 GHz.

For digital equipment, different limits are set for both radiated emissions as well as susceptibility depending on the target customer for the equipment. Equipment intended only for use in an industrial or business environment is classified as Class-A, while equipment that may be used in the home is

classified as Class-B. The radiated emission limits for Class-B are shown in *Figure 16* (Reference 9).



**Figure 16. FCC and EN 55022 Class-B Emission Limits at 3 Meters**

Under both of these classifications, it is necessary to test up to the 5th harmonic of the highest frequency clock present in the system. For HOTLink-based systems this could require testing up to 1.7 GHz.

### Coupling to Copper

There are three primary ways of coupling HOTLink to copper media: direct coupled, capacitor coupled, and transformer coupled. Each of these methods has different bias and termination requirements for the high-speed ECL signals.

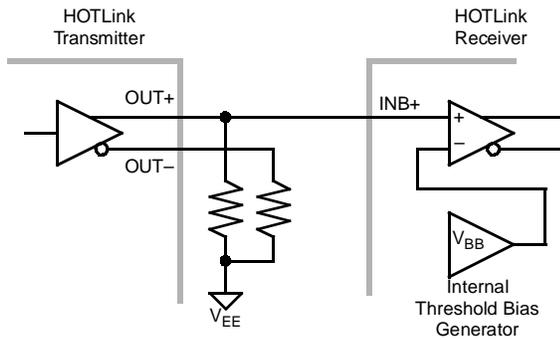
#### Direct Coupling

Direct coupling is where a DC path exists between the HOTLink Transmitter and Receiver on the high-speed serial interface. This coupling is used for those cases where both the transmitter and receiver operate from the same power supply and are in (relatively) close proximity to each other.

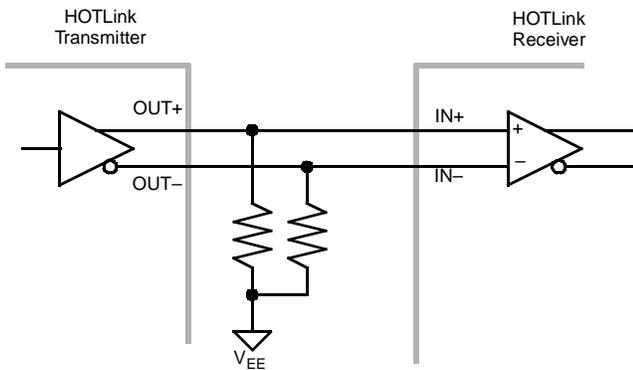
There are many subsets within this direct coupled area. These are differentiated by how far the signal must travel and the quantity of loads present.

#### Direct Coupled: <3 cm Length

For link distances under 3 cm, the serial signals do not have to be treated like transmission lines. In these cases all that is necessary is to bias the ECL signals so that they may properly switch. Because the transmission distance is so short, the signal may be assumed to be digital in nature. This allows the analog transmission concerns of longer distances to be minimized. A single-ended connection schematic is shown in *Figure 17*, while a differential connection is shown in *Figure 18*. Typical values for the pull-down loads are from  $250\Omega$  to  $510\Omega$ .



**Figure 17. Single-Ended Connection**



**Figure 18. Differential Connection**

Because the HOTLink Receiver does not provide an external  $V_{BB}$  reference, a single-ended connection may only be implemented using the INB+ input of the receiver. A differential connection may be implemented using either of the  $INA_{\pm}$  or  $INB_{\pm}$  differential inputs.

The ECL bias in both of these configurations is implemented with a single pull-down to  $V_{EE}$  on each driver output. While this bias configuration does generate more jitter than either a Thévenin or Y-bias, the amount is well under the jitter tolerance limits of the HOTLink Receiver for all supported frequencies.

*Direct Coupled: From 3 cm to 1 m Length*

Once the length of the connection becomes longer than 3 cm, the connection *must* be treated as a transmission line. This requires a termination network at the end of the transmission line. Because the connection is DC coupled, the termination network may also be used to bias the ECL output.

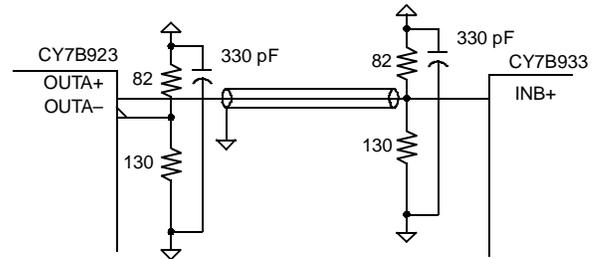
Unlike the previously described bias-only pull-down load, the network here must actually match the impedance of the transmission line. If it does not, a portion of the signal delivered into the transmission line is reflected off the termination and returned to the source. The amount of the reflection is determined by the voltage-reflection coefficient of the load,  $\rho_L$ , which is calculated using Equation 9.

$$\rho_L = \frac{\text{reflected voltage}}{\text{incident voltage}} = \frac{R_L - Z_0}{R_L + Z_0} \quad \text{Eq. 9}$$

Since this type of connection is only terminated at the destination, any signal reflected from the load will be returned to the source. However, because the source is *not* impedance matched to the transmission line ( $\rho_L \cong 1$ ), a large portion of the reflected signal it sees will again be reflected back down the transmission line.

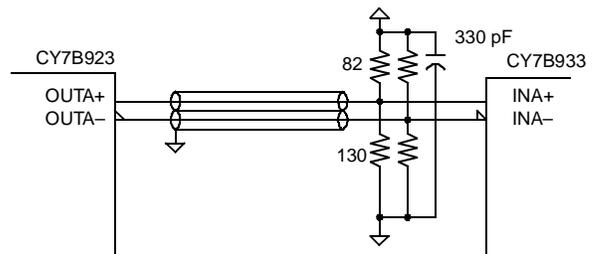
A reflection of this type will continue to travel back and forth between the two ends of the transmission line, being attenuated in amplitude both by the transmission line losses (very low for these short lines) and by the amount of signal absorbed in the terminations.

Figure 19 shows a single-ended connection using a Thévenin bias network. This network is sized for termination to  $V_{CC} - 2V$  of a  $50\Omega$  transmission line, and should be changed if other impedance transmission lines are used. A similar network is added to the OUTA- driver to keep a matched load on the differential driver. While shown in the schematic as a coaxial line, this would in most cases be implemented either as microstrip or stripline. Just as in Figure 17, the INB+ receiver is used for the single-ended connection.



**Figure 19. Direct-Coupled, Single-Ended Interface**

When implemented with two transmission lines (as shown in Figure 20), the signals may be examined differentially by the receiver. While not a true balanced transmission system, this configuration doubles the noise immunity of the single-ended configuration.



**Figure 20. Direct-Coupled, Differential Receiver Interface**

This type of connection is often called a balanced transmission line, but it is not. What actually exists are two single-ended (unbalanced) transmission lines that are examined differentially. Because the electromagnetic waves propagate independently down the two transmission lines, it is very important to make sure that both lines are the same electrical length from the driver to the receiver to allow the two signals to arrive in the same phase relationship they were sent.

### Direct-Coupled Bus

A common usage for HOTLink is as a data-mover on a backplane. In this configuration, the HOTLink Transmitters and Receivers are used to replace some of the wide buses on the backplane, along with their associated drivers, receivers, and connector pins. This usually provides a lower cost, lower power, and more reliable solution than the parallel interface it replaces.

### Single-Ended Bus

A bus of this type utilizes the wired-OR capability of ECL outputs to allow multiple sources on a common bus. Transmission line terminations are still necessary, and in fact must now be placed at *both* ends of the transmission line. Figure 21 shows a sample configuration of a single-ended multiple source and destination bus.

In this configuration, a HOTLink Transmitter and Receiver are located on a card plugged into a backplane. All the receivers are enabled at all times, and the transmitters are controlled using the FOTO signal such that only one of them is allowed to transmit at a time. Because of the single-ended operation of the bus, the INB+ input of the receiver should be used for serial data input.

The transmission line must be terminated at both ends to allow signals to be driven at any point along the transmission line. When the signal is launched into the line it effectively splits, with part of the signal traveling in each direction on the line. When the signal reaches the end of the transmission line it is absorbed into the termination networks. This double termination places a higher current burden on the driver. It sees two 100Ω transmission lines in parallel, which present a load of 50Ω.

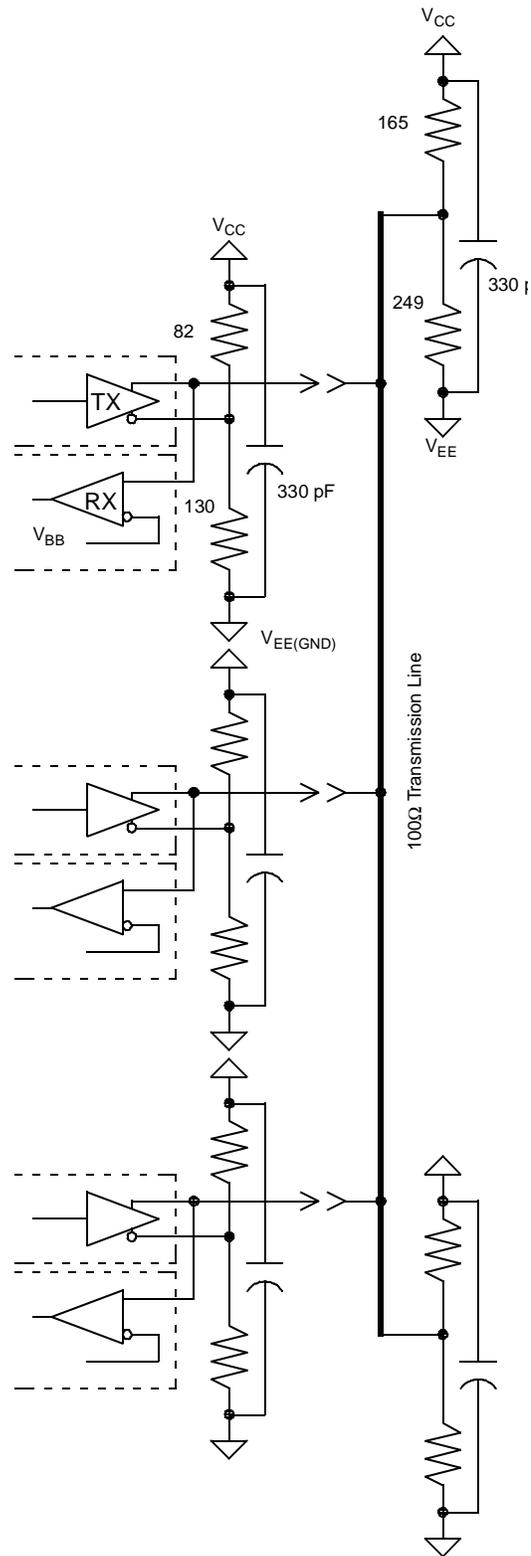
The complementary output of each differential driver must also see the same load as the true output to provide a balanced load for the driver. This requires adding a 50Ω Thévenin bias network for each driver present.

While implemented here with a 100Ω transmission line, other impedances may also be used. The lowest recommended transmission line impedance is 50Ω. This presents a 25Ω effective load on each attached driver.

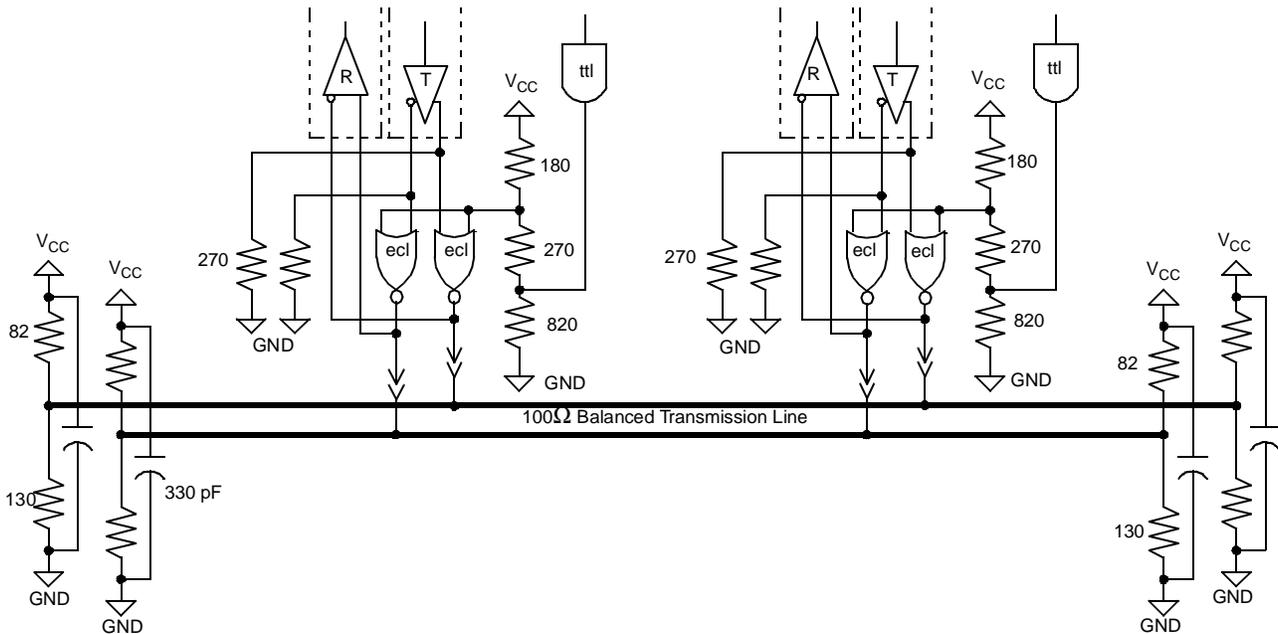
The physical implementation of a single-ended bus does have a few limitations. One of these is how many drivers/receivers can actually be on the bus. This is not a driver current limitation (HOTLink input currents are  $\ll 1$  mA), but is instead due to capacitive and stub effects. Each card on the backplane adds from 3 pF to 10 pF of capacitance to the bus. This added capacitance slows down the rising and falling edges of the signals. When operating in a single-ended environment, the maximum number of driver/receiver pairs should be limited to 20.

The physical placement of each driver/receiver pair is also critical to proper operation. Due to the construction of a backplane and its associated cards, each driver/receiver pair also adds a stub to the transmission line. The longer each stub, the more reflection/distortion it causes on the backplane.

These reflections are limited by placing the driver/receiver directly adjacent to the board/backplane connector. The signal route from the connector to the driver/receiver pair should be kept to under two centimeters in length.



**Figure 21. Single-Ended, Multi-Source Bus**



**Figure 22. Differential, Multi-Source Bus**

### Differential Bus

A single-ended bus of this type may be reliably used when the system noise is understood and within the margins of a single-ended ECL connection. For systems with more loads, more noise present, or those that may be exposed to large external noise sources, the bus may be implemented in a differential form. The transmission line may be created as either a 100Ω balanced transmission line, or two separate 50Ω transmission lines.

In the single-ended bus implementation, bus access is controlled using the HOTLink Transmitter FOTO pin. The FOTO pin was designed to disable the light output of optical modules by driving a differential logic-0 (OUT+=LOW, OUT-=HIGH) when the FOTO input is HIGH. Because the OUT- pin is still sourcing current when FOTO is HIGH, access for a differential bus must be controlled externally. This requires the addition of an external ECL multiplexer or differential driver with output disable capability, as shown in Figure 22. This driver operates by effectively disabling *both* sides of the differential driver from a single control input.

**NOTE:** With each transmission line being 50Ω, each single-ended driver will see a 25Ω load. While HOTLink outputs are designed to drive loads of this size, most ECL drivers are not. It may be necessary to change the transmission line impedance (and the resistor values in the associated termination networks) to operate correctly with other drivers.

The biggest problem in implementing such a structure is that true differential ECL multiplexers are rare, and those capable of disabling both outputs are fewer still. This function may be created from separate gates (requiring two ECL gates for each differential driver present). Being separate gates, these drivers also do not maintain the close current balance normally present in a true differential driver. To keep delays and cur-

rents as matched as possible, both gates should be in the same physical package.

These ECL parts are operated in PECL mode; i.e., they use the same V<sub>CC</sub> and ground as the HOTLink Transmitter and Receiver. Unlike the HOTLink Receiver A/B select pin (an ECL input), which may be controlled from a TTL environment using only two external resistors, these external ECL parts must use a three resistor divider. The third resistor is necessary to limit the V<sub>IH</sub> of the ECL input to no more than V<sub>CC</sub>-0.6V.

Some care must be exercised when selecting these external ECL parts. Because of the switching speeds present on the serial interface (>150 MHz) these parts must be 100K ECL or faster. In addition, because the connections between the HOTLink transmitter and these parts are effectively single-ended connections, the external ECL gates must also be temperature compensated to maintain noise margins.

One final concern deals with drive current. Unlike the HOTLink Transmitter, which can drive 25Ω loads, most ECL drivers can only handle 50Ω loads. If the backplane transmission line impedance is less than 100Ω, special bus drivers (e.g., F100123) or drivers with multiple outputs (e.g., F100313 with outputs tied in parallel) must be used to provide the necessary current. If these parts have differential outputs, the unused (complement) outputs should be attached to bias networks to provide a similar load as that seen by the used (true) output of the driver.

### Capacitive Coupling

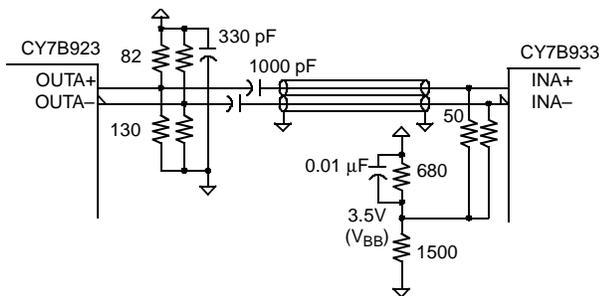
Capacitive coupling may be used for those connections where some reference difference may exist between the source (transmitter) and destination (receiver). This difference may be planned (e.g., true ECL communicating with PECL), or merely anticipated (e.g., possible ground or V<sub>CC</sub> differences). In both of these cases the capacitor is used to

block the DC signal component while allowing the AC components to propagate to the receiver.

This capacitively coupled interface is not recommended for cabling systems that leave a cabinet or extend for more than a few meters. This is primarily due to:

- Limited voltage breakdown of the coupling capacitors under ESD situations
- ESD susceptibility of the receiver due to transients induced in the cable
- Limited common-mode rejection at the receiver end

In a capacitive-coupled system, such as that shown in *Figure 23*, a bias network is still necessary at the driver to allow the output to switch. The preferred location for the DC-block capacitors is adjacent to the transmitter, immediately after the output bias network. This location is necessary due to the reactive nature of capacitors.



**Figure 23. Capacitive-Coupled, Copper Interface**

At the receiver end of the transmission line, the line must be terminated in its characteristic impedance. This is implemented using the two 50Ω resistors in *Figure 23*.

In addition to terminating the transmission line, the receive end must perform a DC restoration to place the received signals within the normal operating range of the HOTLink PECL receiver. This is done using a voltage-divider network.

In this configuration, the receiver reference point is set slightly different from that of a standard ECL receiver. Part of this is due to the HOTLink Receiver being designed for operation at +5V rather than -5.2V or -4.5V. The other is that the HOTLink Receiver has a wider common-mode range than standard 100K ECL parts. To allow operation over the widest range of signal conditions the external bias network on the receive end of the transmission line is set to the center of the HOTLink Receiver 3V common-mode range at  $V_{CC} - 1.5V$ .

While it is possible to bias and terminate the differential inputs with two Thévenin networks, this should not be done. The tolerance differences, even using 1% resistors, are enough to introduce offsets of >50 mV between the inputs. This offset lowers the system noise margin and increases the duty cycle distortion (DCD) jitter in the link. The bypass cap is used to keep the bias point stable by supplying current during any minor transients.

The transmission line in *Figure 23* is shown as two 50Ω unbalanced transmission lines. If the interconnect is implemented using microstrip, stripline, or coaxial cables, this is the type of connection that actually exists. In this dual-unbalanced

connection, the same equal-length restrictions of direct-coupled interfaces still exist.

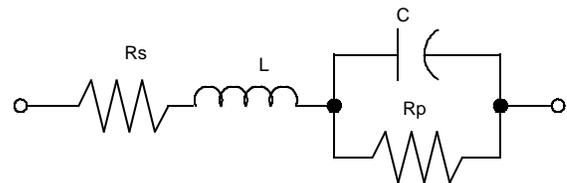
By replacing the two unbalanced transmission lines with a single balanced transmission line (unshielded twisted-pair, shielded twisted-pair, or twinax), it is possible to remove most of the equal-length concern of the conductors in the transmission line. In this configuration, the transmitter and receiver circuits remain the same, but the mode of propagation is now balanced (i.e., conductor-to-conductor, ground path not required). It is possible to also implement balanced transmission lines on circuit boards, but due to their shorter length this is not usually necessary.

A capacitively-coupled link may also be operated using a single piece of coaxial cable, but only with single-ended drive and reception. This requires giving up half of the received signal amplitude (only one driver is used), and connecting the INA- receiver input directly to the reference voltage.

### DC-Block Capacitor

While the desired affect of a DC-block capacitor is to block DC and pass all AC signal components (without loss), real life components don't operate in this fashion. Instead, a real capacitor blocks *most* of the DC, and passes frequency selective amounts of the AC signal components.

An equivalent model of a real capacitor is shown in *Figure 24*. In addition to the pure capacitance C, a number of parasitic resistive and inductive elements are also present. These parasitic elements determine the amount of leakage current, the ESR (equivalent series resistance), and where (in terms of frequency) the capacitor stops acting like a capacitor, and starts acting like an inductor. This frequency point is called the series-resonant frequency of the capacitor.



**Figure 24. Capacitor Equivalent Model**

The very small amount of DC current passed through a capacitor is called leakage current. For most designs this leakage is so small that it is undetectable relative to the AC signal components. The amount of AC signal passed varies with frequency, and is limited on the low end of the frequency spectrum by capacitance, and on the high end by parasitic inductance. This gives a capacitor a passband characteristic.

The amount of AC signal that is passed is controlled by the reactive characteristics of the capacitor, relative to that of the attached transmission line. For those frequencies below the series-resonant frequency of the capacitor, the reactance can be calculated using *Equation 10*. To allow efficient signal transfer, the  $X_C$  should be kept below 1Ω for the frequencies of interest.

$$X_C = \frac{1}{2\pi fC} \quad \text{Eq. 10}$$

Because the reactance of a capacitor varies greatly with frequency, placement of such a component between the receive

end of the transmission line and its termination network is not recommended. This is due to the reflections that would be caused by not terminating the transmission line in its characteristic impedance at all frequencies.

Placing such a capacitor directly adjacent to the driver removes much of this reflection problem. While the reflections still occur, they are absorbed as part of the rise and fall times of the source signal.

Good low-loss, RF-grade capacitors should be used for this application. These parts are available in many different case types and voltage ratings. They must be able to withstand not only the voltage of the signals sent, but any DC difference between the transmitter and receiver and the maximum ESD expected. A typical 1000-pF 50- $\Omega$  V COG/NPO capacitor is available in an 0805 surface mount case size (0.08"L x 0.05"W x 0.02"H). For on-board applications, a 50- $\Omega$  rating should be sufficient. While capacitors with much higher breakdown voltages are available, both cost and space make their use prohibitive. This same 1000-pF COG capacitor at 5-kV breakdown is almost a half cubic inch in size (Reference 7).

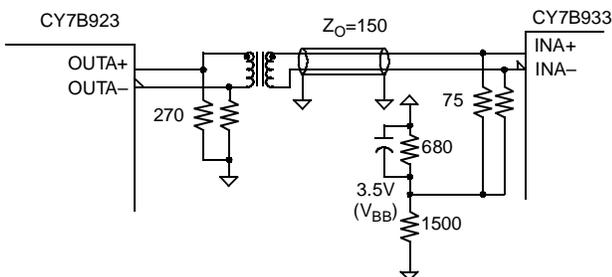
### Transformer Coupling

Transformer coupling is the preferred method for attachment to copper cables that extend for more than a few meters, or are operated between enclosures. Transformers have multiple advantages in copper-based interfaces. They provide:

- High primary-to-secondary isolation
- Common-mode cancellation
- Balanced-to-unbalanced conversion

The transformer is similar to a capacitor in that it also has passband characteristics, limiting both low and high frequency operation. Proper selection of a coupling transformer allows passing of the frequencies necessary for HOTLink serial communications.

The configuration shown in *Figure 25* uses only a single transformer, and either 150 $\Omega$  twinax or twisted pair as the transmission line. This can be done because the transmission system remains balanced end-to-end. Here the primary functions of the transformer are to provide isolation and common-mode cancellation.



**Figure 25. Transformer-Coupled Copper Interface**

In a single transformer configuration the transformer should be placed at the source end of the cable. Unlike the HOTLink differential receiver, which has a full 3V common mode range, an ECL output (when sourcing a zero or LOW-level) will respond to high-going signals picked up on the transmission

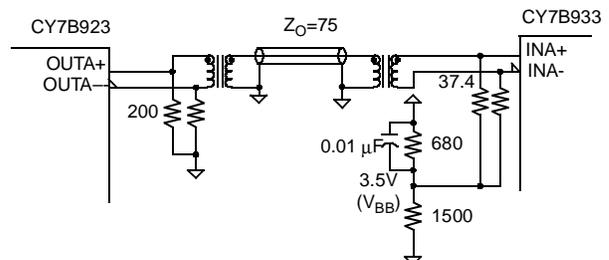
line. If a shield is present, it should be grounded at both ends to an earth or chassis (not signal) ground.

The transmitter shunt-bias network shown in *Figure 25* was selected to provide the maximum signal amplitude into the transmission line, rather than the most symmetrical edges. This configuration gives the highest signal-to-noise ratio at the receiver, but has different slopes of the rise and fall times at the transmitter.

These asymmetric rise and fall times do not add to the system jitter. Instead, the true and complement outputs combine in the transformer to provide a single signal with symmetrical rise and fall times. This insures matched transmission line currents for balanced transmission lines. This bias arrangement also has the advantage of delivering the entire transmitter output signal-swing into the transformer, rather than part into the transformer and part into the bias network. In a standard Thévenin bias or bias to  $V_{TT}$ , the source signal amplitude divides across the load (transformer) and the bias network, causing a significant amplitude loss.

This transformer-coupled configuration has many similarities to the capacitively coupled interface. It still provides DC isolation between the HOTLink Transmitter and Receiver, and requires the  $V_{BB}$  bias (DC-restoration) and termination network at the receiver.

In *Figure 26* a second transformer is added to the transmission system at the destination end of the cable. This configuration allows use of either balanced or unbalanced (coaxial) transmission lines. The configuration shown here is a 75 $\Omega$  coaxial cable system. Here, the first transformer is used for balanced-to-unbalanced conversion, while the second transformer provides unbalanced-to-balanced conversion. With transformers at both ends of the cable, much larger amounts of common-mode noise may also be handled.



**Figure 26. Dual Transformer-Coupled Copper Interface**

The size of the transmitter bias resistors are reduced here to handle the larger current requirements of the load. When driving a common load from a differential source, each driver sees a load impedance of half the actual load present. With a 75 $\Omega$  cable present each driver sees a 37.5 $\Omega$  load.

### Quantitative Interface Comparison

The transformer-coupled interface is the only one recommended for all cable lengths and types. This configuration operates equally as well with very short (<1 meter) lengths as it does with tens or hundreds of meters. Numerous configurations of transformer coupling and biasing were evaluated to determine both how to best configure a HOTLink-to-transformer interface, and to find out how cable impedance affects these configurations.

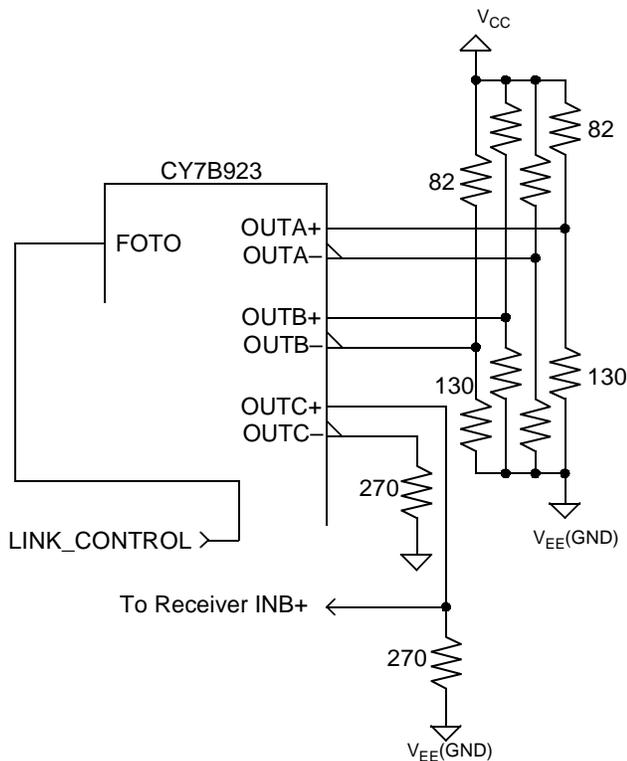
**Test Equipment**

The following equipment was used for the different evaluations:

- HP54100D 1-GHz Bandwidth Digital Sampling Scope
- HP8091A Rate Generator
- HP10240B DC Blocking Capacitor
- HP54002A 50Ω Pods
- Philips PM8919/09 500Ω 10:1 Probes (1.5-GHz Bandwidth)
- Technitrol/Pulse Engineering Transformers
- Cypress CY9266-C HOTLink Evaluation Boards

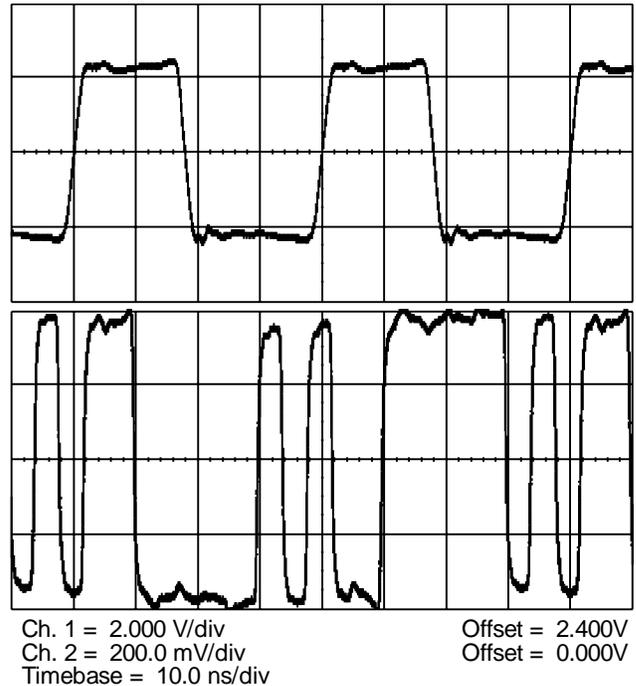
The primary goals of this testing were to determine how ECL operates when driving transformers, and what cable/coupling methods provide the best signal characteristics.

To get a good baseline for the following measurements, a HOTLink CY7B923 Transmitter was connected as shown in *Figure 27*. Measurements were made at the OUTB+ pin of the transmitter with the CY7B923 receiving a 25-MHz TTL clock. This clock is up-multiplied by ten inside the HOTLink Transmitter to generate a serial bit-time of 4 ns.



**Figure 27. Baseline Test Configuration**

The baseline waveforms for this configuration are shown in *Figure 28*. The top trace shows the TTL-level clock into pin 21 of the transmitter, while the lower trace shows the PECL-level signal on pin 28. Both enable signals on the transmitter (ENN and ENA) are disabled, causing the part to generate a continuous stream of alternating disparity K28.5s. This pattern is good for evaluating serial links because it contains the four



**Figure 28. Baseline Clock and Data**

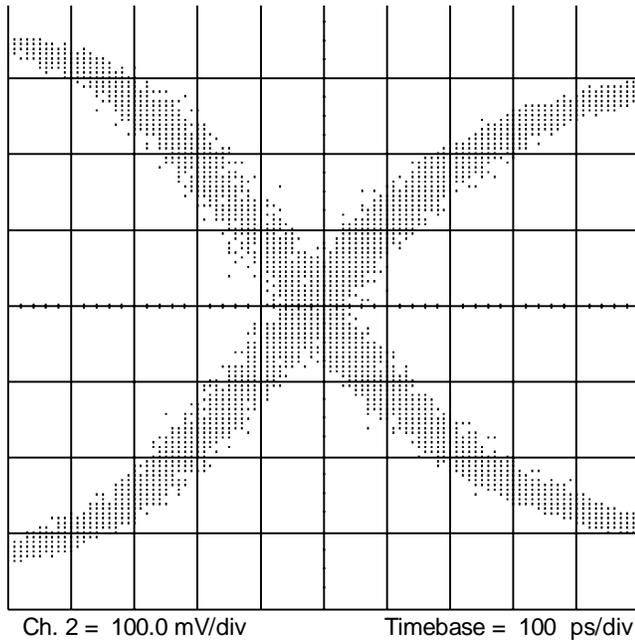
combinations of 1s and 0s necessary to test the characteristics of an 8B/10B code.

At this resolution it is difficult to see any real detail other than amplitude and period. To see the critical edge jitter it is necessary to zoom in on the rising and falling edges of the data. This is shown in *Figure 29*. Here the scope sweep rate has been increased by a factor of 100, going from 10 ns/division to 100 ps/division. The data crossover at the center of the figure is approximately 100 ps wide.

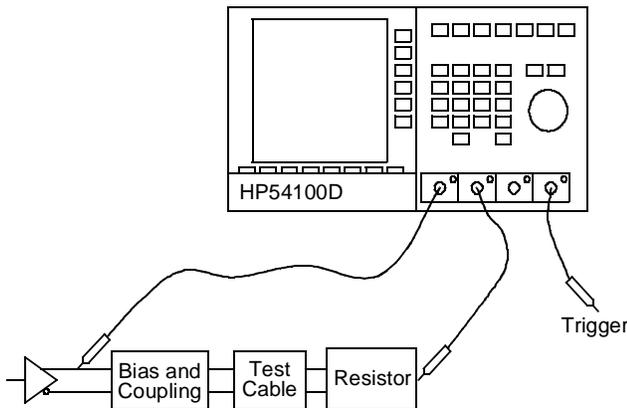
This 100 ps should not be assumed to be the output jitter of the HOTLink Transmitter (it is substantially less than this). It does not take into account the trigger accuracy of the scope, any jitter present in the trigger waveform, or any power supply ripple that the scope may view as additional jitter. However, since all the following measurements are taken with the same set-up and under similar trigger accuracy conditions, this value can be used to provide relative comparisons of different types of media and coupling.

**Test Set-Up**

The test set-up is shown in *Figure 30*. Low-impedance (500Ω) probes were used for all the high-frequency measurements. These probes, when combined with the scope amplifier, provide a measurement bandwidth of approximately 900 MHz. The probe impedance was factored into the bias and termination networks (where possible) to maintain the desired impedances.



**Figure 29. Baseline Jitter**



**Figure 30. Test Set-Up**

All probe connections were made using shielded probe-tip adapters to eliminate any measurement errors caused by probe ground-lead length.

All cable tests were performed using a single 30.4-meter segment (100 feet) of the specified cable. For those tests performed with a cable length of zero, the same test set-up as that shown in *Figure 30* was used, except that the termination resistor was placed directly on the output (secondary) of the coupling transformer.

### Test Configurations

The following test configurations were selected to determine how best to couple to coaxial media using transformers. Additional tests were added to either prove or disprove specific assumptions made in early ANSI Fibre Channel documents about how to couple using transformers. The selected configurations were:

- Thévenin bias, direct-coupled to transformer
- Thévenin bias, AC-coupled to transformer
- Transformer core saturation test
- Shunt bias, direct-coupled to transformer
- Shunt bias, high-frequency AC-coupled to transformer
- Single output, Thévenin bias, direct-coupled
- Single output, Thévenin bias, high-frequency AC bypass
- Single output, Thévenin bias, low-frequency AC bypass
- Dual transformers

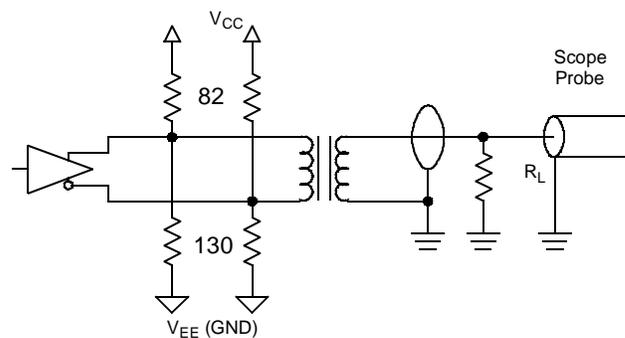
These different configurations (where applicable) were tested with three different impedance coaxial cables:

- 50Ω RG-58 (Belden 8219)
- 75Ω RG-59 (Belden 9259)
- 93Ω RG-62 (Belden 9269)

These specific cables were chosen because they provide the three primary cable impedances in a similar category of cable; i.e., they are all made with similar diameters and dielectric materials. This allows a better comparison to be made of the affect of cable impedance on jitter and attenuation.

### Thévenin Bias, Direct Coupled

The equivalent circuit for a Thévenin Bias differential driver, directly coupled to a transformer, is shown in *Figure 31*. At first glance this may appear to be the best way to couple a cable through a transformer. The bias voltage here is set by the pull-up/pull-down resistor ratio.

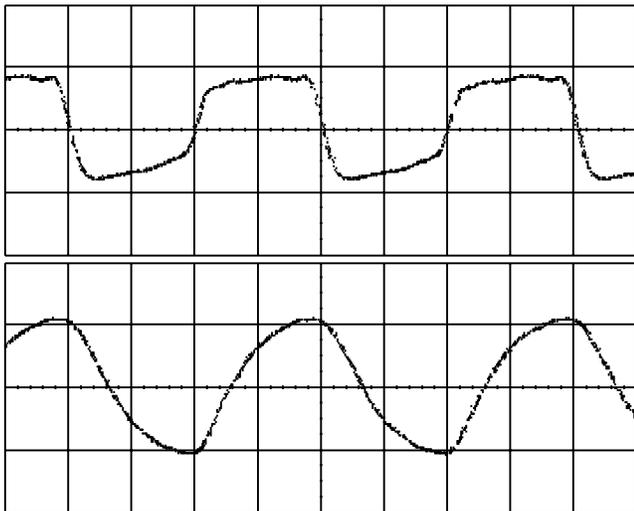


**Figure 31. Thévenin Bias, Direct-Coupled**

*Figure 32* shows the output of one driver on the top trace, and the output of the transformer secondary (when connected to a 50Ω resistive load) on the bottom trace. The primary observation to be made here is that the transformer secondary amplitude is almost equal to that of a single ECL driver. Since two drivers are actually present (differential drive), half of the signal is being lost somewhere.

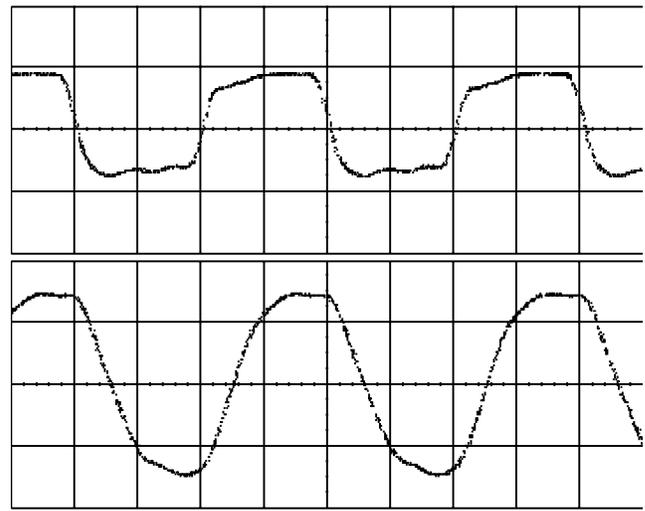
*Figure 33* shows the results when the load on the transformer was changed from 50Ω to 75Ω. Here the driver amplitude remains the same, while the secondary amplitude increases by approximately 50%.

*Figure 34* shows the results with a 93Ω resistive load. Now only a small improvement in output amplitude is seen, while the driver output becomes much closer to a square wave.



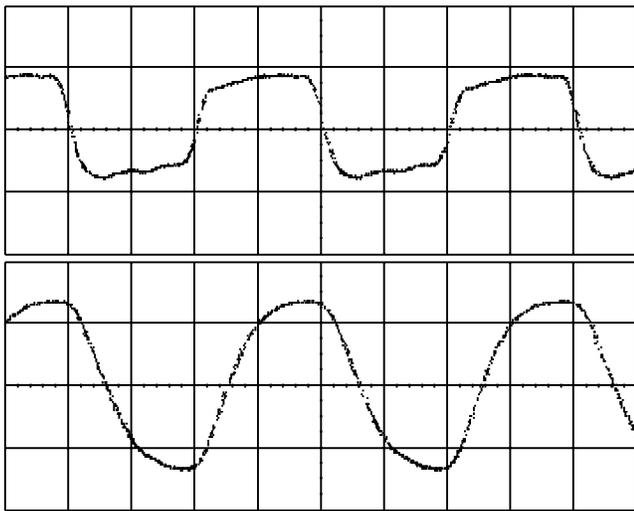
Ch. 1 = 400.0 mV/div  
Ch. 2 = 400.0 mV/div  
Timebase = 2.00 ns/div

**Figure 32. Thévenin Bias, Direct-Coupled, No Cable, 50Ω Load**



Ch. 1 = 400.0 mV/div  
Ch. 2 = 400.0 mV/div  
Timebase = 2.00 ns/div

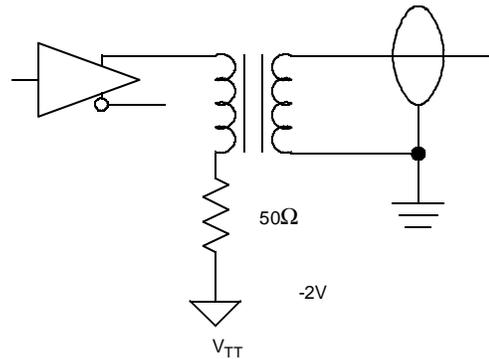
**Figure 34. Thévenin Bias, Direct-Coupled, No Cable, 93Ω Load**



Ch. 1 = 400.0 mV/div  
Ch. 2 = 400.0 mV/div  
Timebase = 2.00 ns/div

**Figure 33. Thévenin Bias, Direct-Coupled, No Cable, 75Ω Load**

The reason for these changes in output voltage with the different loads can be seen in *Figure 35*. Here the Thévenin bias network is converted into a resistor to a specific bias voltage. Under DC conditions, the impedance of the transformer primary approaches zero, while under AC conditions the impedance of the primary reflects that present on the secondary.

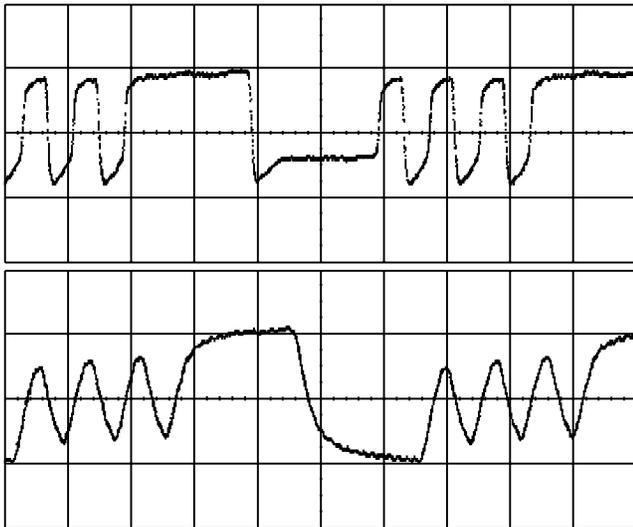


**Figure 35. Thévenin Bias Equivalent Circuit**

Placing a 50Ω load on the transformer secondary is equivalent to replacing the transformer primary with a 50Ω load. Because the Thévenin bias network is effectively in series with the primary, a voltage divider is created. Since both drivers are switching, half the amplitude of both of them is delivered to the load. With other load impedances, other divider ratios exist. The net effect of this type of biasing is that higher load impedances receive larger amounts of the total signal amplitude.

*Thévenin Bias with Cables*

Other effects can be seen when a terminated cable is attached to the transformer secondary instead of just a resistive load. *Figure 36* again shows the driver signal on the top trace, and the signal present at the end of 30.4 meters of RG-58 cable (50Ω) on the bottom trace. To see the effect of the run length limit of the 8B/10B code, a different pattern was selected that contains both long (5 zeros, 5 ones) and short (single-bit) pulses.



Ch. 1 = 400.0 mV/div                      Timebase = 10.0 ns/div  
 Ch. 2 = 400.0 mV/div

**Figure 36. Thévenin Bias, Direct-Coupled, with 50Ω Cable**

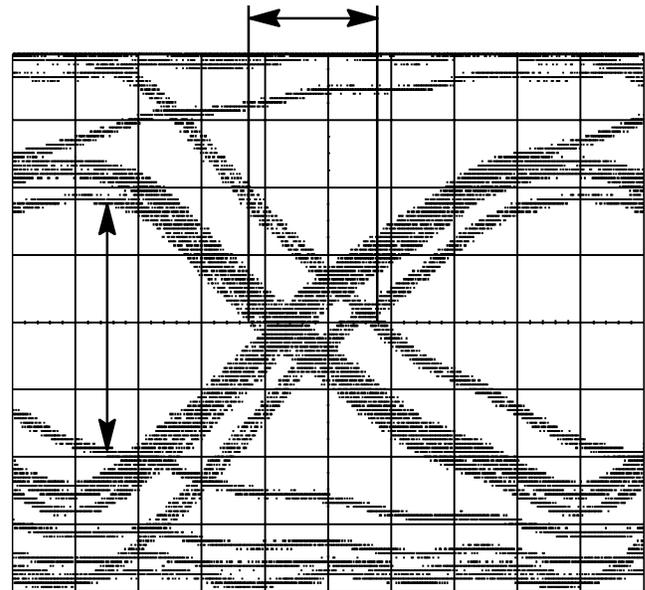
At the end of the cable (shown on the lower trace) the signal is quite different. Now the individual bit-transitions no longer remain centered vertically around the receiver threshold (center line of the lower waveform). This is due to a small DC offset built-up in the cable during the long-0 and long-1 pulses. During these long pulses, the transmission line has time to charge/discharge to near its maximum potential. During the shorter intervals, there is not sufficient time to fully charge or discharge the line. Under these conditions the transmission line is considered a long time-constant line.

Because the  $dv/dt$  rate for all transitions is effectively the same (regardless of the starting voltage), while the voltage change necessary to reach the receiver threshold is not, these long and short pulses are received shifted in time from nominal. This time shift is viewed at the receiver as a form of jitter called data-dependent jitter (DDJ).

As the length of the cable is increased, this difference in ending voltage between long and short transitions continues to increase. At some length of cable this difference becomes so great that the short transitions no longer cross the receiver threshold and the link becomes unusable. DDJ is one of the primary length-limiting factors of a copper-cable-based link.

Figure 37 shows the same signal as the bottom trace of Figure 36. The triggering and timebase have been changed to allow viewing of the individual bits in an overlay format called an “eye” pattern. The normal viewing of eye patterns has the eye opening (marked with the vertical arrow) in the center of the screen. This is used to see how large this opening is relative to a single bit time. The eye patterns shown in this (and following) figure is slightly time shifted, to allow central viewing of the signal crossing area.

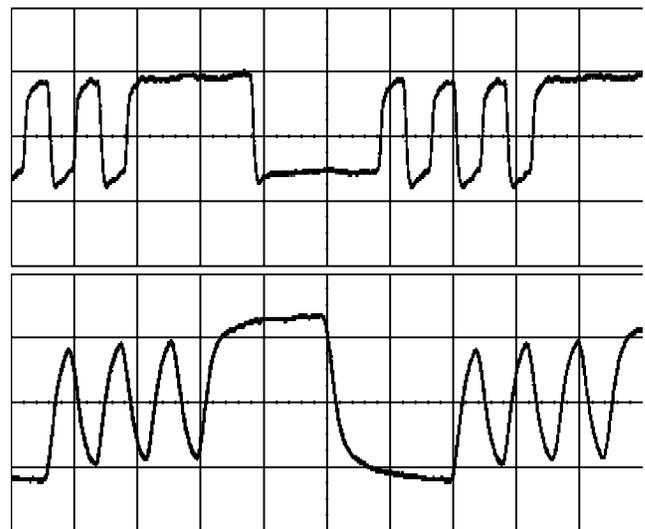
Figure 37 shows that the maximum usable amplitude of the eye is around 350 mV (marked with the vertical arrows). The jitter per bit (marked by the horizontal arrows) is around 1000 ps (25% of a single bit).



Ch. 2 = 100.0 mV/div                      Timebase = 500 ps/div

**Figure 37. Eye Diagram, Thévenin Bias, Direct-Coupled, with 50Ω Cable**

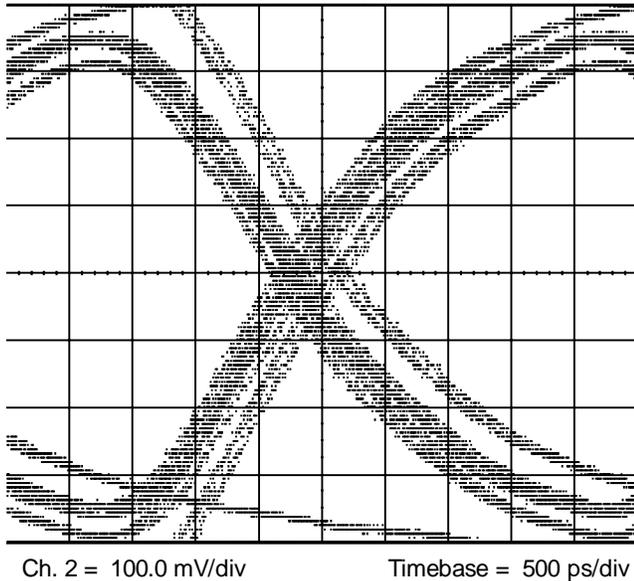
In Figure 38, this same configuration is tested using a 75Ω cable and termination. The signal amplitude at the end of the cable (bottom trace) has increased significantly from that of the 50Ω system. Taken as a percentage of the signal delivered to the destination, there is much less variation of peak signal amplitude from the short to long transitions.



Ch. 1 = 400.0 mV/div                      Timebase = 10.0 ns/div  
 Ch. 2 = 400.0 mV/div

**Figure 38. Thévenin Bias, Direct-Coupled, with 75Ω Cable**

Figure 39 shows the eye diagram for this 75Ω system. The usable amplitude here has increased to almost 600 mV; a 70% improvement over the 50Ω system. The amount of jitter present has also been substantially reduced, going to 700 ps or about 17% of a bit time.



**Figure 39. Eye Diagram, Thèvenin Bias, Direct-Coupled, with 75Ω Cable**

Figures 40 and 41 show the source and destination signals for a 93Ω system. The signal at the end of the cable has increased again up to 700 mV, while the jitter has been reduced to 500 ps (12%).

By comparing these three systems in Table 4, certain relationships become apparent. First, that as the cable impedance is increased, the signal amplitude delivered to the load is also increased. This amplitude increase also provides a better signal-to-noise ratio (SNR) at the receiver.

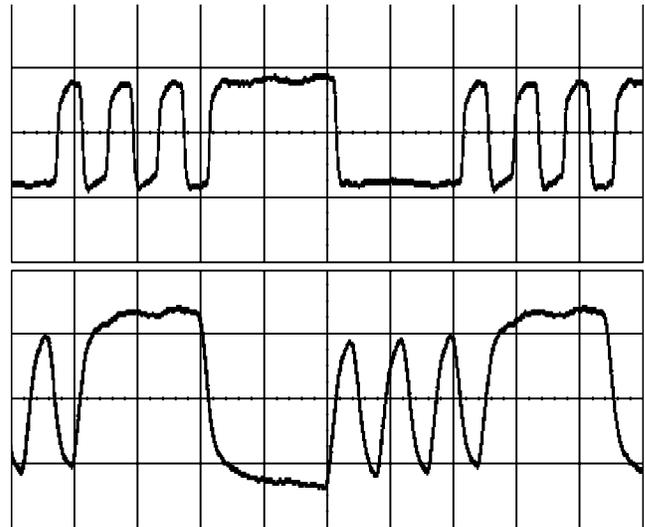
The second relationship is that as the impedance increases, the amount of jitter in the system is reduced. The ANSI Fibre Channel standard allows for links with up to 70% jitter at the receiver (Reference 8). While this standard only currently supports 75Ω coaxial cables (and 150Ω STP cables), these measurements show that 30.4-meter segments of 50Ω and 93Ω cable would also satisfy the maximum jitter specification.

*Thèvenin Bias, AC (Capacitive) Coupled*

The equivalent circuit for a Thèvenin bias differential driver, capacitively coupled to a transformer, is shown in Figure 42. Just as with the direct coupled system, the output bias voltage is set by the pull-up/pull-down resistor ratio. The capacitors now insure that no DC path exists through the transformer that might cause a core saturation that could limit both the bandwidth and energy transfer through the transformer.

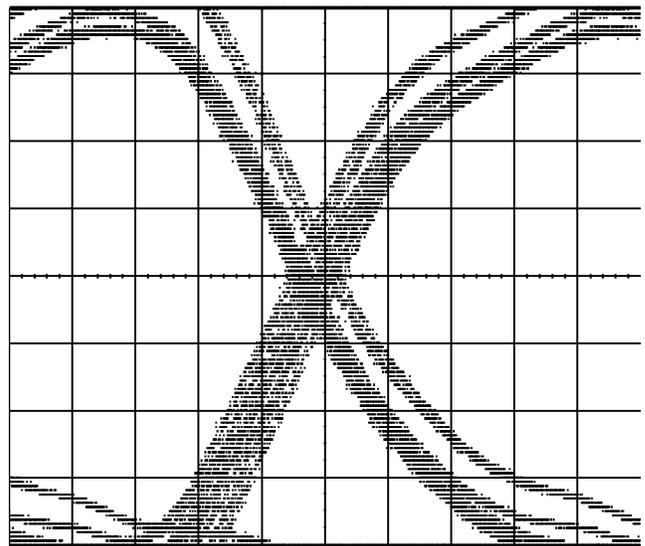
**Table 4. Cable Impedance Comparison**

Configuration	50Ω		75Ω		93Ω	
	Amplitude	Jitter	Amplitude	Jitter	Amplitude	Jitter
Thèvenin Bias, Direct-Coupled	350 mV	25%	600 mV	17%	700 mV	12%



Ch. 1 = 400.0 mV/div      Timebase = 10.0 ns/div  
Ch. 2 = 400.0 mV/div

**Figure 40. Thèvenin Bias, Direct-Coupled, with 93Ω Cable**



Ch. 2 = 100.0 mV/div      Timebase = 500 ps/div

**Figure 41. Eye Diagram, Thèvenin Bias, Direct-Coupled, with 93Ω Cable**

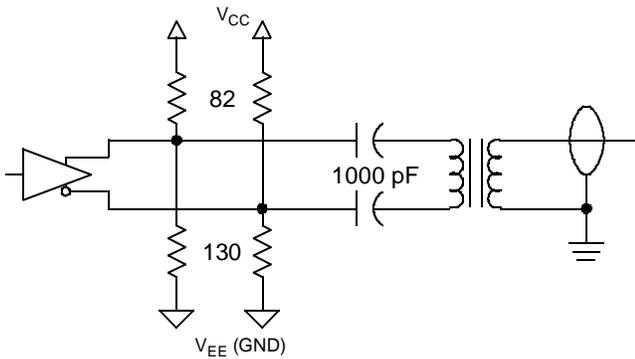


Figure 42. Thévenin Bias, AC-Coupled

The 1000-pF capacitors used here are RF-grade NP0-type parts. The passband of these parts (in this system) is such that they act like a high-pass filter. This limited low-end bandwidth can be seen on the long-0 and long-1 pulses in *Figure 43*. This figure shows the signal characteristics of all three impedances with a resistive load shown on the left, and 30.4-m cable and load on the right.

On a direct-coupled connection (at the transformer secondary), these long-1 and long-0 pulses switch to their HIGH or LOW state and remain there. In this AC-coupled configuration, these same pulses switch to the same HIGH and LOW levels, but slowly lose amplitude over the duration of the pulse. This amplitude loss is called *droop*.

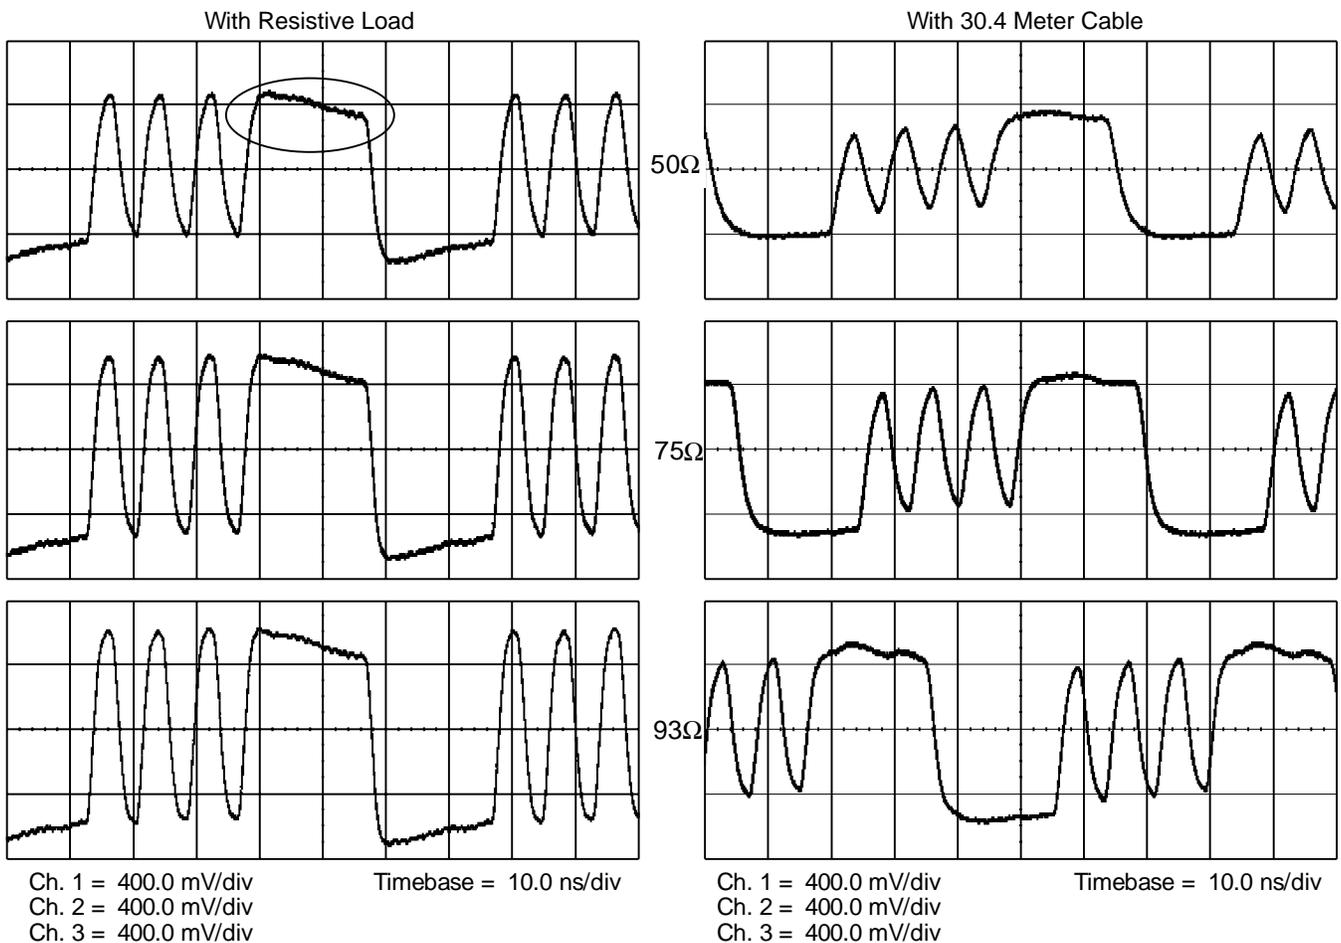


Figure 43. Thévenin Bias, AC-Coupled, with Resistive Load and Cable

This droop in many cases improves the signal characteristics at the load (receiver) end of the cable. Comparing the top right column trace in *Figure 43* with the bottom trace in *Figure 36* shows that the AC-coupled signal has a smaller peak amplitude for the long-duration pulses. This translates directly into a larger usable amplitude and smaller jitter percentage.

The capacitors in this link perform a rudimentary frequency-spectrum equalization. Because this equalization is performed prior to the signal being placed on the transmission line, it is called pre-compensation. A similar spectrum correction, when applied at the receiver end of the transmission line, is referred to as post-compensation or equalization.

These same signals are shown as eye patterns in *Figure 44*. *Table 5* compares the amplitude and jitter in these AC-coupled waveforms with the previous direct coupled configuration. The key observation made here is that the AC-coupling in all cases improves the amplitude and jitter. This improvement in all cases (with the specific coupling transformer and biasing evaluated here) is due to the limited bandwidth of the capacitor, not because there is no DC-path through the transformer. This was confirmed by actually forcing controlled amounts of DC through the transformer to determine where core saturation occurs.

### Transformer Core Saturation Testing

To validate that a small DC current flow (caused by a possible small mismatch in the ECL driver/load circuits) does not affect

the signal coupled through the transformer, a small modification was made to the previous AC-coupled test set-up (see *Figure 45*). This change involved the addition of two resistors (labeled R in *Figure 45*), attached to the primary of the transformer, to force a DC current through the primary. All tests were performed with a 50Ω resistive load on the transformer secondary.

To better see the effect, the data pattern was changed to use maximum run-lengths of six bits. While this is beyond the limits of the 8B/10B code, it serves to put the interface under greater stress.

The results of these tests are shown in *Figure 46*. The top trace shows the transformer secondary output with 13 mA of DC in the primary. The middle trace shows the same circuit with 30 mA of DC in the primary. The bottom trace shows 50 mA of DC in the primary. Notice that the secondary waveform starts to change around 30 mA, and is quite distorted at 50 mA. This means that the transformer core starts to saturate with around 30 mA of DC in the primary.

A normally biased and loaded ECL output can never have this much of a DC imbalance. This means that unless some type of pre-compensation is desired, there should be no need to AC-couple to the transformer primary.

**Table 5. Driver Coupling Comparison**

Configuration	50Ω		75Ω		93Ω	
	Amplitude	Jitter	Amplitude	Jitter	Amplitude	Jitter
Thèvenin Bias, Direct-Coupled	350 mV	25%	600 mV	17%	700 mV	12%
Thèvenin Bias, AC-Coupled	400 mV	20%	650 mV	15%	750 mV	11%

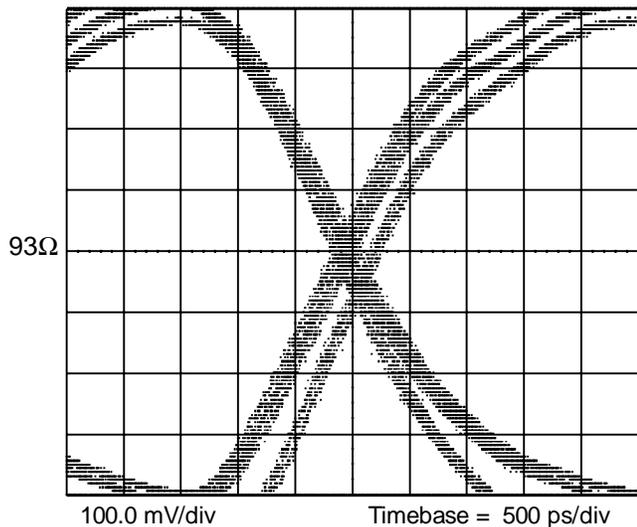
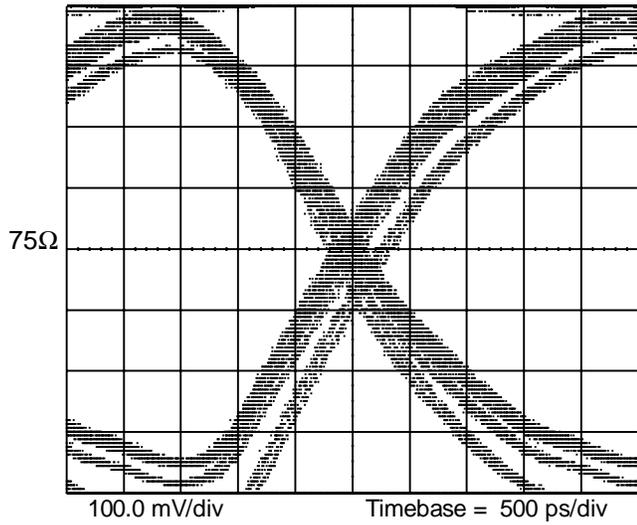
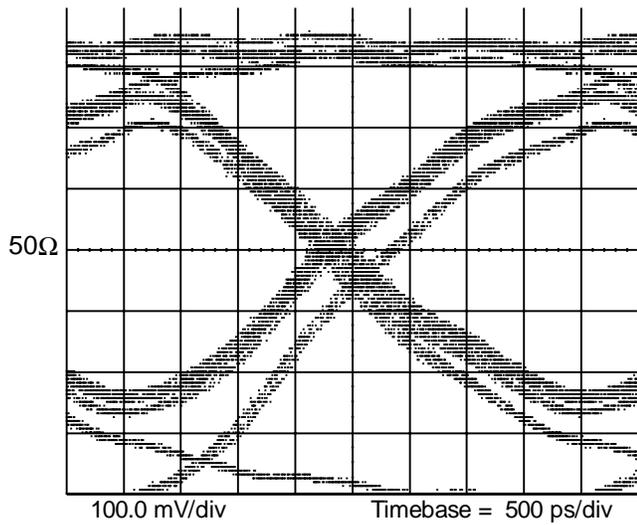


Figure 44. Eye Diagram, Thévenin Bias, AC-Coupled, with Cable

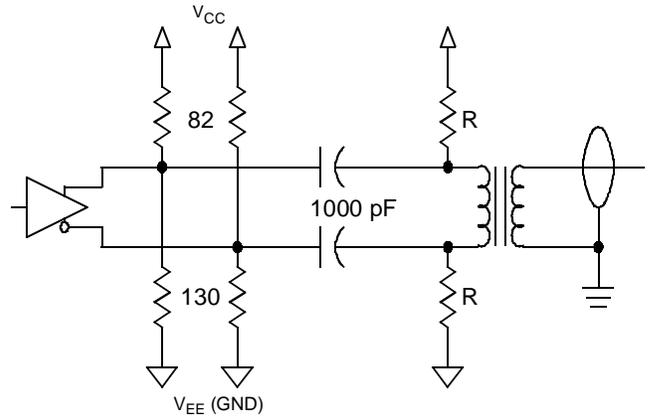
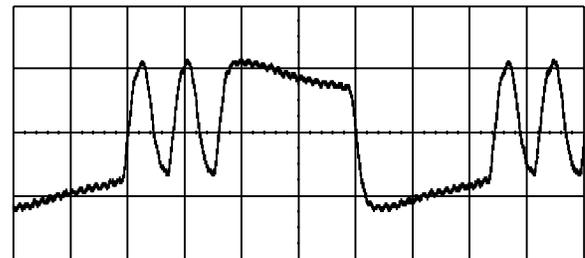


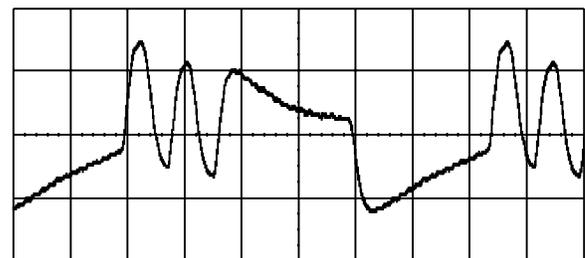
Figure 45. Transformer Core Saturation Test Fixture



13mA Primary Current

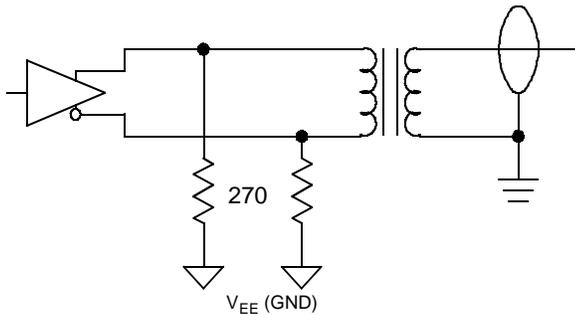


30mA Primary Current



50mA Primary Current

Figure 46. Transformer Core Saturation Test



**Figure 47. Shunt Bias, Direct-Coupled**

*Shunt Bias, Direct-Coupled*

A shunt bias, where a single resistor is attached from each PECL output to  $V_{EE}$  (ground), is normally used only for digital logic applications. This is due to the slightly different rise and fall times generated as the outputs switch. When used to drive a wideband transformer, as shown in *Figure 47*, this bias method has some distinct advantages.

First, it only requires a single resistor per driver, unlike the Thévenin bias which requires two resistors and a bypass capacitor. Second, and probably more important, this configuration allows much more of the ECL driver's signal swing to be seen on the transformer secondary.

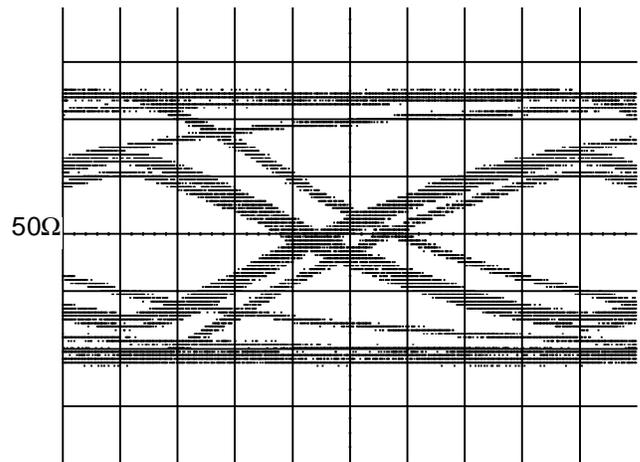
The signal transmission characteristics of this type of coupling are shown in *Figure 48*. This details the eye patterns for all three cable impedances. Unlike the previous eye diagrams, which could be displayed at a 100 mV/div scale, these signals are now shown at 200 mV/div.

Because these signals are all direct coupled, the jitter measurements are back around where they were with the direct-coupled Thévenin bias setup. The received signal amplitude however has increased around 100 mV over that of a Thévenin bias. This shows that the system jitter is independent of the signal drive level.

*Shunt Bias, AC-Coupled*

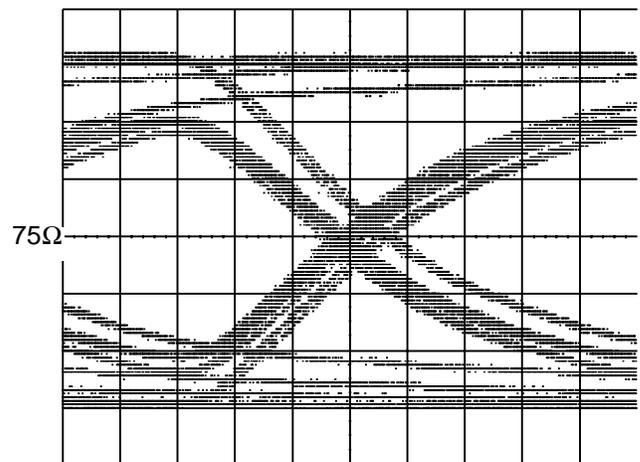
By combining the improved amplitude of a shunt-bias coupling with the limited frequency response of a capacitively coupled system, it is possible to squeeze out a slightly better signal.

The circuit for this configuration is shown in *Figure 49*. Here the capacitors again serve to block some of the lower frequency spectral components, which are not as severely attenuated by the transmission line.



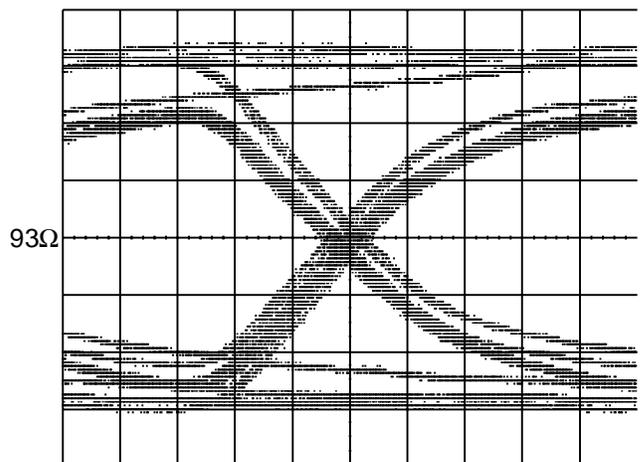
200.0 mV/div

Timebase = 500 ps/div



200.0 mV/div

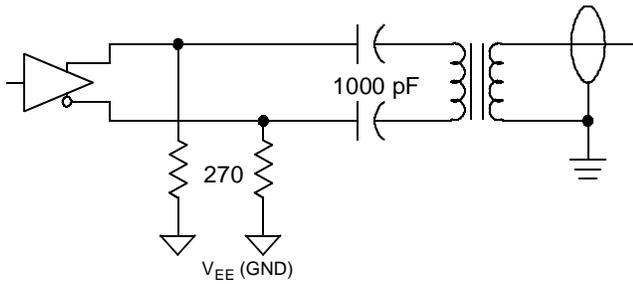
Timebase = 500 ps/div



200.0 mV/div

Timebase = 500 ps/div

**Figure 48. Eye Diagrams, Shunt-Bias, Direct-Coupled, with Cable**


**Figure 49. Shunt Bias, AC-Coupled**

The signal transmission characteristics of this type of coupling are shown in *Figure 50*. This figure details the eye patterns for all three cable impedances. These eye diagrams are again displayed at 200 mV/div.

The received signal amplitude in this shunt bias, AC-coupled configuration continues to operate as a function of cable impedance. As the cable impedance is increased, the received signal amplitude grows larger, and with less jitter.

The effect of the coupling capacitor on the circuit is more prominent on the lower impedance cables. On the 93Ω cable, the jitter improving effect (with this short length of cable) is basically nonexistent. With longer cables it is expected that this should have a much larger effect.

A quantitative comparison of all four configurations is shown in *Table 6*.

### Single Transformer Configurations

All of the previous coupling circuits were based on a differential driver working into a common load. These configurations allow the amplitude swing of both drivers to be presented to

the load. While this is expected to be the primary coupling mode for copper interconnect, it is also possible to drive these same connections through a single driver.

#### Single Driver, Direct-Coupled

A Thévenin-biased direct-coupled configuration is shown in *Figure 51*. When coupled in this mode it is possible to double the number of connections driven from a single source, at the expense of approximately 6 dB of amplitude on the cable.

This loss of amplitude has minimal affect on how far a signal can be driven on a copper cable. HOTLink-based copper links (when used without equalization) are generally limited by jitter accumulation rather than attenuation. The amplitude loss may effect the bit-error-rate for the link due to the reduced noise margins.

*Figure 52* shows the signal characteristics for this configuration when driving a 50Ω resistive load. Here the top trace shows the output of the driver while the bottom trace is at the transformer secondary. While the traces may look similar, the bottom trace is shown at a different vertical resolution. In effect only half of the driven signal is appearing at the load. This is again due to the voltage divider that exists between the transformer and the Thévenin bias network

Comparing the top trace in this figure with the same trace in *Figure 36* shows that the low side distortion is now gone. The pulses also are much more squared-off in this single driver configuration.

#### Single Driver, Direct-Coupled, AC Bypass

With the Thévenin bias network, both AC and DC signal components are dissipated in the network. By capacitively shunting the Thévenin network, it is possible to drop the DC signal component across the bias network, and drop the AC component across the transformer's primary. This configuration is shown in *Figure 53*.

**Table 6. Shunt vs. Thévenin Bias Comparison**

Configuration	50Ω		75Ω		93Ω	
	Amplitude	Jitter	Amplitude	Jitter	Amplitude	Jitter
Thévenin Bias, Direct-Coupled	350 mV	25%	600 mV	17%	700 mV	12%
Thévenin Bias, AC-Coupled	400 mV	20%	650 mV	15%	750 mV	11%
Shunt Bias, Direct-Coupled	400 mV	25%	700 mV	16%	800 mV	11%
Shunt Bias, AC-Coupled	500 mV	17%	800 mV	15%	850 mV	12%

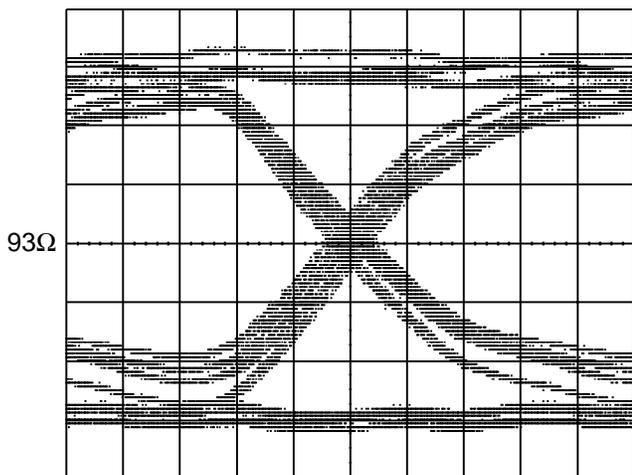
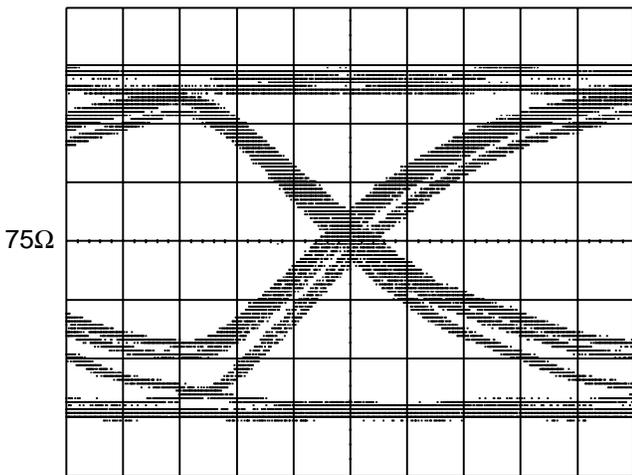
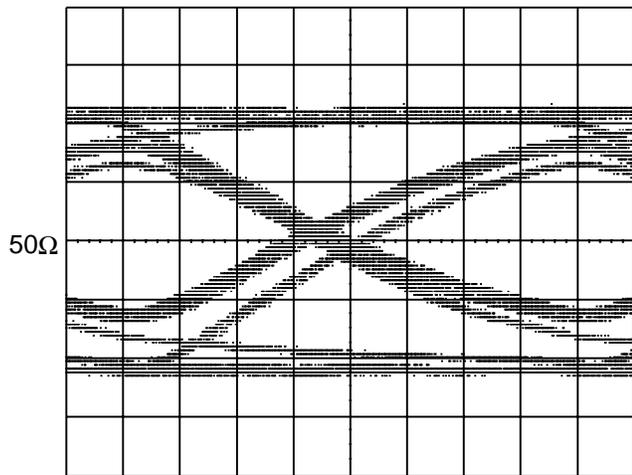


Figure 50. Eye Diagrams, Shunt Bias, AC-Coupled, with Cable

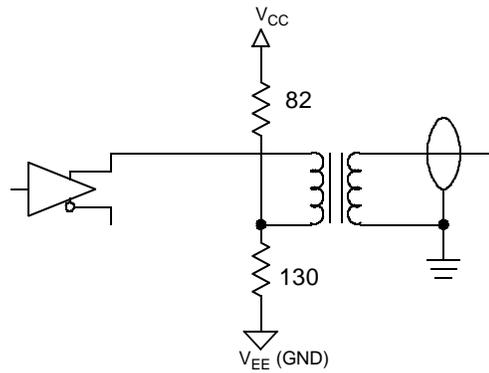
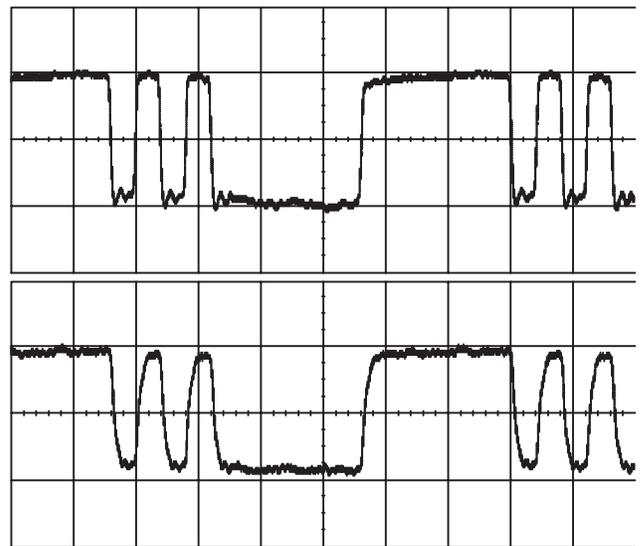


Figure 51. Single Driver, Thévenin Bias



Ch. 1 = 400.0 mV/div  
Ch. 2 = 200.0 mV/div  
Timebase = 10.0 ns/div

Figure 52. Single Driver, Direct-Coupled, 50Ω Resistive Load

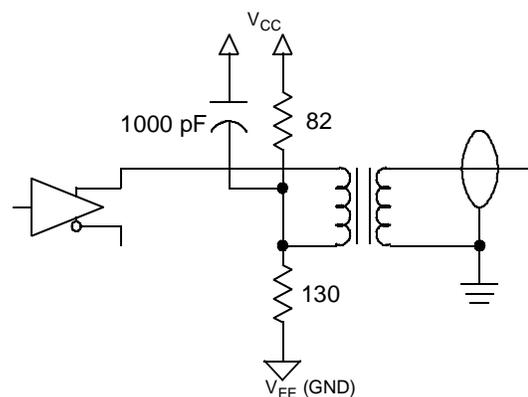
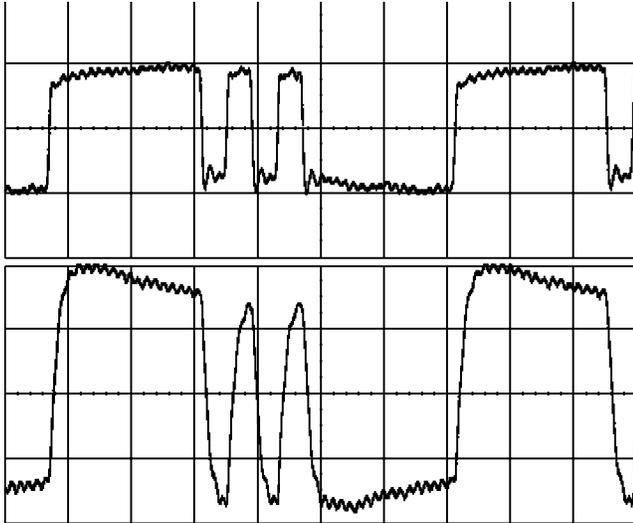


Figure 53. Single Driver, Direct-Coupled, High-Frequency Bypass

The added capacitor effectively doubles the signal delivered to the load. Because of the size of capacitor selected here, there is some limiting of the low-frequency signal components. These affects are shown in *Figure 54*.



Ch. 1 = 400.0 mV/div Timebase = 10.0 ns/div  
Ch. 2 = 200.0 mV/div

**Figure 54. Single Driver, Direct-Coupled, High-Frequency Bypass, 50Ω Resistive Load**

The capacitor again provides a small amount of precompensation to the circuit. This configuration tends to increase the source end jitter, while decreasing the jitter at the end of the cable.

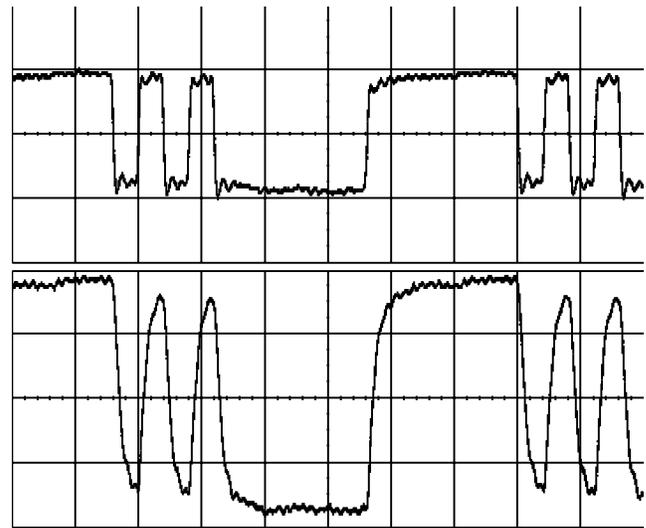
Replacing the 1000-pF capacitor with a 0.027-μF part significantly changes the AC passband characteristics of the coupling network, as shown in *Figure 55*. Now the low-frequency signal components that were blocked by the small 1000-pF capacitor are allowed to couple through the transformer. This configuration provides minimal jitter at the transformer secondary, but has more at the end of the cable than the high-frequency bypass configuration.

**Dual Transformers**

In the previous differential coupling configurations where a single transformer was driven at both ends, the possibility existed of one driver having an effect on the other. To see if any such affect was present, tests were performed that used separate transformer primaries to drive a common load.

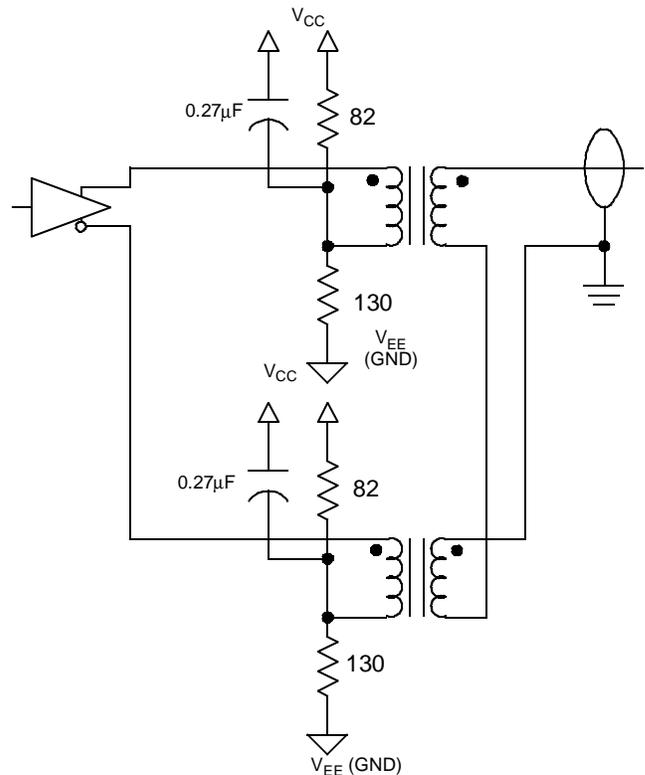
Based on the excellent waveform results achieved from a single driver/transformer configuration, the configuration in *Figure 53* (with the larger 0.27-μF capacitor) was duplicated on the complement output of the differential driver. With each of these circuits operated into separate 50Ω resistive loads, the waveforms remain the same as those shown in *Figure 55*.

When connecting the secondaries of these two transformers in series (as shown in *Figure 56*), remember that the polarity of the signals of the transformer attached to the complimentary driver are 180° out of phase with those of the true driver.



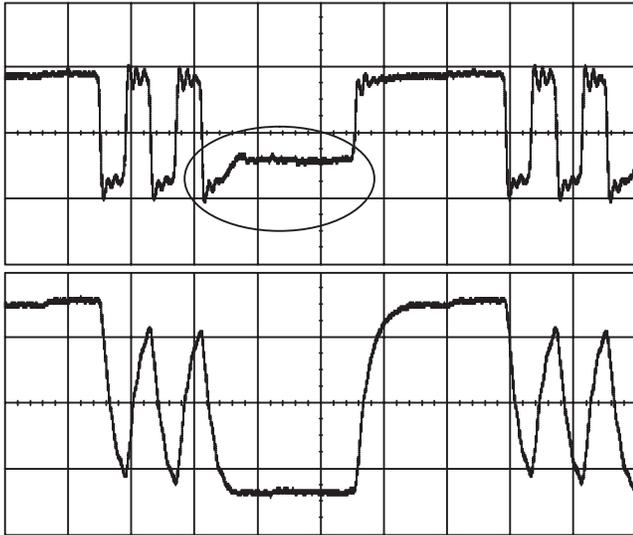
Ch. 1 = 400.0 mV/div Timebase = 10.0 ns/div  
Ch. 2 = 200.0 mV/div

**Figure 55. Single Driver, Direct-Coupled, Low-Frequency Bypass, 50Ω Resistive Load**



**Figure 56. Dual Transformer, Series Secondaries**

This allows their signal amplitudes to add. *Figure 57* shows the net result of this circuit. Note the LOW-level distortion present.

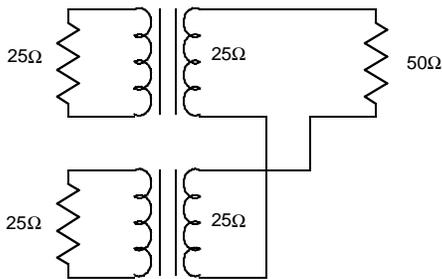


Ch. 1 = 400.0 mV/div  
Ch. 2 = 400.0 mV/div  
Timebase = 10.0 ns/div

**Figure 57. Dual Transformer, Series Secondaries, 50Ω Resistive Load**

The major changes that have occurred in the circuit are the amount of inductance present in the transformer(s) and the current run through them. With a single driver switching 800 mV into a 50Ω load, 16 mA of current are present. Doubling the output swing into the same 50Ω load by using both drivers, also doubles the current.

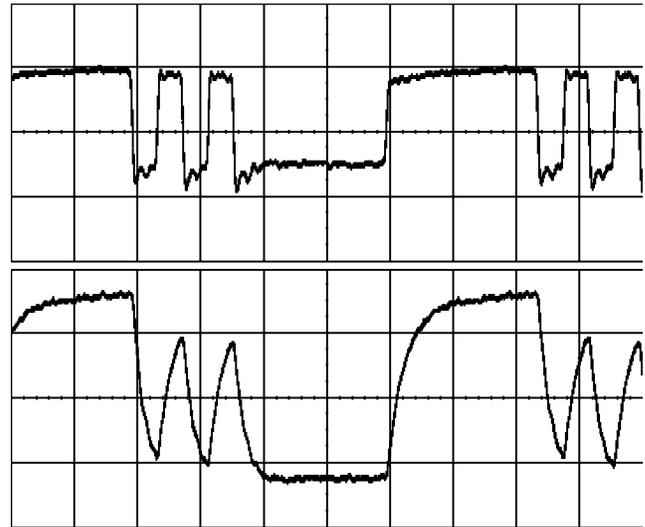
In these dual-driver configurations each driver must source twice as much current as a single driver configuration. The reason for this can be seen in *Figure 58*. With a 50Ω load on the secondary of a transformer, this same load is reflected on the primary. With dual transformers, half of the load is present on each transformer.



**Figure 58. Dual-Transformer Equivalent Loading**

To confirm this, the single driver circuit in *Figure 53* (with the larger 0.27-μF capacitor) was tested with a 25Ω resistive load. The results of that test are shown in *Figure 59*. This shows that the single driver configuration also generates the zero-level offset when presented with a low-impedance load, but at half the amplitude of a dual-transformer configuration.

This low side distortion is caused by the biasing network being sized for too large of a load impedance. An ECL driver



Ch. 1 = 400.0 mV/div  
Ch. 2 = 200.0 mV/div  
Timebase = 10.0 ns/div

**Figure 59. Single Driver, Direct Coupled, 25Ω Resistive Load**

sources current to set the HIGH or 1-level, while the bias network must sink sufficient current to set the LOW or 0-level.

The need to drive low-impedance loads places specific requirements on the current capability of the drivers. To differentially drive a 50Ω load (or transmission line) each driver must be capable of driving 25Ω single-ended loads. The line-bias networks must also be capable of sinking these large currents. This drive capability is beyond that of most ECL components, which are usually designed for only 50Ω loads. Only a few parts specifically identified as line drivers are made for operation with 25Ω loads.

The HOTlink transmitter PECL drivers are high-current line drivers and are designed specifically for driving 25Ω transmission lines. The use of standard ECL outputs designed for only 50Ω loads requires the addition of series current-limiting resistors in each primary leg of the transformer.

#### Long Cable Observations

When interfacing HOTLink to long cables:

- Higher cable impedances exhibit lower losses and less DDJ induced jitter.
- DC-block capacitors are not necessary but may be used to provide some precompensation to lower the destination jitter.
- Lower transformer inductance values provide less distortion and better high-frequency bandwidths.

#### Conclusions

The HOTLink family of data communications parts are designed to work optimally with either fiber-optic or copper-based interconnect. When interfaced to copper media, they may be interfaced to short, medium, and long-distance connections using only low cost passive components.

## References

1. Orr, William I., *Radio Handbook*, 23rd Edition, SAMS, 1992
2. Blood Jr., William R., *MECL System Design Handbook*, Fourth Edition, 1988
3. Trompeter, Ed, *Electronic Systems Wiring & Cable*, technical paper, Trompeter Electronics, Inc.
4. *The Radio Amateur's Handbook*, 50th Edition, ARRL, 1973
5. Hess, David & Goldie, John, *AN-916 A Practical Guide to Cable Selection*, National Semiconductor/Berk-Tek, 1994
6. Fowler, Bill, *Transmission Line Characteristics*, AN-108, National Semiconductor
7. *1990-91 Resistor/Capacitor Data Book*, Philips Components, 1990
8. *Fibre Channel Standard*, ANS X3.230-1994, American National Standards Institute, 1994
9. *Compliance Engineering Reference Guide*, Compliance Engineering, 1994

HOTLink is a trademark of Cypress Semiconductor.