



HOTLink™ Design Considerations

Application Note Overview

The HOTLink™ family of data communications products provides a simple and low-cost solution to high-speed data transmission. While these products are easy to use, the methods used to connect them to high-speed serial interfaces are often not intuitive. This document provides a basic level of explanation of the parallel and serial interface characteristics, and provides some cookbook solutions for interfacing them to different types of parts and media.

Primary Topics

The primary topics covered in this application note are:

- HOTLink Overview
- HOTLink Serial Signal Characteristics
- Terminating HOTLink Serial Signals
- Interfacing to HOTLink
- Serial Link Support Components

HOTLink Overview

HOTLink Features

- Fibre Channel compliant
- IBM® ESCON™ compliant
- DVB-ASI compliant
- ATM compatible
- 8B/10B-coded or 10-bit unencoded
- 150- to 400-Mbps data rate
- TTL-synchronous I/O
- No external PLL components
- Triple ECL 100K serial outputs
- Dual ECL 100K serial inputs
- Low power: 350 mW (Tx), 650 mW (Rx)
- Compatible with fiber-optic modules, coaxial cable, and twisted-pair media

- Built-In Self-Test (BIST)
- Single +5V supply
- 28-pin SOIC/PLCC/LCC
- 0.8μ BiCMOS

Functional Description

The CY7B923 HOTLink Transmitter and CY7B933 HOTLink Receiver are point-to-point communications building blocks that transfer data over high-speed serial links (fiber-optic, coax, and twisted/parallel-pair) at 150- to 400-Mbits/second. *Figure 1* illustrates typical connections to host systems or controllers.

Eight bits of user data or protocol information are loaded into the HOTLink Transmitter and are encoded into a 10-bit transmission character. The transmission character is shifted serially out of the three differential positive ECL (PECL) serial ports at the bit-rate (which is ten times the character-rate).

The HOTLink Receiver accepts the serial bit stream at its differential line receiver inputs and, using a completely integrated phase-locked-loop (PLL) clock synchronizer, recovers the timing information necessary for data reconstruction. The bit stream is deserialized, decoded, and checked for transmission errors. The recovered character is presented in parallel to the receiving host along with the synchronized character-rate clock.

The 8B/10B encoder/decoder (Reference 1, 2) can be disabled in systems that already encode or scramble the transmitted data. Signals are available to create a seamless interface with both asynchronous FIFOs (i.e., Cypress's CY7C42X) and clocked/synchronous FIFOs (i.e., Cypress's CY7C44X/46X). A BIST pattern generator and checker allows testing of the transmitter, receiver, and the connecting link as a part of a system diagnostic check.

HOTLink devices are ideal for a variety of applications where a parallel interface can be replaced with a high-speed point-to-point serial link. Applications include interconnecting workstations, servers, mass storage, and video transmission equipment.




Figure 1. HOTLink System Connections

CY7B923 HOTLink Transmitter Description

The function of the HOTLink Transmitter is to convert character-rate parallel data into a high-speed serial data stream. A logic block diagram of the transmitter is shown in *Figure 2*.

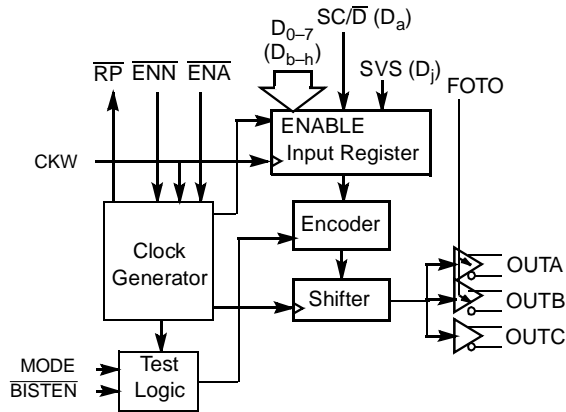


Figure 2. CY7B923 Transmitter Logic Block Diagram

Input Register

The Input Register holds the data to be processed by the HOTLink Transmitter and allows the input timing to be made consistent with standard FIFOs. The Input Register is clocked by CKW (clock write) and loaded with information on the D_{0-7} , SC/\overline{D} (special character/data select), and SVS (send violation symbol) pins. Two enable inputs (ENNA and ENN) allow the user to choose when data is to be sent. Asserting ENNA (enable, active LOW) causes the inputs to be loaded on the rising edge of CKW. If ENN (enable next, active LOW) is asserted when CKW rises, the data present on the inputs at the *next* rising edge of CKW is loaded into the input register. These two enable inputs allow proper timing and function for compatibility with either asynchronous FIFOs or clocked/synchronous FIFOs without external logic.

In BIST mode, the Input Register becomes a pattern generator by logically converting into a linear-feedback shift-register (LFSR). When enabled, this LFSR generates a 511-character sequence that includes all Data and Special Character codes, including the explicit violation symbols. This pattern provides a predictable but pseudo-random sequence that can be matched to an identical LFSR in the HOTLink Receiver. For additional information see the Cypress Semiconductor application note "HOTLink Built-In Self-Test."

Encoder

The Encoder transforms the input data, held by the Input Register, into a form more suitable for transmission on a serial interface link. The code used is specified by the ANSI X3.230-1994 Fibre Channel standard (Reference 3) and the IBM ESCON channel (Reference 4) (code tables are available in the CY7B923/CY7B933 data sheet). The eight D_{0-7} data inputs are converted to a 10-bit Transmission Character, as directed by the state of the SC/\overline{D} input. If SC/\overline{D} is HIGH, the data inputs represent a control code and are encoded using the Special Character code tables. If SC/\overline{D} is LOW, the data inputs are converted using the Data code table. If a character-time passes with the inputs disabled, the Encoder outputs a Special Character Comma (K28.5 or SYNC) to maintain link

synchronization. The SVS input forces the transmission of a specified Violation symbol to allow the user to check error handling logic in a system controller.

The 8B/10B coding function of the Encoder can be bypassed for systems that include an external encoder or scrambler as part of the controller. This bypass capability is controlled by setting the MODE select pin HIGH. When in Bypass mode, D_{a-i} (note that bit order is specified by the Fibre Channel 8B/10B code) become the ten inputs to the Shifter, with D_a being the first bit to be shifted out.

Shifter

The Shifter accepts parallel data from the Encoder once each character-time and shifts it to the serial interface output buffers using a PLL multiplied bit-clock that runs at 10 times the character-clock (CKW) rate. Timing for the parallel transfer is controlled by the counter included in the Clock Generator, and is not affected by signal levels or timing at the input pins.

OutA, OutB, OutC

The serial interface ECL output buffers (100K signal levels referenced to +5V) are the drivers for the serial media. They are all connected to the Shifter and present the same serial data. Two of the output pairs ($OUTA_{\pm}$ and $OUTB_{\pm}$) are controlled by the FOTO input and can be disabled by the system controller to force a logical zero (i.e., "light off") at the outputs. The third output pair ($OUTC_{\pm}$) is not affected by FOTO and supplies a continuous data stream suitable for loop-back testing of the subsystem.

$OUTA_{\pm}$ and $OUTB_{\pm}$ respond to FOTO input changes within a few bit-times. However, since FOTO is not synchronized with the transmitter data stream, the outputs are forced off or turned on at arbitrary points in a transmitted character. This function is intended to augment an external laser safety controller and as an aid for Receiver PLL testing.

In wire-based systems, control of the outputs may not be required, and FOTO can be strapped LOW. The three output pairs are intended to add system and architectural flexibility by offering identical serial bit-streams with separate interfaces for redundant connections or for multiple destinations. Unneeded outputs can be left open or wired to V_{CC} to disable and power down the unused output circuitry.

Clock Generator

The clock generator is an embedded phase-locked loop (PLL) that takes a character-rate reference clock (CKW) and multiplies it by ten to create a bit-rate clock for driving the serial shifter. The character-rate reference comes from CKW, the rising edge of which clocks data into the Input Register. This clock must be a crystal-referenced (or other high-stability) pulse stream that has a frequency between the minimum and maximum specified for the HOTLink Transmitter/Receiver pair. Signals controlled by this block form the bit-clock and the timing signals that control internal data transfers between the Input Register and the Shifter.

The read pulse (\overline{RP}) is derived from the feedback counter used in the PLL multiplier. It is a character-rate pulse stream with the proper phase and pulse widths to allow transfer of data from an asynchronous FIFO. Pulse width is independent of CKW duty cycle, since proper phase and duty cycle is maintained by the PLL. The \overline{RP} pulse stream ensures correct data transfers between asynchronous FIFOs and the transmitter input latch with no external logic.

Test Logic

Test Logic includes the initialization and control for the built-in self-test (BIST) generator, the multiplexer for Test mode clock distribution, and control logic to properly select the data encoding. Test Logic is discussed in more detail in the CY7B923/CY7B933 HOTLink data sheet.

CY7B933 HOTLink Receiver Description

The function of the HOTLink Receiver is to convert a high-speed serial data stream into character-rate parallel data. A logic block diagram of the receiver is shown in Figure 3.

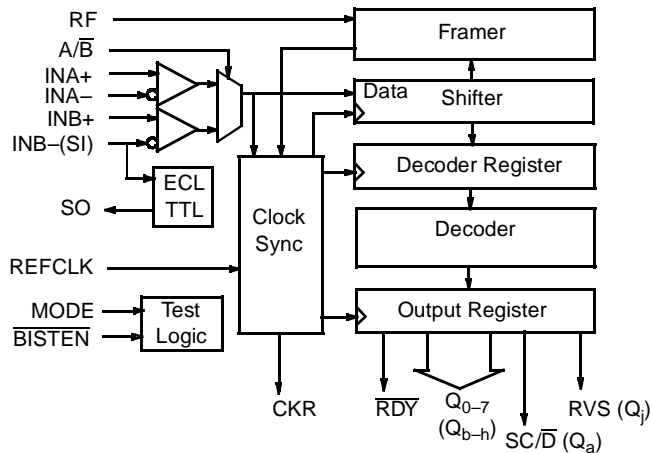


Figure 3. CY7B933 Receiver Logic Block Diagram

Serial Data Inputs

The HOTLink Receiver has two differential line receivers (INA_{\pm} and INB_{\pm}) that can be selected as inputs for the serial data stream. INA_{\pm} or INB_{\pm} is selected with the A/\bar{B} input. INA_{\pm} is selected when A/\bar{B} is HIGH and INB_{\pm} is selected when A/\bar{B} is LOW. The threshold of A/\bar{B} is compatible with ECL 100K signals. TTL logic elements can be used to select the INA_{\pm} or INB_{\pm} inputs by adding a resistor voltage divider to a TTL driver connected to A/\bar{B} (see Figure 35). The differential sensitivity of INA_{\pm} and INB_{\pm} will accommodate wire interconnect with filtering losses or transmission line attenuation greater than 20 dB ($V_{DIF} \geq 50$ mV). These inputs can alternatively be directly connected to fiber-optic interface modules (any ECL logic family, not limited to ECL 100K) with up to 1.2V of differential signal. The common-mode tolerance accommodates a wide range of signal termination voltages. The highest HIGH input that can be tolerated is $V_{IN} = V_{CC}$, and the lowest LOW input that can be interpreted correctly is $V_{IN} = GND + 2.0V$.

ECL/TTL Translator

The function of the $INB(INB+)$ input and the $SI(INB-)$ input is determined by the connection on the SO output pin. If the ECL/TTL translator function is not required, the SO output is wired to V_{CC} . A sensor circuit detects this connection and causes the inputs to become INB_{\pm} (a differential line-receiver for serial-data input). If the ECL/TTL translator function is required, the SO output is connected to a normal TTL load (typically one or more TTL inputs, but no pull-up resistor) and the inputs become INB (single-ended ECL 100K-level serial-data input) and SI (single-ended ECL 100K-level status input).

This positive-referenced ECL-to-TTL translator is provided to eliminate external logic between an ECL carrier-detect or link status signal and a TTL input in the control logic. The input threshold is compatible with ECL 100K levels (+5V referenced).

Clock Sync

The Clock Synchronizer function is performed by an embedded phase-locked loop (PLL) that tracks the frequency of the incoming serial bit-stream and aligns the phase of its internal bit-rate clock to the serial data transitions. This block contains the logic to transfer the data from the Shifter to the Decode register once every character. The counter that controls this transfer is initialized by the Framer logic. CKR is a buffered output derived from the bit counter used to control Decode register and Output register transfers.

Clock output logic is designed such that when reframing causes the counter sequence to be interrupted, the period and pulse width of CKR is never less than normal. Reframing may stretch the period of CKR by up to 90%, and either CKR pulse width HIGH or pulse width LOW may be stretched, depending on when reframe occurs.

The REFCLK input provides a character-rate reference frequency to improve PLL acquisition time and limit unlocked frequency excursions of CKR when no data is present at the serial inputs. The frequency of REFCLK is required to be within $\pm 0.1\%$ of the frequency of the clock that drives the transmitter CKW pin.

Framer

Framer logic checks the incoming bit-stream for the pattern that determines the character boundaries. This combinatorial logic filter looks for the ANSI Fibre Channel symbol defined as a Special Character Comma (K28.5) (Reference 3). When framing is enabled and this character is found, the free-running bit-counter in the Clock Sync block is synchronously reset to its initial state, thus framing the data on the correct character boundaries.

Random errors that occur in the serial data can corrupt some data patterns into a bit-pattern identical to a K28.5, and thus cause an erroneous data-framing error. The RF input prevents this by inhibiting reframing during times when normal message data is present. When RF is held LOW, the HOTLink Receiver deserializes the incoming data without trying to reframe the data. When RF rises, \overline{RDY} is inhibited until a K28.5 is detected, after which \overline{RDY} resumes its normal function. While RF is HIGH, it is possible that an error in the serial data could cause misframing, after which all data would be corrupted. Likewise, a K28.7 followed by D11.x, D20.x, or an SVS (C0.7) followed by D11.x will cause erroneous framing. These sequences must be avoided while RF is HIGH.

If RF remains HIGH for greater than approximately 2048 characters, the framer switches to multi-byte framing, requiring two K28.5 Special Characters within a five character span.

Shifter

The Shifter accepts serial data from one of the serial data input receivers, one bit at a time, as clocked by the Clock Sync logic. Data is examined by the Framer on each bit, and is transferred to the Decode Register once per character.

Decode Register

The Decode Register accepts data from the Shifter once per character as determined by the logic in the Clock Sync block. It is presented to the Decoder and held until it is transferred to the Output Register.

Decoder

Parallel data is transformed from ANSI Fibre Channel 8B/10B codes (Reference 3) back to “raw data” in the Decoder. This block uses the standard decoder patterns found in the Valid Data Characters and Valid Special Character Codes and Sequences (code tables are available in the CY7B923/CY7B933 data sheet). Data Characters are signaled by a LOW on the SC/ \overline{D} output and Special Characters are signaled by a HIGH on the SC/ \overline{D} output. Unused patterns or disparity errors are signaled as errors by a HIGH on the RVS (Received Violation Symbol) output and by specific Special Character codes.

Output Register

The Output Register holds the recovered data (Q_{0-7} , SC/ \overline{D} , and RVS) and aligns it with the recovered character clock (CKR). This synchronization ensures proper timing to match a FIFO interface or other logic that requires glitch-free and specified output behavior. Outputs change synchronously with the rising edge of CKR.

In BIST mode, this register becomes a pattern generator and checker by logically converting itself into a linear-feedback shift-register (LFSR). When enabled, this LFSR generates a 511-character sequence that includes all Data and Special Character codes, including the explicit violation symbols. This provides a predictable but pseudo-random sequence that can be matched to an identical LFSR in the transmitter. When synchronized, it checks each character in the Decoder with each character generated by the LFSR and indicates errors using RVS. Patterns generated by the LFSR are compared after being buffered to the output pins and then fed back to the comparators, allowing a test of the entire receive function.

In BIST mode, the LFSR is initialized by the first occurrence of the transmitter BIST loop start code D0.0 (D0.0 is sent only once per BIST loop). Once the BIST loop has been started, RVS goes HIGH for pattern mismatches between the received sequence and the internally generated sequence. Code rule violations or running disparity errors that occur as part of the BIST loop do not cause an error indication. RDY pulses high once per BIST loop and can be used to check test pattern progress. The receiver BIST checker can be reinitialized by leaving and re-entering BIST mode.

Test Logic

Test Logic includes the initialization and control for the built-in self-test (BIST) checker, the multiplexer for Test mode clock distribution, and control logic for the decoder. Test Logic is discussed in more detail in the CY7B923/CY7B933 HOTLink data sheet.

HOTLink Serial Signal Characteristics

The serial interfaces on the HOTLink Transmitter and Receiver are based on the standard for high-speed digital logic called emitter-coupled-logic or ECL. This form of logic has been used commercially in integrated circuits since the early 1960s, and prior to that it was implemented in discrete form.

ECL is a non-saturating form of digital logic. ECL gets its name from how the emitters of a differential amplifier in the circuit are connected. The main features of this logic family are very high speed, low noise, and the ability to drive low-impedance transmission lines.

In the past, many engineers have avoided ECL as a logic family because it was different from the TTL-compatible families with which they were more familiar. Proper use of ECL requires the understanding and application of transmission lines, line termination, and power supply bypassing. Because of the faster speeds present in the newer TTL-compatible families, these same disciplines are now required for TTL circuits.

ECL Signal Level Reference

The primary differences between ECL and other logic families are the signal levels used to represent the HIGH and LOW logic levels.

In the TTL and CMOS logic families, a LOW is usually some level close to V_{SS} , and a HIGH is usually some level close to V_{CC} . The ground or reference point for these measurements is usually the V_{SS} point, with V_{CC} set to +5V from that ground reference.

In standard ECL this changes significantly. Instead of having the ground reference at V_{SS} , it is placed at V_{CC} . This means that both HIGH and LOW logic levels exist at potentials that are negative with respect to ground. Standard ECL is specified as operating with a negative supply (–4.5V to –5.2V for V_{EE}). Since ground is only a reference point, it is also possible to operate ECL with a positive supply. When used in this mode ECL is usually referred to as PECL which means Positive ECL.

ECL Basic Switch

Internally, ECL gates (or switches) operate using a current source whose current is directed through one of two paths to V_{CC} . A schematic of this basic ECL switch is shown in Figure 4 (Reference 5).




Figure 4. Basic ECL Switch

In this ECL switch, the state of the switch is determined by the voltage drop across R1 and R2. The output signal swing is set by the size of these resistors and the magnitude of the current passed through them.

The base of Q2 is biased at a fixed voltage called V_{BB} . This voltage determines at what level of V_{IN} on Q1 that the majority of the current flowing in the switch changes from R1 to R2. If V_{IN} is set to the same voltage as V_{BB} , the current divides

equally between R1 and R2. Increasing V_{IN} by 125 mV above V_{BB} causes essentially all the current to flow through Q1 (and hence R1). Lowering V_{IN} to 125 mV below V_{BB} causes essentially all the current to flow through Q2. This means that an input swing of as little as 250 mV can cause the ECL gate to switch completely from a 0 to a 1. To provide noise immunity and allow operation over a wide variety of conditions, the actual signal swing specified for ECL signals is around 800 mV.

Emitter-Follower

The switch shown in Figure 4 can react very quickly but, because of its high-value resistor pull-ups (R1 and R2), its switching delay varies directly with load capacitance. To allow larger loads to be driven, and to make the output voltages compatible with the input of subsequent gates, additional transistors are added in an emitter-follower configuration as shown in Figure 5.

These emitter-follower transistors have a very low on-impedance (5–7 Ω). This allows ECL gates to drive transmission lines having impedances at or below 50 Ω , and can supply load currents of up to 50 mA.

The emitter-follower transistors have an uncommitted emitter as their output. This allows the transistor to source, but not sink, current. This is effectively the opposite of an open-collector output in a TTL part. To allow the output to function correctly, it requires a load that operates as a pull-down.




Figure 5. Buffered ECL Switch

ECL Signal Levels

ECL signals operate over a very narrow and tightly controlled range. These signal levels are referenced from the V_{CC} pins of the parts. Figure 6 shows the relationships of the different output and input levels for ECL gates. The names of these levels are detailed in Table 1.

ECL Output Signal Levels

ECL outputs are all referenced from V_{CC} . A typical ECL driver has an output-HIGH level (V_{OH}) of $V_{CC} - 0.85V$ and an output-LOW level (V_{OL}) of $V_{CC} - 1.7V$. These typical values are seldom specified for parts because a good design must be done using the range limits for these signals as listed in Table 1. Actual values for these levels vary by individual part type and ECL family.




Figure 6. ECL Signal Levels

Table 1. ECL Signal Level Names

Name	Description
V_{OHH}	Highest Output HIGH Voltage
V_{OHL}	Lowest Output HIGH Voltage
V_{OLH}	Highest Output LOW Voltage
V_{OLL}	Lowest Output LOW Voltage
V_{IH}	Lowest Input HIGH Voltage Threshold
V_{IL}	Highest Input LOW Voltage Threshold
V_{NH}	High Input Noise Margin ($V_{OHL} - V_{IH}$)
V_{NL}	Low Input Noise Margin ($V_{OLH} - V_{IL}$)

ECL Input Signal Levels

ECL Inputs are also referenced from V_{CC} . A typical ECL receiver has an input-HIGH (V_{IH}) threshold of $V_{CC} - 1.1V$ and an input-LOW (V_{IL}) threshold of $V_{CC} - 1.47V$. These differences between the output and input HIGH and LOW values translate directly into the usable noise margin (V_{NH} and V_{NL}) of a system.

Viewing ECL Signals

Proper viewing of ECL signals requires use of an oscilloscope and probes with sufficient bandwidth to see the important features of the waveforms. Depending on the speed of the signals being viewed, different scope and probe characteristics are required.

Oscilloscope Bandwidth

Oscilloscope bandwidth is not a simple number; it is based on the combined bandwidths of multiple pieces of the measurement system. These can include the oscilloscope, the scope probe amplifier, the probe itself, and possibly other components.

The calculation for bandwidth is based on an inverse sum-of-squares as shown in Equation 1.

$$bw = \frac{1}{\sqrt{\left(\frac{1}{bw_1}\right)^2 + \left(\frac{1}{bw_2}\right)^2}} \quad \text{Eq. 1}$$

Thus a scope with a 1-GHz bandwidth probe using a 1-GHz bandwidth amplifier would only have a usable bandwidth of 700 MHz.

The current ANSI Fibre Channel standard specifies the minimum system bandwidth for testing as 1.8 times the signaling rate. For testing with the HOTLink parts (400 MBaud), this translates to a minimum system bandwidth of 720 MHz. This is translated into a viewable rise time using Equation 2 (Reference 6).

$$t_r = \frac{0.35}{bw} \quad \text{Eq. 2}$$

This means that the oscilloscope and probes, having a 700 MHz bandwidth, can display signals with rise-times no faster than 480 ps, without having more than 3 dB of attenuation.

Note: Various scope manufacturers use different conventions to specify bandwidth for their equipment; i.e., specified bandwidth is not necessarily where the displayed waveforms are 3 dB down in amplitude.

Scope Probes

Scope probes are available with many different characteristics. The three main types are referred to as passive high-impedance, active high-impedance, and passive low-impedance.

Passive high-impedance probes usually range from as low as 10-k Ω to 10-M Ω load impedance. This number identifies the loading effect of the probe when attached to a circuit. The best feature of high-impedance probes is that their impedance is usually much larger than those of the circuit under test and thus do not present any appreciable DC-load to the measured signal when present.

Passive high-impedance probes do suffer one major drawback: significant capacitive loading. Most high-impedance probes present from 5 pF to 20 pF of capacitance at the probe tip. This capacitance affects measurements in two ways; it slows down the circuit being measured, and it degrades the rise-time of the probe. The upper bandwidth limit for passive high-impedance probes is around 400 MHz.

Active high-impedance probes combine a high bandwidth amplifier with the probe to improve the overall bandwidth of the system. These probes usually exhibit load impedances of 10 k Ω to 10 M Ω but have load capacitances of less than 3 pF. This type of probe has a typical upper bandwidth limit of around 1 GHz.

Care should be taken when using active probes as the manufacturers specified bandwidth may not be where the signal measured is 3 dB down. To achieve higher bandwidths, some active probes have non-linear responses to equalize the probe response. When presented with edge rates or frequency components beyond the specified probe bandwidth, the probe and scope may actually display a distorted waveform having more high-frequency components present than are actually in the measured signal.

Passive low-impedance (resistive divider or transmission line) probes are used for the highest frequency work. These probes are available in load impedances from 50 Ω to 5 k Ω , and present load capacitances of 1 pF or less. A typical upper bandwidth limit for these probes is around 3 GHz. Unlike the high-impedance probes, low-impedance probes are de-

signed to connect to a 50 Ω transmission line system and do not require compensation. The probe itself is an extension of the 50 Ω transmission line present in the scope, and contains a precision resistive-divider at the probe tip.

The main drawback of passive low-impedance probes is the load impedance they present to the circuit. The rule of thumb for probes is that the probe impedance needs to be an order of magnitude greater than the impedances present around it to avoid any appreciable distortion. To get around this the probe is often designed as part of the system under test, such that its impedance is factored into the design. When the probe is not present it may be necessary to change component values or configurations to compensate for the absence of the probe (Reference 7).

Table 2 shows a summary of typical oscilloscope probe characteristics. For proper viewing of HOTLink ECL signals, an oscilloscope (and probes) should have a minimum system bandwidth of 720 MHz. In most cases this requires use of low-impedance probes.

Table 2. Typical Probe Characteristics

Probe Type	Z	C _{load}	BW (MHz)
Passive High-Z	10 k– 10 M Ω	5–20 pF	400
Active High-Z	10 k– 10 M Ω	3 pF	1000
Passive Low-Z	50–5 k Ω	1 pF	3000

Probe Grounding

As with any measurement, a good ground is mandatory. What is often misunderstood is just what *is* a good ground. At the frequencies used with HOTLink, a long looping ground lead is about as good as no ground at all. Three factors come into play that degrade the signal: the reflections caused by the scope probe, and the ground inductance and parasitic capacitance limiting the probe's bandwidth. A simple rule of thumb for ground leads is that they exhibit about 1 nH of inductance for each millimeter of length. As the length of the probe's ground lead increases, the probe's resonance point decreases.

To view a signal with minimal distortion, the probe's resonant frequency must remain above the highest frequency signal component of interest. The graph in Figure 7 shows how a scope probe's resonant frequency varies for different lengths of ground loop inductance and tip capacitance. This graph is based on Equation 3 with the diagram of a low-impedance probe shown in Figure 8.

$$\omega = 2\pi f = \frac{1}{\sqrt{LC}} \quad \text{Eq. 3}$$

From this graph it is quite apparent that a ground lead of only 10 mm cuts the resonant frequency of the probe by 75%. For signal viewing at HOTLink serial data rates it is usually necessary to use coaxial scope-tip sockets soldered directly to a circuit board, or some other probe type that probes for signal and ground without a loose ground lead (Reference 8).




Figure 7. Scope Probe Resonant Frequency




Figure 8. Scope Probe Tip Schematic

Probing From V_{CC}

The normal mode for probing ECL is to use V_{CC} as the ground reference. In this mode the signal being viewed is below ground and is relatively close to the ground reference. If the overall circuit design uses TTL parts in a mix with the negative referenced ECL, the TTL signals all exist above ground. If the ECL parts are operated in a PECL mode where they share a common V_{CC} supply with other TTL or CMOS parts, all probing should be done from TTL ground, which is the V_{EE} side of the ECL parts.

Probing From V_{EE}

When V_{EE} is used as the scope ground, other issues may come into play. In this mode the ECL signal is now positioned almost 4V above the reference level. While many scopes are able to perform a DC offset to make the ECL signal viewable, some do this at the expense of sensitivity. In other words, a signal that is viewable at 100 mV/div when offset less than 2V, may only be viewable at 500 mV/div when offset by 4V. Since the total signal swing for ECL signals is only 800 mV, it may be difficult to see a detailed representation of the waveform at this resolution.

Another problem with measuring from V_{EE} is that all the references in the ECL part are regulated from V_{CC} , not V_{EE} . This means that any amplitude changes or ripple in the power supply are now added into the displayed waveform.

One way around the offset problem is to AC couple the signal into the scope. Some scopes offer this as a front panel set-up

selection, while others require the addition of a wide-bandwidth DC-blocking capacitor in-line with the scope probe. Either of these remove all DC components from the signal under test, and allow the signal to be displayed at the maximum resolution of the oscilloscope.

Wide-bandwidth capacitors designed for this function are available from most test equipment manufacturers for use with existing probes and scope amplifiers. Some common capacitor types for SMA connector probes are the Tektronix 015-1-13-00 and Hewlett-Packard 11742A. For BNC connected probes the Hewlett-Packard 10240B is also available.

Sample ECL Waveforms

ECL signals, when properly biased, terminated, and bypassed, are very clean and stable. Any noticeable overshoot on signals is usually caused by reflections from improperly terminated transmission lines or improper probing. *Figure 9* shows what a pristine single-ended ECL waveform should resemble when viewed on a scope.




Figure 9. Good ECL Waveform, Single-Ended vs. V_{CC} Ground

Both the rising and falling edges are quite symmetrical and approximate an RC charge/discharge curve. The peak-to-peak range of the transition covers approximately 800 mV and is centered around $V_{CC} - 1.3V$. This signal was measured using a 500Ω, 1.5-GHz bandwidth low-impedance probe, on a scope having 1-GHz bandwidth, with V_{CC} as the probe ground. The probe load impedance (500Ω) was combined with other bias resistors to present a 50Ω to $V_{CC} - 2V$ load on the signal.

With incorrect termination, a waveform such as that in *Figure 10* can result. Here the spike in the middle of a low area may cross the receiver V_{IH} threshold and cause the receiver to start to switch.

ECL Logic Families

Just as the TTL compatible world has its 7400, 74LS, 74H, 74S, 74AS, 74ALS, etc. logic families that have evolved over time, so does ECL. The most common families still in use are referred to as 10K (e.g., SL10104), 10KH (e.g., MC10H116), and 100K (e.g., F100150). These ECL families differ in terms




Figure 10. Bad ECL Waveform

of speed, signal levels, noise margins, and temperature and voltage stability.

10K ECL

The 10K ECL family has been around since 1971. It provides propagation delays of 2 ns with slow 3.5-ns edge rates (10%–90%). The voltage swings and switching thresholds of this logic family are relatively insensitive to variations in the power supply voltage but are affected by operating temperature (–30°C to +85°C). The V_{BB} bias network is fixed at $V_{CC} - 1.29V$, and is compensated for voltage and temperature. In the basic 10K ECL switch the current source is unregulated and consists of a single resistor between V_{EE} and the tied emitters of the differential amplifier. The transfer curves of a simple 10K gate are shown in Figure 11 and detail how this family is sensitive to temperature variations in both inputs and outputs (Reference 19).

10KH ECL

To improve system speeds, the 10KH ECL family was introduced in 1981. It reduced propagation delays to 1 ns while edge rates were set to 1.8 ns. Because the thresholds and voltage swings remain the same in 10KH as in 10K, these two ECL families are fully compatible with each other. The temperature- and voltage-compensated V_{BB} reference network from 10K parts was replaced with a fully compensated and regulated supply. To improve the V_{OL} levels, the resistor current source was replaced with a regulated current source. This allowed the collector resistors in the ECL switch to be matched and have similar switching characteristics. The transfer curves of a simple 10KH gate (see Figure 12) illustrate how this family improves noise margins over 10K ECL, yet remains sensitive to temperature variations. The 10KH family also is specified to operate over a narrower temperature range (0°C to 75°C) than 10K ECL (Reference 19).

100K ECL

The 100K ECL family is a faster and easier to use ECL logic family. Introduced in 1973, this family improved on the internal structures to provide 750-ps propagation delays and 1-ns edge rates. In addition to speed improvements, the 100K ECL family was the first to introduce full compensation. This




Figure 11. 10K ECL Transfer Functions




Figure 12. 10KH ECL Transfer Functions

means that all the critical structures in the parts are now compensated for variations in voltage and temperature. This minimizes differences in propagation delays from one stage to the next that limit the maximum operating rate of a system. This stability is illustrated in the transfer curves in Figure 13 (Reference 5).

In the 100K ECL family the operating temperature range is expanded to 0°C to 85°C but the nominal operating voltage changes from –5.2V to –4.5V.




Figure 13. 100K ECL Transfer Functions

HOTLink ECL Outputs

All ECL outputs of the HOTLink Transmitter are ECL 100K-level compatible. This means that these outputs meet or exceed all voltage, current, and edge rates specifications of 100K ECL and will interoperate with other 100K ECL parts. This signal level compatibility is required by the ANSI Fibre Channel standard (Reference 3).

The HOTLink ECL outputs actually are substantially better than the 100K ECL specification, allowing operation with $5V \pm 10\%$ supplies over the full -55°C to $+125^{\circ}\text{C}$ temperature range. This allows the HOTLink parts to be used in TTL, PECL, or ECL environments.

The HOTLink Transmitter has six ECL outputs configured as three differential pairs: OUTA_{\pm} , OUTB_{\pm} , and OUTC_{\pm} (see Figure 2). These differential outputs may be used to communicate with ECL-compatible receivers in either single-ended (strongly discouraged) or differential (preferred) modes.

HOTLink Transmitter Single-Ended Connections

A single-ended connection is used most often for logic functions. In this type of a connection, a single output of a driver is attached to a single input of a receiver. The receiving element is thus dependent on the driver and interconnect for maintaining the input signal in the narrow voltage bands specified for a valid logic 1 or 0.

Figure 14 illustrates the basic components of a single-ended connection. The driver differential pair outputs are biased to allow them to switch. The receiver, as with all ECL gates, is based on a differential amplifier. In the case of a single-ended receiver, the second input into the differential amplifier is not present at an external pin on the chip, but is instead connected internally to a V_{BB} reference voltage. The receiver switches as the signal present on IN+ goes either above or below the internal threshold set by V_{BB} .

While connections of this type are perfectly fine for logic functions, they should be avoided for a communications link. In a single-ended environment, any signal level differences

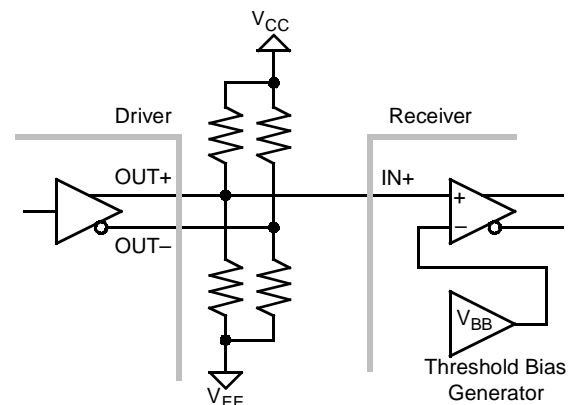


Figure 14. Single-Ended Connection

(caused by temperature, logic family, transients, power supply noise, etc.) directly affect the received signal timing. In a logic function this timing variation limits a design both in determining how fast the system may operate, and in how much noise margin is present.

In a communications link these variations in timing translate directly into jitter in the serial data stream. Jitter affects a serial link by limiting not only how fast the link can operate (data rate) but also how far the data can be sent. Jitter is discussed in detail later in this document.

The only expected single-ended connection on a HOTLink Transmitter is for a local loopback function to a HOTLink Receiver (when the INB- input is not available for a differential connection because it has been used as an ECL-to-TTL translator). In this connection it is expected that the transmitter and receiver are in relatively close proximity, such that the connection between them is more on the order of a logic connection than a communications link. The small amount of jitter caused by the single-ended connection will be far below the jitter susceptibility of the HOTLink Receiver.

HOTLink Transmitter Differential Connections

A differential connection is the preferred attachment for HOTLink Transmitter serial outputs. In a differential connection both outputs of a driver are connected to the true and complement inputs of an ECL-compatible receiver. When connected in this fashion the majority of the interconnect dependencies are removed. The main advantages of a differential connection are insensitivity to the logic family, operating temperature, and power supply variations. In addition, the connection is now immune to most common-mode noise.

Figure 15 shows the basic components of a differential connection. The driver differential pair outputs are biased to allow them to switch. Now both true and complement inputs of the receiver differential amplifier are available at external pins and are connected to the complementary outputs of the driver.

Some ECL differential receivers may also provide an external V_{BB} reference. This reference is provided for those cases where a driver is connected single-ended to one of the differential receiver inputs. The other receiver input must then be connected to the V_{BB} reference to allow the receiver to switch. With a true differential connection this V_{BB} output should remain open.

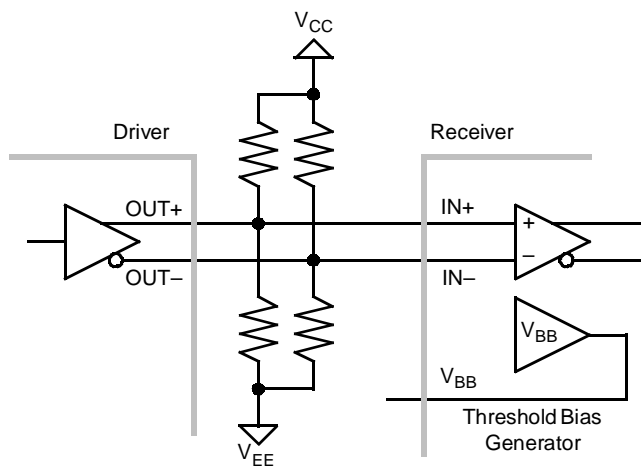


Figure 15. Differential Connection

The main concerns in a differential connection are signal skew and crosstalk. Skew is the difference in arrival time of the OUT+ and OUT- signals at the receiver. Crosstalk is the coupling of energy into these same two signals.

As the amount of signal skew present in a differential connection is increased, the effective signal rise and fall times at the differential receiver also increase. In systems with large amounts of signal skew, it is possible for short pulses to never be detected by the receiver.

The main cause of signal skew is asymmetric routing of the true and complement signals between the driver and the receiver. A 1-inch difference in routing length is equal to about 150 ps of signal skew. This problem is corrected by maintaining matched signal runs between the HOTLink Transmitter and the ECL differential receiver.

The main cause of crosstalk is long parallel traces carrying other signals. The adjacent lines act as coupling transformers and transfer energy from one to another. One cure for this is to limit the length of the connection by placing the ECL differential receiver as close to the HOTLink Transmitter as possible. Other possibilities are to route the differential signals between the power planes as stripline. If routing is to remain on the same plane, the crosstalk effects can be minimized by horizontally separating the adjacent signals as far as possible or by routing a ground trace (with many vias to attach the ground trace to the ground plane) between the signals.

HOTLink ECL Inputs

The ECL inputs on the HOTLink Receiver are also ECL 100K-level compatible. Similar to the transmitter, these inputs have also been enhanced to operate over a wider range than standard 100K ECL.

The differential INA± and INB± inputs offer improved minimum sensitivity of 50 mV, compared to 150 mV for the few 100K ECL differential receivers available. These inputs may be connected directly to either power rail without damage to the part, or changing the internal thresholds of other sections of the receiver. These same differential inputs also operate with a 3V common-mode rejection range (V_{CC} down to $V_{CC} - 3V$) that is twice the 1.5V range of standard 100K ECL differential receivers ($V_{CC} - 0.5V$ down to $V_{CC} - 2V$).

Note: While differential outputs are quite common on ECL parts, true differential inputs are rare. The most common usage for differential inputs is on line receivers and clock drivers. The common-mode range on some parts with differential inputs is quite limited and should not be expected to operate over even a narrow range unless explicitly stated in the manufacturer's data sheet.

The INA± inputs of the HOTLink Receiver should always be connected to a differential signal source. Since there is no V_{BB} reference output on the receiver, there is no way to properly bias the second input of the differential receiver.

The INB± inputs may be configured to operate either as a differential receiver (in which case it should be connected to a differential signal source) or as two single-ended receivers. When operated as two single-ended receivers (as configured using the SO pin) the INB+ input operates as a 100K ECL single-ended receiver for serial data, while the INB-(SI) input operates as a 100K ECL single-ended receiver for an ECL-to-TTL level translator. The V_{BB} reference for these signals is available only inside the HOTLink Receiver and is not brought to an external pin. Signals connected to these single-ended inputs must now ensure operation within the 100K threshold levels.

Mixing ECL Logic Families

It is often desirable to use ECL parts of different families together in the same design. This can be done if certain rules are followed. The main reasons for these rules are the variability in signaling levels in ECL 10K family parts. Figure 16 shows a DC-level comparison for 100K ECL outputs driving single-ended 10K ECL inputs.

In this configuration there is only 20 mV of margin between the 100K V_{OHL} and the 10K V_{IH} at the upper end of the temperature range. With 10K parts driving other 10K parts (assuming a common operating temperature) this is not a problem as the internal reference in each part follows a similar temperature shift. If the case temperature of the receiving 10K part can be kept below 35°C (100 mV margin), it can safely be used with 100K ECL parts for logic functions.

While the V_{OLH} specification appears to also have a noise margin problem, it does not. What occurs here is a condition where the receiver may be operated outside its linear region; i.e., 1s and 0s are detected properly but the timing response may not match the manufacturer's data sheet.

Figure 17 shows the opposite configuration with 10K ECL logic driving either a single-ended 100K ECL receiver or a HOTLink Receiver. Here there are no tight margin areas between input and output thresholds. This means that 10K ECL parts can safely be used to drive 100K ECL inputs over their full temperature range.

Figure 17 also highlights the enhanced input range for the HOTLink Receiver. Unlike the narrow input range present on standard ECL families, the ECL inputs on the HOTLink Receiver maintain normal operation over the entire V_{CC} to $V_{CC} - 3V$ range.




Figure 16. 100K ECL Driving 10K ECL




Figure 17. 10K ECL Driving 100K ECL

Single-Ended Connections

Both of these comparisons are based on single-ended connections, where only a single ECL output is used to drive the receiving internally-referenced single-ended gate. In these cases, the other input to the receiving differential amplifier is connected internally to a V_{BB} reference. This type of connection should not be used to drive the INA_{\pm} or INB_{\pm} differential inputs of the HOTLink Receiver.

Differential Connections

One of the biggest advantages of ECL is the ability to communicate in a differential mode. This mode is relatively rare on logic parts (most commonly used for clock drivers and line receivers), as it requires both the driving and receiving parts to have both true and complement outputs and inputs respectively. When connected in this manner, the receiving part is no longer comparing the input signal to its V_{BB} reference, but instead compares the true and complement inputs to each other.

When used in this mode there is no problem using 100K ECL with 10K ECL at any temperature. Because an ECL receiver only requires around 250 mV of difference to fully switch, and the difference between the outputs of a differential driver remains near 800 mV, any differential connection has a minimum of twice the noise margin of a single-ended connection.

This type of connection is also immune to minor differences in the reference voltages between parts. Because the connection is differential, any common-mode voltages present on the received signals (due to power supply differences, AC coupling, ground shift, etc.) within the common-mode range are canceled out in the receiving differential amplifier. Some ECL parts with differential inputs can accept up to 1V of common-mode offset on the received signal without degradation of performance. The enhanced 100K ECL compatible inputs of the HOTLink Receiver can accept inputs between V_{CC} and $V_{CC} - 3V$, offering a common-mode range of 3V.

HOTLink Transmitter Connections

Unlike conventional negative-referenced ECL, the high-speed outputs on the HOTLink Transmitter are implemented in 100K positive-referenced ECL (PECL). This allows the TTL and ECL interfaces on the transmitter to operate from a common +5V power supply.

The HOTLink Transmitter has three differential output sections: $OUTA_{\pm}$, $OUTB_{\pm}$, and $OUTC_{\pm}$. In addition to operating as 100K ECL-compatible signals, these outputs have been enhanced with additional features.

Power Saving Mode

A standard ECL output structure uses a constant current source at the base of a differential amplifier (see *Figure 5*). In these standard parts, this current source is enabled and dissipating power even when the outputs are not used.

The HOTLink Transmitter ECL outputs, while still operating as true 100K ECL outputs, incorporate some additional structures (see *Figure 18*) to save power when the outputs are not used. The differential amplifier (D1) under normal conditions directs the I_S current from the current source through its internal transistors. As this current is switched, the output driver transistors (Q1 and Q2) change their operation point and the amount of current they source (a properly biased ECL output sources current in both 1 and 0 states; i.e., it never enters

cutoff). Each output driver (Q1 and Q2) contains a high value pull-up resistor ($RO+$ and $RO-$) and a voltage comparator (C1 and C2).




Figure 18. HOTLink Transmitter ECL Output

When both voltage comparators of a HOTLink differential output detect a voltage above a 100K ECL output-high level (V_{TH}), the current source (I_S) for that differential output pair is disabled. This results in a current savings of around 5 mA (25 mW) for each unused output pair.

FOTO Control of $OUTA_{\pm}$ and $OUTB_{\pm}$

The HOTLink Transmitter $OUTA_{\pm}$ and $OUTB_{\pm}$ differential outputs have an additional control input not present in the $OUTC_{\pm}$ output pair. While the $OUTC_{\pm}$ outputs are always enabled to follow the serial data stream generated in the HOTLink Transmitter shifter, the $OUTA_{\pm}$ and $OUTB_{\pm}$ outputs are not. These outputs are also controlled by a TTL-level input called FOTO (fiber-optic transmitter-off). While $OUTA_{\pm}$ and $OUTB_{\pm}$ are disabled, the $OUTC_{\pm}$ pair remains active and can be used for a local loopback source.

This FOTO signal is used to force the differential outputs of the $OUTA_{\pm}$ and $OUTB_{\pm}$ drivers to a state where a logical 0 is being driven. This state corresponds to a condition on optical modules where no (or minimal) light is transmitted. While not required for LED-based optical modules, this capability may be required for laser-based links (see ANSI Z136.1 and Z136.2, F.D.A regulation 21 CFR subchapter J, and IEC 825) (References 9, 10, 11, 12, 13).

ECL Output Biasing

ECL outputs have specific loading requirements to ensure proper operation. Because of the open-emitter structure of an ECL output, it can source current but cannot sink current. To allow the output to switch, some form of pull-down is required on the output. This pull-down usually takes the form of a resistive load; either to V_{EE} or $V_{CC} - 2V$.

Most ECL outputs are specified for driving load impedances as low as 50Ω. Because an ECL output does not swing rail-to-rail, this load is usually specified at $V_{CC} - 2V$, a point slightly below the ECL V_{OL} . At this point, when the ECL gate is driving a logic-0 signal, a small current is running through the load resistor to keep the output transistor in the active region. Typical currents sourced when driving a logic-1 (I_{OH}) and logic-0 (I_{OL}) are calculated using Equations 4 and 5 re-

spectively, where R_T is the effective load impedance and V_{TT} is the effective bias voltage.

$$I_{OH} = \frac{V_{OH} - V_{TT}}{R_T} = \frac{(-0.9) - (-2.0)}{50\Omega} = 22mA \quad \text{Eq. 4}$$

$$I_{OL} = \frac{V_{OL} - V_{TT}}{R_T} = \frac{(-1.7) - (-2.0)}{50\Omega} = 6mA \quad \text{Eq. 5}$$

These I_{OH} and I_{OL} values are the basis for the timing and signal levels in the HOTLink data sheet. For other values of I_{OH} and I_{OL} , the transmitter exhibits slightly different characteristics. These current flows can be achieved in many ways. The four most common methods are:

- Shunt bias to V_{TT} bias voltage
- Shunt bias to V_{EE} bias voltage
- Thévenin bias to V_{TT} bias voltage
- Y-bias to V_{TT} bias voltage

Shunt Bias to V_{TT}

In shunt bias, as shown in *Figure 19*, a single resistor is used as a pull-down load on an ECL output to some bias voltage. When biased to V_{TT} , a single 50Ω resistor (R_T) from the ECL output to V_{TT} is all that is necessary. This type of biasing requires an additional power supply to provide the ($V_{CC} - 2V$) V_{TT} -level. This bias type dissipates the least average-power (13 mW) of any output load type. It is often used in large ECL systems, in systems where overall power dissipation is a major concern, or where there is enough ECL present to warrant its design and implementation.




Figure 19. Shunt Bias to V_{TT}

Shunt Bias to V_{EE}

ECL outputs may also be biased to the V_{EE} supply as shown in *Figure 20*. Here a load resistance (R_T) of around 270Ω is connected to the V_{EE} supply to provide a similar current load for the ECL output driver. This value is determined by taking the average current flow for both a 1 and a 0, at the midway point (V_{BB}) in the output swing. The calculation for this is shown in Equation 6.




Figure 20. Shunt Bias to V_{EE}

$$R = \frac{V_{EE} - V_{BB}}{\frac{I_H + I_L}{2}} = \frac{5 - 1.3}{\frac{22 + 6}{2}} = 264\Omega \quad \text{Eq. 6}$$

Unlike the shunt bias to V_{TT} , this bias arrangement dissipates a significant amount of power in both the 1 and 0 states (47 mW average). This bias type (due to mismatched RC charge and discharge rates) exhibits a faster falling edge than rising edge. Because of this, its use is usually limited to logic functions, and is discouraged for serial links and for biasing differential output pairs. This is discussed in detail later in this document.

Thévenin Bias to V_{TT}

In a Thévenin bias network, a pair of resistors (R_1 and R_2) are used to create a load whose Thévenin equivalent matches that of a single resistor attached to a specific bias voltage (V_{TT}). For ECL this voltage is usually $V_{CC} - 2V$. These resistors are connected as shown in *Figure 21*. The values of R_1 and R_2 are solved using Equations 7 and 8.

$$R_1 = \frac{V_{EE} \times R_T}{V_{EE} - V_{TT}} \quad \text{Eq. 7}$$

$$R_2 = \frac{V_{EE} \times R_T}{V_{TT}} \quad \text{Eq. 8}$$




Figure 21. Thévenin Bias Equivalent

Solving for 50Ω and $V_{CC} - 2V$ yields values of 83Ω and 125Ω for a 5V system. While this combination does provide a similar dynamic load to the shunt bias to V_{TT} , it dissipates nearly an order of magnitude more power (138 mW) than its shunt to V_{TT} equivalent.

The capacitor shown in *Figure 21* is needed to allow R_1 and R_2 to provide the proper load for AC signals. In a Thévenin equivalent circuit, the power supply is assumed to be a short circuit. While this may be accurate for DC or very low frequency AC signals, the power supply appears as a near infinite impedance at RF frequencies. The bypass capacitor across R_1 and R_2 is used to create an AC short. This capacitor must be sized to operate as a short near the frequencies in use. For HOTLink-based systems this capacitor should probably be in the range of 300 pF to 0.01 μF .

Y-Bias to V_{TT}

Unlike the three previously described bias configurations, the, Y-bias can only be used with differential outputs. In this configuration the current from both ECL outputs are summed together through a single load to V_{EE} . Since the net sum of these currents remains constant as the output switches, the

voltage drop across this load remains constant. A schematic of this bias network is illustrated in *Figure 22*.




Figure 22. Y-Bias Network

Here R_T is the desired load impedance, usually 50Ω to $V_{CC} - 2V$ for ECL systems. R_L is determined by summing the currents of a logic 1 and a logic 0 (as shown in Equations 4 and 5), and calculating the resistance necessary to drop the remaining voltage. This calculation is shown in Equation 9 and solved for a 50Ω R_T .

$$R_L = \frac{V_{CC} - V_{EE} + V_{TT}}{I_{HI} + I_{LOW}} = \frac{3V}{28mA} = 107\Omega \quad \text{Eq. 9}$$

This type of bias provides a significant power savings over a Thévenin bias because only a single pull-down resistor is used to dissipate power for two outputs. For a 50Ω equivalent load the power dissipation is only 110 mW for two outputs (55 mW for one). Just as with the Thévenin bias, a capacitor is necessary to create an AC short.

Matched Loading

Just as the differential amplifier in an ECL switch directs current flow, so do the emitter-follower output transistors. As these transistors are turned on an off, large amounts of current are switched through the driver's V_{CC} package pins. Because of the inductance present in these pins, transients can be induced in the internal V_{CC} supply.

Fortunately the effects of this lead-inductance only manifest themselves when the current through the V_{CC} supply pin *changes*. If the current is kept stable, no transients are induced. Due to the differential configuration of many ECL outputs, it is possible to keep this current stable by having matched loads on the true and complement outputs of the differential driver. This means that if a design uses either one or both outputs of a differential driver, they both should drive loads of the same magnitude.

Figure 23 shows a differential output driver connected to a load including the package inductance present on the V_{CC} power pin. As the differential driver changes state, the overall current through $L1$ remains the same (assuming that both R_T loads are the same value).

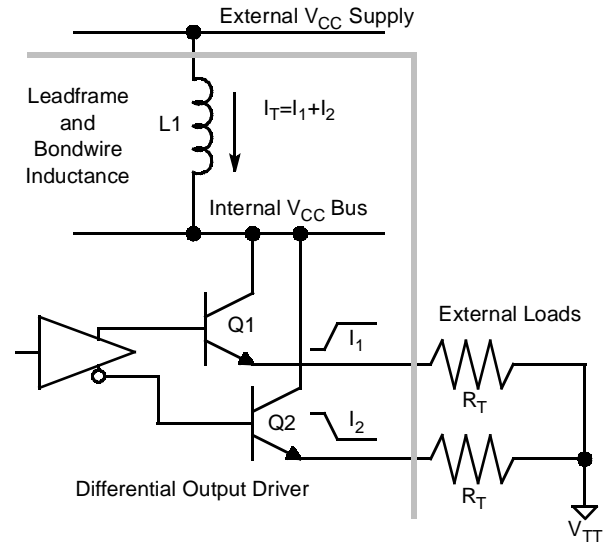


Figure 23. Loaded Differential Driver

If one of the two R_T load resistors is removed, some very undesirable things start to happen. First, the external power supply must now react to a dynamic rather than a static need for current. This increases the amount of power-supply bypassing that is needed next to the ECL driver V_{CC} pin. The second is a variation in the internal and external V_{CC} supplies caused by the dynamic current flow. This effect is examined in the following approximation.

For a single ECL output the current difference from a logic 1 to a logic 0 (into a 50Ω to $V_{CC} - 2V$ load) is 16 mA (see Equations 4 and 5). The ECL 100K family data sheets document that signal transition times may be under 500 ps. By assuming the rise and fall portions of the signal are related to a triangular waveform, this transition may be roughly converted to a fundamental frequency using Equation 10.

$$F = \frac{1}{2 \times T_r} = \frac{1}{2 \times 500E^{-12}} = 1GHz \quad \text{Eq. 10}$$

The Fourier series for a triangular waveform is listed in Equation 11. This illustrates that most of the energy content is present at the fundamental frequency with much smaller components present at the higher odd harmonics. To simplify the following calculations only the fundamental frequency is assumed to be present (Reference 24).

$$\frac{8V}{\pi^2} \left(\cos \omega_0 t + \frac{1}{9} \cos 3\omega_0 t + \frac{1}{25} \cos 5\omega_0 t + \dots \right) \quad \text{Eq. 11}$$

If a package pin inductance of 4 nH is assumed (typical for many surface mount components), Equation 12 can be used to determine the impedance of the package at this frequency.

$$X_L = 2\pi FL = 2\pi \times 1E^9 \times 4E^{-9} = 25\Omega \quad \text{Eq. 12}$$

Using Ohm's Law we can now convert this change in current into an internal voltage change, as illustrated in Equation 13.

$$V = I \times X_L = 16mA \times 25\Omega = 400mV \quad \text{Eq. 13}$$

This temporary difference between the internal V_{CC} and the external V_{CC} supply is the same phenomenon known in a TTL environment as ground bounce.

All of this, of course, is based on the assumption that the output is able to switch at this speed (500 ps) and provide the specified current (16 mA) when presented with a high-impedance source. What actually occurs is that the output edge slows down to match the current transfer permitted by the on-resistance of the output driver transistor and the package reactance.

Most ECL parts use a couple of different techniques to combat this problem. Both are quite simple to implement. The first is to use a separate package pin to provide power to the emitter-follower output transistors. This prevents any V_{CC} shift caused by the output drivers from affecting the sensitive differential amplifiers and voltage references present in other parts of the device.

The second method is to maintain a balanced load on the differential output drivers. Since the rising and falling edge rates of ECL are very symmetrical, $\Delta I_1 = \Delta I_2$. Because these changes in output current are symmetrical, $\Delta I_T \cong 0$. From Equation 13 we know that any induced ΔV is directly proportional to ΔI ; thus as ΔI goes to 0, so does ΔV .

AC Characteristics of Output Drivers

In an ECL driver, the time it takes for the signal to rise is largely determined by its internal resistors and parasitic capacitors (C_{int} and R_{int} in Figure 24), since the emitter-follower can supply large currents to charge the load capacitance. The DC voltage to which the output rises is determined by the emitter-follower transistor characteristics and the internal driver resistor (R_{int}) value. However, the AC voltage (overshoot, ringing, etc.) is determined primarily by the load characteristics. A capacitive load (along with the inductance found in the package, printed circuit traces, and other load components) causes the output to rise significantly beyond the anticipated DC output level, since the emitter-follower cannot supply any compensating current at the top of its transition.

Unlike the output rise time, the fall time is primarily determined by the time constants of the load capacitance and

pull-down circuit. The output LOW voltage (V_{OL}), is determined by R_{int} , I_S , and the characteristics of the emitter-follower transistor. In a properly designed system the load circuit has time constants comparable to (or shorter than) the internal fall time, such that the emitter-follower can source a small amount of current during the entire time it is switching from HIGH to LOW. If this is not true, the emitter-follower transistor will be shut off for part of the transition time, and the output will follow the time constant of the load.

Figure 25 illustrates the effects of two different load or bias circuits. The assumption in both of these examples is that the load circuit controls the fall time of the signal, and that the pull-down current is being supplied by a resistor to a V_T of either $V_{CC} - 2V$ or V_{EE} (+3V or ground for a PECL environment). In the dashed curve, the standard ECL load of 50Ω to $V_{CC} - 2V$ is used, causing an output current of approximately 20 mA when the output is HIGH, and 5 mA when the output is LOW. This load (or its equivalent) can be created using all of the previously described bias networks except shunt the bias to V_{EE} (shown in the solid curve).




Figure 24. ECL Output Driver with Loading




Figure 25. Falling Edge Rate Comparison for Bias to V_{TT} and V_{EE}




Figure 26. Expanded Detail of Falling Edge Rate Comparison

The same amount of pull-down current can be realized with a single resistor (R_L in Figure 24) in a shunt bias to V_{EE} configuration. To get a comparable output current (and assure comparable voltages at the output) the pull-down resistor is chosen to sink approximately the average of I_{OH} and I_{OL} when connected to a voltage midway between V_{OH} and V_{OL} (see Equation 6). The I_{OH} and I_{OL} currents listed here yield a pull-down resistor of around 300Ω. This type of bias is perfectly correct and adequate for ECL logic circuits where the mismatch between rise and fall times is absorbed into the normal logic delays and set-up times. In a data transmission system the effects of this type of output bias can be unpredictable and will often degrade performance.

Figures 25 and 26 illustrate the difference in output fall time assuming a constant load capacitance, with the only variation being the bias resistor and voltage. The 50Ω load resistor (dashed line) follows an RC discharge curve which ends at $V_{CC} - 2V$. For normal loading this soft edge rate more closely matches the rise time of the output as controlled by the emitter-follower, and is less affected by variations in load capacitance and reflection currents.

The 300Ω load resistor (solid line) follows an RC discharge curve which would normally end at V_{EE} (ground). While this appears to have a crisper edge rate, it is more severely affected by load capacitance variation and transmission line reflection currents that must be accommodated.

Figure 26 shows that with either pull-down the total voltage swing is the same, and is determined by the internal voltage swing of the driver, as buffered by the emitter-follower transistor. While the RC curve for the 300Ω pull-down continues to V_{EE} , the emitter-follower is turned on and sourcing current at the V_{OL} point and does not allow the output to continue farther down the curve.

In either configuration the signal delays match, since both falling edges cross the mid-swing line at approximately the same time, but the rise and fall times are different. These rise and fall times determine the higher frequency spectral components of the waveform. Differences in these spectral components affect the termination efficiency and waveform distortion caused by cable attenuation (Reference 14), and also effect the current distribution or “balance” on the output transmission line.

Transmission Line Termination

While often confused with ECL output biasing, termination of transmission lines is something quite different. Because of the reactive characteristics of transmission line termination, the resistors used for termination are often used as part of the output bias network, but they perform different functions.

Due to the high switching speeds of ECL, most of the interconnect between parts cannot be treated as simple connections. They must instead be treated as transmission lines. The distance between parts, in conjunction with the signal loading and rise and fall times, is used to determine at what point the interconnect must be treated as a transmission line. The general assumption is that short lines do not require termination, while long ones do. The determination of what is a long line is made using Equation 14 (Reference 5).

$$l_{\max} = \frac{1}{2} \sqrt{\left(\frac{C_L}{C_O}\right)^2 + \left(\frac{T_r}{\delta}\right)^2} - \frac{C_L}{2C_O} \quad \text{Eq. 14}$$

The values for this equation for microstrip construction on G10/FR4 type board would be:

- l_{\max} – maximum unterminated line length
- T_r – source 20% to 80% rise time
- C_L – load capacitance (2 pF assumed for a load)
- δ – delay per unit length (0.148-ns/inch)
- C_O – capacitance per inch

Running this calculation for various impedance and rise-time combinations yields the lengths listed in Table 3. Lengths beyond those listed here require termination.

Table 3. 100K ECL Maximum Unterminated Line Length (In Inches), Microstrip Construction

Z_O	Line Length (in inches)		
	0.5 ns	1 ns	1.5 ns
50Ω	1.38	3.06	4.74
62Ω	1.32	2.99	4.67
75Ω	1.25	2.91	4.59
90Ω	1.18	2.82	4.50
100Ω	1.14	2.76	4.44

The lengths listed in *Table 3* assume digital switching characteristics. The HOTLink ECL serial signals are, for the most part, analog in nature. This effectively shortens the maximum unterminated length. For HOTLink serial signals, any ECL trace greater than one inch (2.5 cm) in length should be terminated.

The objective of transmission line termination is to prevent reflection of power from the destination back to the source. This is accomplished by terminating the transmission line in its characteristic impedance (Z_0). The two basic types of line termination are referred to as series and parallel termination.

The actual amount of the source signal reflected is based on how well the line impedance matches the destination impedance. This determines how much voltage is reflected back into the transmission line. This ratio of reflected voltage to incident voltage is called the reflection coefficient ρ (rho) and is shown in Equation 15 (Reference 5).

$$\frac{V_r}{V_i} = \rho = \frac{R_T - Z_0}{R_T + Z_0} \quad \text{Eq. 15}$$

Series Termination

Series termination (sometimes referred to as source termination) requires that the load be high-impedance to properly operate. This type of line termination is **not** recommended for use with HOTLink because of the reactive nature of all parts at the high frequencies present on the HOTLink ECL signals.

Parallel Termination

In parallel termination the desired characteristic is to terminate the *end* of the line (rather than the source) in its characteristic impedance. This results in a reflection coefficient of zero; i.e., no signal is reflected. This type of termination is implemented the same as shunt bias networks. *Figures 27 and 28* show the two equivalent forms of parallel termination.

Parallel termination offers the advantages of allowing distributed loads on the transmission line, and of having the termination network also operate as the bias network.

In the single-resistor form of parallel termination shown in *Figure 27*, the R_T resistor is sized to match the Z_0 impedance of the transmission line. This termination form has the same advantage as the single resistor shunt bias because it dissipates less overall power than the Thévenin equivalent termination. It also has the same drawback of requiring a separate power supply.




Figure 27. Parallel Termination to V_{TT}

In a Thévenin equivalent termination (shown in *Figure 28*) two resistors ($R1$ and $R2$) are used to form an equivalent circuit to that in *Figure 27*. *Table 4* lists the $R1$ and $R2$ resistor val-




Figure 28. Thévenin Equivalent Parallel

ues for a number of common transmission line impedances. This table assumes operation with a 5V source and a termination voltage of $V_{CC} - 2V$, and selects the nearest standard 1% resistor value when an exact match is not available. These values are calculated using the same Equations 7 and 8 as used for calculating a Thévenin bias network (Reference 15).

Table 4. Thévenin Bias Resistor Values

Z_0	$R1$	$R2$
50 Ω	82.5	124
70 Ω	118	174
75 Ω	124	187
80 Ω	133	200
90 Ω	150	226
100 Ω	165	249
120 Ω	200	301
150 Ω	249	374

Terminating HOTLink Transmitter ECL Signals

The HOTLink CY7B923 Transmitter has three different ECL differential output pairs named $OUTA_{\pm}$, $OUTB_{\pm}$ and $OUTC_{\pm}$ (see *Figure 2*). How (or if) these outputs are terminated depends on what the output is used for.

$OUTC_{\pm}$

The $OUTC_{\pm}$ outputs of the HOTLink Transmitter are not controlled by the transmitter FOTO signal and are thus always enabled to drive serial data. While fully capable of driving either optical modules or copper cables, it is expected that the most common usage of this differential output will be as a local loopback to a HOTLink CY7B933 Receiver INB_{\pm} inputs.

This signal may be connected to the HOTLink Receiver either differentially or single-ended. When connected differentially, the $OUTC+$ output is connected to the $INB+$ input, and the $OUTC-$ output is connected to the $INB-$ input. When connected single-ended, the $OUTC+$ output is connected to the $INB+$ input.

Note: For the $INB+$ input to be used differentially, the SI/SO ECL-to-TTL translator (mapped through the $INB-$ input) must

be disabled. This is done by connecting the SO output directly to V_{CC} .

Once the connection is made, the type of termination required is determined by the distance between the HOTLink Transmitter and the HOTLink Receiver. If the distance is kept short enough (under one inch) no termination is required and the output only needs to be biased (Reference 5). This can be done with a single pull-down resistor to V_{EE} . While this type of termination does induce some jitter into the serial data stream (due to mismatched rise and fall times), the amount is well within the receiver limits.

If the distance is greater than one inch, the line should be terminated (Reference 5). To do this correctly requires determination of the characteristic impedance of the board traces used to connect the source and destination. Please see the Cypress Semiconductor application note "Driving Copper Cables with HOTLink" for information on how to determine the characteristic impedance of various types of transmission lines (Reference 16).

For local connections that do not travel through external transmission media (i.e., coax, twisted-pair, optical fiber, etc.) parallel termination may be used. The important consideration here is that both the OUTC+ and OUTC- outputs must be terminated/biased into the same size of load to maintain a current balance inside the HOTLink Transmitter.

If neither of the OUTC± outputs are used, both outputs should be left open or pulled up to V_{CC} to disable the current source for the differential driver (see Figure 18).

OUTA± and OUTB±

The OUTA± and OUTB± outputs of the HOTLink Transmitter are both controlled by the FOTO signal which is required to meet laser safety regulations for communications links (References 9, 10, 11, 12, 13). Other than this special enable signal, these outputs operate the same as the OUTC± outputs.

Driving Optical Modules

When connecting to optical modules, it is best to drive the optical module data inputs differentially. This provides the highest noise immunity for the system, and the lowest signal jitter. When used with *de facto* standard optical modules this becomes mandatory because the optical modules have a differential data input, yet do not provide a V_{BB} supply to bias the other input of the differential amplifier of the optical transmitter. Because this interface is intended for driving some external segment of optical cable, series termination (which uses shunt bias to V_{EE} and increases jitter) should not be used. Since the HOTLink parts are most probably be the only ECL parts in the system, the recommended termination is a Thévenin or Y-termination.

Both the Thévenin and Y-terminations provide the bias necessary for the ECL signal to switch, and the impedance necessary to terminate a transmission line. One of these types of termination/bias should be used even when the distance from the HOTLink Transmitter to the optical transmitter is short. This is necessary to maintain symmetrical rise and fall times for the OUTx± differential outputs.

PECL Optical Modules

Interfacing to optical modules in PECL mode is quite simple, requiring only a few passive parts. The schematic in

Figure 29 illustrates the connections and parts necessary for this type of connection.




Figure 29. HOTLink Transmitter-to-PECL Optical Module

One of the key items often missed in this type of connection is proper bypassing of the termination/bias networks. The theory behind a Thévenin network is that the power supply is considered as a short for AC. While this may be true for near DC applications, the base frequencies and harmonics present in the HOTLink Transmitter output are far beyond any frequency the power supply itself could pass.

To make the power supply a short, a capacitor must be placed across the Thévenin pair. The size of the capacitor is determined by the frequency of operation of the serial link. A good rule-of-thumb is to pick the largest value capacitor whose series resonant frequency is 30% above the highest baseband frequency of the Baud rate of the serial data (Reference 17). Since the data is sent using an NRZ modulation (non-return-to-zero), the highest baseband frequency is one half the serial bit-rate (Reference 18).

Another important characteristic is the dielectric type of the capacitor. For this type of analog operation, a good high-frequency RF-type capacitor must be specified. This means specifying either NP0 or COG type dielectric.

Standard ECL Optical Modules

Those optical modules with the case connected to V_{CC} are designed for use in a negative DC supply system. These types of modules may also be driven by a HOTLink Transmitter.

By far the simplest method is to connect the module the same as a PECL module, with the exception of the Case pins. Here, instead of attaching the Case pins to ground (V_{EE}), they are attached to V_{CC} . If the case is metallic in nature, care must then be exercised such that it does not come into direct contact with ground.

If the optical module is to be used below ground, it must be AC coupled to the HOTLink Transmitter. This type of connection is shown in Figure 30.




Figure 30. HOTLink Transmitter-to-Negative-Referenced ECL Optical Module

The HOTLink Transmitter outputs are biased the same as for a PECL optical module. AC-coupling capacitors are used to connect the HOTLink Transmitter positive-referenced ECL outputs to the negative-referenced ECL inputs of the optical module. These coupling capacitors actually operate as a bandpass filter, centered around their series resonant frequency. To pass additional low- or high-frequency components, additional capacitors should be placed in parallel with the coupling capacitors.

Capacitively coupled signals require DC restoration and, if the connection length warrants, transmission line termination. DC restoration is necessary to place the signal swings in the input range of the ECL receiver. Unlike ECL outputs, which are biased to a level slightly below their $V_{OL}(\text{min})$ -level ($V_{CC} - 2V$), AC-coupled ECL inputs need to be biased to the center of the receiver input range. This is the same as the V_{BB} reference point of $V_{CC} - 1.3V$. In Figure 30, this reference point is created from a resistive divider network, and bypassed with a $0.01\text{-}\mu\text{F}$ capacitor to provide the dynamic current response needed for the differential inputs.

While many optical modules or ECL gates generate a V_{BB} level, this output *must not* be used to bias this reference point because it cannot provide sufficient dynamic current. The V_{BB} output of an optical module, or other ECL gate, is an unbuffered tap of the internal V_{BB} reference. While fully capable of delivering the few μA of current necessary to drive an input, it cannot tolerate the transient currents present at the end of a low-impedance transmission line. Because the V_{BB} source is unbuffered, this also means that any external transients applied to it will move the V_{BB} reference inside the receiver, with unpredictable consequences.

While it is possible to create a V_{BB} power amplifier (by using multiple ECL buffers in parallel) to create a buffered form of V_{BB} , such amplifiers should not be used with HOTLink. They are prone to oscillation and ringing. Such amplifiers should also not be used for DC restoration (as needed here) because the V_{BB} amplifier is not quite DC stable; i.e. its output usually contains a low-level (10–50 mV) oscillation whose frequency is set by the delay through the part. This low-level noise is not a problem for logic applications, but for analog applications it causes increased jitter on the biased signals.

In this example, the V_{EE} for the optical module is set to $-5.2V$. This is a common supply voltage for ECL circuits. If a different supply voltage is used, the values in the resistive divider must be changed to maintain the V_{BB} reference point at $V_{CC} - 1.3V$.

One drawback of this circuit is the inability to react to a DC state in the data stream. If the HOTLink Transmitter is set to transmit all 1s or all 0s (e.g., FOTO is set to disable transmitting), the optical module inputs will both return to a V_{BB} level. At this level the optical module's output will probably oscillate due to the high gain present in the optical module's ECL-to-optical translator. In this AC-coupled configuration (when operated with laser-based optical drivers) it is necessary to use some method other than FOTO to meet the laser safety restrictions (References 9, 10, 11, 12, 13).

Driving Copper Media

The ANSI Fibre Channel Standard currently identifies both coaxial cable and shielded twisted-pair as supported copper media types. The HOTLink Transmitter easily interfaces to these and many other types of copper media, and allows communication on them at distances well beyond the lengths called out in the ANSI Standard (Reference 3).

Numerous characteristics determine how far a signal can be transmitted on copper media. The most important of these are:

- Voltage amplitude of the signal fed into the cable
- Jitter and ringing on the source signal
- Attenuation characteristics of the cable
- Length of the cable
- What (if any) equalization is used in the system
- Receiver loading and sensitivity

Coupling to the cable (transmission line if on a backplane) may be done in multiple ways, depending on the media type and distances involved.

Direct Coupled

For those instances where the signal never leaves the same chassis (or even the same board) it is possible to directly couple to the media. Here the media is effectively the circuit board traces, runs of twisted-pair, twinax, or dual coax. The main criteria here is that there must be no chance for a significant V_{CC} reference difference (transient or DC) between the HOTLink Transmitter and HOTLink Receiver, including any common-mode induced noise. For the HOTLink Receiver, this maximum difference is around 1V. Under these conditions the HOTLink Transmitter and Receiver may be connected as shown in Figure 31.




Figure 31. Direct-Coupled, Copper Interface

While *Figure 31* shows a pair of 50Ω transmission lines, the actual impedance can be higher or lower than this. For other impedance values it is necessary to change the Thévenin termination networks.

When sent through twin coaxial cables (as shown in *Figure 31*) or two separate transmission lines, care must be taken to make sure that both lines are electrically the same length. Any difference in length causes one of the two transmitted signals to arrive at the receiver input either leading or lagging the other. This difference manifests itself as jitter in the receiver. If twisted-pair or twinax is used instead, both the OUTA+ and OUTA– signals combine to form a single signal sent down a balanced transmission line.

Capacitor Coupled

For configurations where it is possible to have significant ground or reference differences, some form of AC coupling becomes necessary. If the signals remain in a well protected environment (minimal EMI/ESD exposure) this AC coupling can be performed with capacitors. When this is done, bias/termination networks are required at both ends of the cable. A schematic detailing this type of connection is shown in *Figure 32*.

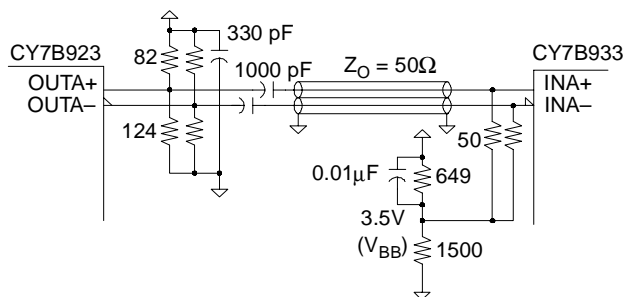


Figure 32. Capacitive-Coupled, Copper Interface

Good low-loss RF-grade capacitors should be used for this application. These parts are available in many different case types and voltage ratings. The capacitors used must be able to withstand not just the voltage of the signals sent, but of any DC difference between the transmitter and receiver and the maximum ESD expected. A typical 1000-pF 50-WV C0G capacitor would be 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-WV 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 C0G capacitor at 5-kV breakdown is almost a half cubic inch in size (Reference 15).

This type of coupling is very similar to that used to drive an optical module that is not at the same reference as the HOTLink Transmitter. Since the HOTLink Receiver and an optical module both operate with ECL 100K-level compatible inputs, this should be expected.

In this configuration, the receiver reference point is set slightly different from that for 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 V_{BB} bias network on the receive end of

the transmission line is set to the center of the HOTLink Receiver 3V common-mode range at $V_{CC} - 1.5V$.

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

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

Addition of a second set of coupling capacitors at the receive end may improve some of these characteristics, but it will not remove them.

Transformer Coupled

The preferred copper attachment method is to transformer couple to the media. 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. A schematic detailing a transformer-coupled interface is shown in *Figure 33*.




Figure 33. Transformer-Coupled, Copper Interface

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 and termination network at the receiver.

The connection at the HOTLink Transmitter is quite different now. The output bias network is now a simple pull-down to V_{EE} . While this causes the transmitter outputs to have asymmetric rise and fall times, it does not add to the system jitter. Instead, the true and complement outputs combine in the transformer primary to provide a single signal with symmetrical rise and fall times. This bias arrangement also has the advantage of delivering the entire transmitter output voltage swing into the transformer, rather than part into the transformer and part into the bias network.

The configuration shown in *Figure 33* uses only a single transformer and either 150Ω twinax or twisted-pair as the transmission line. This can be done because the transmission system remains balanced from end to end. Here the primary

functions of the transformer are to provide isolation and common-mode cancellation.

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.

In *Figure 34* a second transformer is added to the transmission system at the receiver end of the cable. This configuration allows use of either balanced or unbalanced (coaxial) transmission lines. The configuration shown here is a 75Ω coaxial cable system. Here the first transformer is used for balanced-to-unbalanced conversion, while the second transformer provides unbalanced-to-balanced conversion.




Figure 34. Dual Transformer-Coupled, Copper Interface

HOTLink Receiver ECL Inputs

The HOTLink Receiver has five 100K ECL (PECL) compatible inputs: INA+, INA-, INB+, INB-(SI), and A/B (see *Figure 3*). The A/B input is used to select which serial data input (INA± or INB±) is fed to the receiver PLL and shifter.

The INA± differential input is normally used for the primary received data input. This input is only functional as a differential receiver. To use it as a single-ended receiver, a V_{BB} reference would have to be attached to one of the INA± inputs. Since the HOTLink Receiver does not provide a V_{BB} output, this must come from either an external ECL gate or a resistive divider. Because neither of these sources can be guaranteed to be at the exact internal V_{BB} reference of the HOTLink Receiver (and will thus introduce jitter into the system), operation of INA+ in single-ended mode is not recommended. Also, operation in single-ended mode generally takes twice the signal swing (100 mV for HOTLink) for a receiver to properly detect data.

The INB± differential input is expected to be used as the local loopback receiver, but may also be used for external data reception. It is capable of being operated as a differential receiver, or as two single-ended receivers.

To operate the INB± inputs as a differential receiver it is necessary to have the SO output either directly connected to V_{CC} or pulled up to V_{CC} through a low-value resistor (minimum input voltage of V_{CC} - 250 mV). This pin, while normally used as an output, has a voltage comparator on the output to both disable it and to operate the INB± inputs as a differential pair. When used as a differential receiver the INB± inputs operate the same as the INA± inputs.

If the SO pin is instead allowed to remain in the standard TTL output range (below V_{CC} - 850 mV), it is enabled as a

TTL-level driver, and is the output end of an ECL-to-TTL level translator. In this mode the HOTLink Receiver INB+ input is a single-ended ECL receiver for serial data, while the INB- input becomes the input end of the ECL-to-TTL translator. The expected use of this translator is for converting an ECL carrier-detect signal to TTL levels.

ECL Input Levels

Unlike standard 100K ECL logic, the HOTLink ECL inputs are designed to operate, not only over the full 100K ECL voltage and temperature range, but substantially beyond.

Normally, 100K ECL inputs should never be raised above V_{CC} - 700 mV. If this occurs, the input transistor saturates and can damage other internal structures in the gate. Because the HOTLink Receiver is designed for use in a communications environment, its input structures are more robust and can be taken all the way up to V_{CC} with no degradation in performance. This provides a common-mode operating range more than twice that of standard ECL.

The HOTLink ECL Receivers also provide higher gain than that available from standard 100K ECL. The receiver is able to fully detect 1s and 0s with as little as 50 mV of differential signal at the inputs. Those few 100K ECL parts capable of differential operation usually specify this at 150–200 mV.

The HOTLink ECL inputs are also robust on the V_{IL}(min) side. When operated in differential mode these inputs provide full functionality down to V_{CC} - 3V, yielding a full 3V common-mode operating range. For single-ended operations these same inputs can be taken all the way to V_{EE} (ground or 0V).

Controlling A/B from TTL

While the A/B path select on the HOTLink Receiver is a PECL input, it can be controlled from a TTL driver with as few as two resistors. Controlling a traditional PECL input from TTL normally requires a third resistor to limit the HIGH state to the specified V_{IH}(max). Only a two-resistor divider is needed with the HOTLink Receiver (as shown in *Figure 35*) because it can tolerate a full V_{CC}-level on its ECL inputs.




Figure 35. TTL-to-HOTLink PECL Interface

HOTLink Receiver Biasing

Unlike ECL outputs, which always require an output bias to create the output-low level, ECL inputs instead require levels within their input range to allow them to switch. When the HOTLink Receiver is directly connected to the biased output of either a 10K, 10KH, or 100K ECL driver (see *Figures 16* and *17*), these conditions are satisfied.

PECL Optical Modules

Connecting a PECL optical module to the HOTLink Receiver is the same as connecting two ECL parts together. This connection is shown in *Figure 36*.

A bias network is required on the output of the optical module to allow it to switch. A Thévenin (or Y-bias) network should be used on the high-speed serial lines (RO and NRO as shown in *Figure 36*) to keep induced jitter to a minimum. The signal- or carrier-detect output (SIGO) of the module is considered a logic level signal and only requires a pull-down type of biasing to allow the output to switch.

If the distance between the optical module and the HOTLink Receiver is short (see *Table 3*) then the bias network may be placed anywhere between the optical module and the HOTLink Receiver. If this distance is long, then the interconnect traces must be treated as a transmission line and the bias network must be moved to the receiver to also act as line termination. If the transmission line impedance is other than 50Ω, then different values of resistors are necessary (see Equations 7 and 8, and *Table 4*).

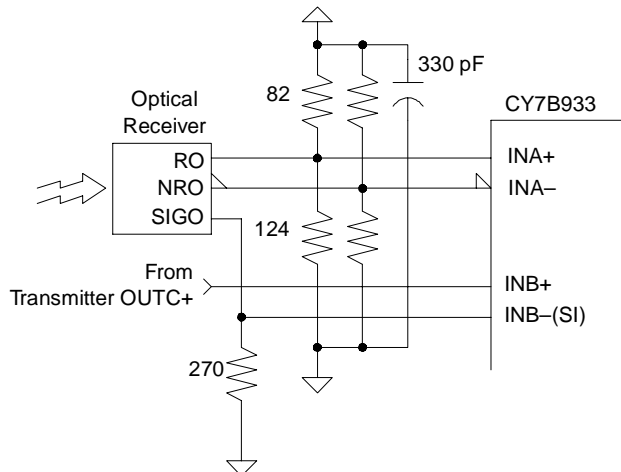


Figure 36. PECL Optical Module-to-HOTLink Receiver

Standard ECL Optical Modules

Optical modules with the Case pins connected to V_{CC} are designed for use in a negative DC supply system. These types of modules may also drive a HOTLink Receiver.

By far the simplest method is to connect the module the same as a PECL module, with the exception of the Case pins. Here, instead of attaching the Case pins to ground (V_{EE}), they are attached to V_{CC} . If the case is metallic in nature, care must then be exercised so that it does not come into direct contact with ground.

If the optical module is used below ground it must be AC-coupled to the HOTLink Receiver. A schematic detailing this type of connection is shown in *Figure 37*.

From a parts-count standpoint this type of connection should be avoided if at all possible. Just as with the HOTLink transmitter-to-negative referenced ECL optical modules, this interface requires biasing on both sides of the AC-coupling capacitors.

Because the signal detect output of the optical module is not an AC signal, capacitive coupling cannot be used to feed this signal into the HOTLink Receiver INB- input. The simplest thing to do here is to use an external ECL-to-TTL translator (as shown in *Figure 37*) to convert the signal-detect output to a positive referenced TTL environment.

The INA± differential inputs must be biased to near the midpoint of the common-mode range of the HOTLink Receiver. The two 50Ω resistors tied to this synthesized reference point are sized to properly terminate the transmission line impedance of the interconnect.

Receiving from Copper Media

The direct-coupled, capacitor-coupled, and transformer-coupled configurations for copper interconnect are covered in the HOTLink Transmitter-to-copper interface section of this document, with schematics of these connections shown in *Figures 31 through 34*.

Signal-Detect for Copper Interface

When interfacing to optical modules, the generation of a carrier- or signal-detect function is a simple connection to an ECL output. With a copper interface, this signal-detect function must be built from other components.

The key to a good signal-detect implementation is to create one that accurately detects the presence or absence of a valid data stream, yet does not significantly load or distort the received signal. A sample carrier-detect circuit is shown in *Figure 38*.

This circuit uses a reference divider-network similar to that in *Figure 37*, except that an additional voltage reference point is created. This new reference point sets a threshold for received amplitude at which the signal detect circuit starts to respond. For this example, this reference point is set to 100 mV above the carrier detect receiver V_{BB} reference point. This 100-mV offset is also necessary to prevent the 10H116 amplifiers from oscillating when no signal is present.

A 10H116 was selected here for numerous reasons. It is small (20-pin PLCC), fast (1 ns), and does not have 50-kΩ pull-down resistors built into its input structures. While these pull-down resistors (present on most ECL parts) are very handy for logic design, they have a significant impact when used for fast analog applications as done here.

Two sections of the 10H116 are used as received signal level comparators. One looks for logic-1 levels while the other look for logic-0 levels. The output of these two comparators are wire-ORed together and feed an RC network. The capacitor in this network is charged when either of the comparators is turned on, and discharges through a bleeder resistor when neither comparator is on.

The third section of the 10H116 also operates as a comparator, evaluating the voltage level on the RC network. Because the level on this capacitor changes so slowly, and ECL operates as an analog amplifier, positive feedback was added to cause the comparator to switch faster and to full ECL levels. The amount of hysteresis is set by the feedback resistor. For slow changing signals of this type, a minimum of 150 mV of hysteresis is recommended.

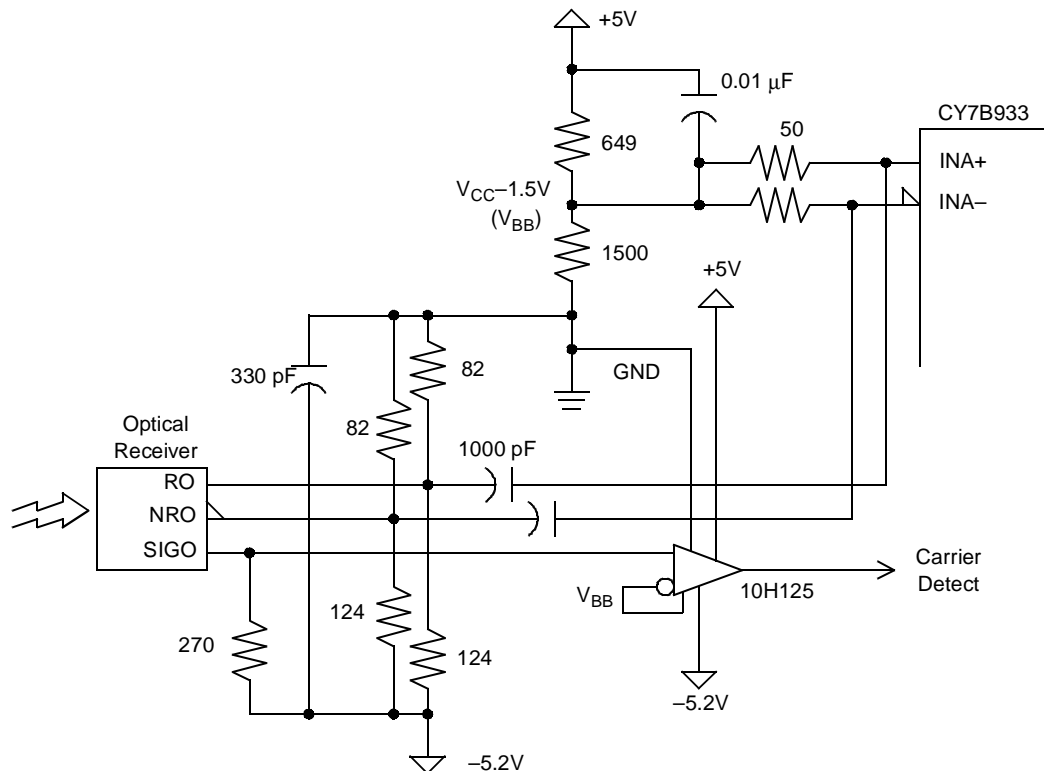


Figure 37. Negative-Referenced Optical Module-to-HOTLink Receiver

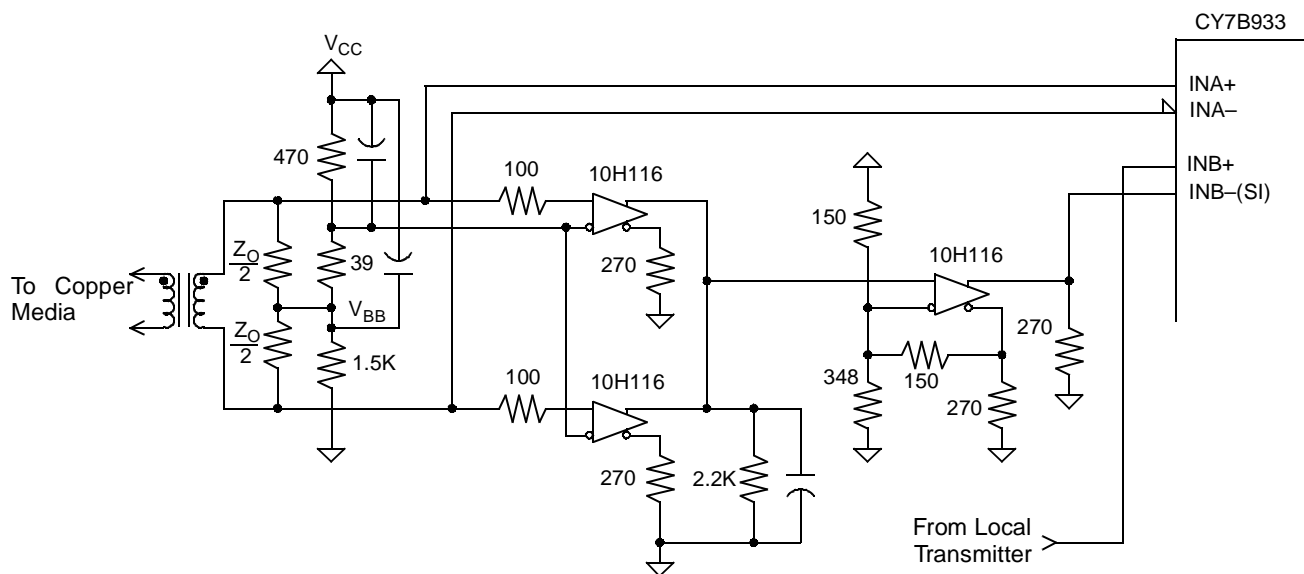


Figure 38. Copper Interface Signal Detect Circuit

Copper Signal Characteristics

Communication on copper-based media is very similar to communication on optical fiber. Both suffer from increasing signal degradation with increasing media length.

The transmitted signal is composed of multiple-frequency components, and requires a fairly wide-bandwidth media to propagate those signal components. A large part of the bandwidth requirement is determined by the 8B/10B code and NRZ modulation used in HOTLink for communication.




Figure 39. Short Time Constant Transmission Line Response

NRZ Modulation

NRZ is an acronym for non-return-to-zero. This is one of the most basic types of data encoding whereby a signal is HIGH for a 1 and LOW for a 0. The upper waveform in *Figure 39* illustrates an NRZ data stream. Other forms of encoding (Manchester, Biphasic, etc.) are used in data communications that encode clock information as part of the 1s and 0s. With an NRZ data stream, a phase-locked-loop is necessary to recover the bit-clock to allow data to be captured (Reference 18).

8B/10B Code Dependencies

A phase-locked loop (PLL) requires transitions meeting specific criteria to allow it to correctly recover a clock. If binary data were sent serially using only an NRZ encoding, long periods could exist where no transitions are sent. During these periods (if they are long enough) the receiving PLL can drift such that it is no longer able to properly recover the data sent. 8B/10B encoding is used to ensure that sufficient transitions are present in the NRZ data stream such that the receiving PLL remains synchronized to the data.

The 8B/10B code is a run-length limited code. This means that there are limits to the maximum and minimum length of a continuous sequence of 1s or 0s in the data stream. The code operates by converting an 8-bit data byte (with uncontrolled transition density) into a 10-bit transmission character (with controlled transitions). The 8B/10B code is referred to as a 1:5 code because the minimum number of consecutive 1s or 0s is one, while the maximum number is five (References 1, 2).

Translating these code limits into frequencies gives the baseband limits of the code. For example, with a serial bit-rate of 300 MHz, a pattern sent with the maximum number of consecutive 1s and 0s (five high, five low) would be equivalent to a 30-MHz square wave. Using the highest rate of alternating bits of 0 and 1 gives a frequency of 150 MHz.

As far as signal propagation goes, these numbers only refer to a sinusoidal frequency. Since square waves are used at the source, there are many additional higher-frequency harmonics present. To propagate a reasonable signal it is recommended that the system bandwidth also include (at a minimum) the 3rd harmonic of the highest baseband frequency, and preferably through the 5th harmonic.

For our previously described example operating at a bit-rate of 300 MHz, the necessary system bandwidth would be:

$$BW = (3 \times F_{MAX}) - F_{MIN} \quad \text{Eq. 16}$$

$$BW = (3 \times 150 \text{ MHz}) - 30 \text{ MHz} = 420 \text{ MHz} \quad \text{Eq. 17}$$

Transmission Line Effects On Serial Data

In a perfect world a perfect square wave could be launched down a perfect transmission line and it would come out the end looking the same as when it went in. Unfortunately, the laws of physics make such a transmission line impossible.

Instead, transmission lines have significant amounts of parasitic capacitance, inductance, resistance, and the terminations are reactive in nature. This means that a lossy system exists. The cable attenuation characteristics of copper cables are such that the higher frequencies have greater losses than the lower frequencies (see *Figure 77* for some sample cable attenuation curves).

When data is sent through such a lossy medium, distortion occurs. The higher-frequency spectral components are significantly reduced in amplitude, while the lower-frequency spectral components are reduced by a lesser amount. In addition, the higher-frequency spectral components propagate faster than the lower-frequency components. The square waves fed into the cable come out looking like RC charge/discharge curves.

These frequency-selective losses are equivalent to a time constant. For very short transmission lines (or very slow data rates) this time constant is short enough that transmitted 1s and 0s can completely charge or discharge the transmission line for each bit sent. The input and output signal waveforms for a transmission line of this type are shown in *Figure 39*.

Because the line can fully charge or discharge on even the fastest possible transition, the time to reach the receiver threshold is always the same. This allows the data out of the receiver to look just like the data sent into the transmission line.

As a transmission line is lengthened, its time constant increases. When the time constant is large enough that the line can no longer be fully charged and discharged in a single bit-time, the received data edges become time-displaced from their desired positions. Since coding theory refers to each transmitted 0 and 1 as a symbol, this type of distortion is called *intersymbol interference* or ISI. For communications systems, distortion of this type is called Data-Dependent Jitter (DDJ).




Figure 40. Long Time Constant Transmission Line Response

Input and output waveforms for a long time constant transmission line are shown in *Figure 40*. The receiver output is added to illustrate the edge displacement. As the transmission line becomes increasingly longer it is even possible for some single-bit transitions to not be detected at all by the receiver (based on the data pattern sent) because they fail to cross the receiver threshold. This may be corrected through use of frequency compensation circuits at either the source (precompensation) or destination (equalization) ends of the transmission line.

8B/10B Code Running Disparity

The 8B/10B code attempts to limit the maximum distance (voltage) from the receiver threshold that a transmitted signal can reach, by controlling the DC signal content of the characters sent and the maximum separations between 1s and 0s used to represent each character. To do this the 8B/10B code provides two 10-bit transmission characters to represent each 8-bit data value. The difference between these characters is the ratio of 1s to 0s. To determine which of the two characters to send, the HOTLink Transmitter evaluates the previously sent character for the density of 1 and 0 bits (when operated with the encoder enabled). If the net result is positive running disparity (more 1s than 0s), the following data byte is encoded using the form with more 0s than 1s. If the net result is negative running disparity (more 0s than 1s), the following data byte is encoded using the form with more 1s than 0s. The goal of this is to maintain as near as possible a net value of DC over time for the serial data sent to minimize baseline wander.

Baseline Wander

Methods of data encoding that are not DC balanced (i.e., 4B/5B as used with FDDI) suffer from a characteristic known as baseline wander. This is a side effect of an AC coupled system attempting to propagate a signal that contains a DC component.

Baseline wander is a (relatively) long-term, low-frequency effect, generated when the average DC-level of a transmitted signal varies with the data sent. This DC component is lost because the transmission system is AC-coupled. At the receiving end of the cable this appears as data that does not remain centered around the receiver threshold. This effect is shown in *Figure 41*.




Figure 41. Baseline Wander Example

If the receiver was actually presented with perfectly square pulses (with transitions that always crossed the receiver threshold) then baseline wander would not be a problem. Unfortunately, what are actually sent and received are more in the form of trapezoids or curves with measurable rise and fall times. The farther a signal drifts from being centered around the receiver threshold, the more that the threshold crossings are time displaced. This time displacement is also known as jitter.

Jitter

Jitter is a high-frequency deviation from the ideal timing of an event. Many different aspects of a serial link can affect the total jitter present in the link. Those based on real and repeatable direct measurements are referred to as *deterministic jitter*. Other effects, which are not directly repeatable and are more probabilistic in nature, are called *random jitter*.

Deterministic jitter itself may be broken into two major components: those based on the accuracy of the duty cycle of the information, and those based on the interaction of the 1s and 0s due to the limited bandwidth of the transmission system. The jitter that affects adjacent edges and duty cycle is called Duty Cycle Distortion (DCD). The jitter based on the data patterns sent is called Data-Dependent Jitter (DDJ).

Data-Dependent Jitter Characteristics

Data-Dependent Jitter (DDJ) is a measurement of intersymbol interference based on the maximum timing deviations caused by a worst-case data pattern. DDJ is affected by many environmental characteristics, in addition to the code used. These include the length of the cable, the attenuation characteristics of the cable, the integrity of the signal launched into the cable, and how well the cable is terminated. Because of the frequency-selective attenuation present in copper cables,

DDJ is one of the main limiting factors on how far a recoverable signal may be sent.

To measure DDJ for a specific configuration, data patterns having specific characteristics need to be repeatedly launched into the cable. These patterns must present the worst-case transition characteristics based on the code used for sending data. This is usually described in terms of sequential combinations of long and short 0s and 1s.

A long 0 or 1 is specified as the longest continuous LOW or HIGH that can be sent. For the 8B/10B code this is five bits in length. The short 0 or 1 is the shortest LOW or HIGH that can be sent. For the 8B/10B code this is one bit in length. The sequences used for testing are diagrammed in *Figure 42*.

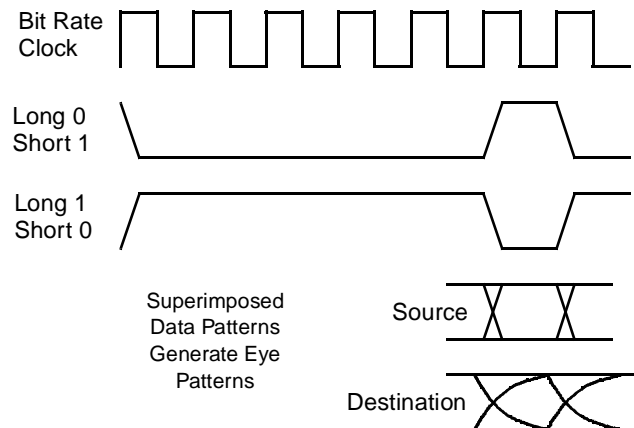


Figure 42. Eye Pattern Generation Waveforms

A design feature of the HOTLink Transmitter is that when neither data enable is active (EN_A and EN_N both HIGH), the part repeatedly sends out both disparity versions of the K28.5 SYNC code. This generates a 20-bit repeating pattern of 00111110101100000101. This alternating pattern contains the necessary combinations of long and short 0s and 1s for performing a proper eye pattern test.

The opening of the “eye” (see *Figure 43*) relative to the width of a bit cell is a good measure of link integrity. As this window gets smaller, it becomes more difficult for the HOTLink Receiver PLL to determine where to sample each bit cell (Reference 5).

The maximum variation, from early to late, of when the received signal crosses the receiver threshold is equal to the amount of jitter present. This jitter is usually expressed as a percentage relative to the width of a bit-cell window. This relationship is shown in Equation 18.

$$\text{Jitter} = \frac{\text{Bit}_{\text{TIME}} - Th_{\text{VAR}}}{\text{Bit}_{\text{TIME}}} \times 100\% \quad \text{Eq. 18}$$

The oscilloscope illustration in *Figure 44* is an actual DDJ measurement based on a 100-foot (30.4m) segment of RG59 cable. The jitter measured in this configuration is approximately 600 ps.

Duty-Cycle Distortion Jitter Characteristics

In most cases duty-cycle distortion (DCD) is caused by the components used to make a link, rather than the data sent




Figure 43. Eye Diagram




Figure 44. DDJ Measurement

across the link. It manifests itself as either differences in the rise and fall times or differences in period for bits sent as a 0 compared to bits sent as a 1. This is measured by sending a pattern down a communications link that does not exhibit DDJ and using an averaging mode on the oscilloscope to filter out any random jitter (RJ) that may be present.

The HOTLink Transmitter has a built-in DCD pattern generator that is activated by placing the transmitter in BIST mode (BISTEN LOW) while both EN_A and EN_N remain HIGH. In this mode the transmitter sends out an alternating 1-0 pattern (D10.2 or D21.5). Since all pulses in a square wave are the same, this pattern does not generate any DDJ. An example measurement of DCD for an optical link is shown in *Figure 45*.




Figure 45. DCD Measurement

When viewed from the receiver threshold (center horizontal line) in *Figure 45*, the timing for a logic 1 is seen to be slightly shorter than that of a logic 0. This difference in time is the DCD jitter present in the link.

Random Jitter Characteristics

Random jitter (RJ) is that portion of jitter that is not repetitive in nature and is caused by external or internal noise in a system (thermal noise, EMI, etc.). It is measured by using a data pattern free of DDJ (i.e., the same pattern used to measure DCD) relative to the transmitter clock. Now, averaging is turned off but infinite persistence is enabled. This captures the maximum variation of a transition relative to the clock. An example measurement of RJ for an optical link is shown in *Figure 46*.




Figure 46. RJ Measurement

In this measurement the amount of random jitter present is measured by how wide the trace is as it crosses the threshold. This particular optical link has approximately 200 ps of ran-

dom jitter present. This measurement was made using a 250-Mbit/second data pattern (4-ns/bit). *Equation 18* yields an RJ of 5% for this link example.

When making measurements of this kind, the tolerances of the signal sources and accuracy of the test equipment must also be taken into account. If the trigger source contains 50 ps of jitter, and the scope trigger accuracy is ± 50 ps, then the actual jitter present may be substantially less than that measured.

The duration of the test is also important. If the RJ is truly random and follows a Gaussian (normal) distribution, then the jitter width is unbounded. When used for bit error rate evaluations (BER), in most cases a scope can only reliably deliver random jitter measurements out to $1E-6$.

Frequency Characteristics of 8B/10B Data

Most digital design engineers are accustomed to viewing signals in the time domain using an oscilloscope. This instrument provides information about how a signal looks referenced to the passage of time. The waveforms in *Figure 47* show the HOTLink Transmitter CKW clock on the upper trace and one of the ECL data output signals on the lower trace. The individual bit cells in the output serial stream may be seen as the eye between the rising and falling output edges.




Figure 47. HOTLink Transmitter Serial Data

In the 8B/10B code, data is sent as a non-return-to-zero (NRZ) waveform. In this waveform the clock information is contained in the edges, while the data is contained in the interval between the edges. While an oscilloscope-type display allows us to see what the output looks like in terms of voltage, rise time, period, etc., it does not present any frequency-specific information. To properly design filters, couplers, or transmission systems, it is necessary to know the spectral characteristics of the signals. This information can only be examined through use of a spectrum analyzer.

A spectrum analyzer could easily be called a frequency domain oscilloscope. A conventional spectrum analyzer operates as a swept frequency, superheterodyne receiver that dis-

plays a signal's amplitude versus its frequency. It operates by sweeping a narrow-band tuned filter across a specified section of the electromagnetic spectrum, and measuring (and displaying) the rms voltage of the signal at each frequency. This swept-filter technique shows the specific frequency components that make up a complex signal, but does not provide any phase related information (Reference 22).

The spectrum analyzer plot in *Figure 48* shows the spectral characteristics of the HOTLink Transmitter serial outputs when sending the 511-character BIST pattern. The data patterns sent in the BIST loop are similar to those sent during normal communications traffic. This figure was made using a 30-MHz character-rate clock (300-MHz bit-rate data). The envelope shows a relatively even distribution of power below the bit-rate of the data, and significant amounts of energy present in the information out to 1 GHz. This illustrates how necessary it is to have a true wideband transmission system to propagate the signals.




Figure 48. BIST Pattern Spectral Characteristics

Figure 48 also shows a large dip in the energy distribution below 30 MHz. This confirms that the 8B/10B code used has no true DC component.

Figure 49 shows the spectral characteristics for the highest frequency data pattern that can be sent, a continuous 0101 (D21.5 character) pattern. With the 30-MHz character-clock used here, this pattern is equivalent to a 150-MHz square wave. Unlike *Figure 48*, most of the energy here is located at the fundamental frequency of 150 MHz, and at odd harmonics of that frequency. Other frequency components present in the signal are at least 30 dB down from the data being sent. These components are either generated by other parts of the HOTLink circuitry as it clocks, encodes, shifts, etc., the users data, or from external sources such as power-supply switching noise.

Figure 50 shows the spectral characteristics for the lowest legal frequency pattern that can be sent, a continuous 0000011111 (K28.7) pattern. This pattern ends up being an exact match in period to the source clock (30 MHz) with a fixed 50% duty cycle. Here, the largest amounts of energy are present at 30 MHz and all odd harmonics above that. The smaller frequency components present at the even harmonics are again due to the internal operation of the HOTLink Transmitter and external system noise. If this figure is compared to *Figure 49*, many of these even harmonic compo-




Figure 49. 0101 (D21.5) Pattern Spectral Characteristics




Figure 50. 0000011111 (K28.7) Pattern Spectral Characteristics

nents can be seen to have almost exactly the same level in both figures.

To verify that these spectral characteristics have some resemblance to theory, these same two source waveforms were generated mathematically and analyzed using an FFT (fast Fourier transform) algorithm. This transform analyzes a source waveform and computes its spectral components.

Because the input waveforms are not true square waves, time constant curves based on a natural logarithm were used to synthesize the rising and falling edges. These rising and falling edge equations are listed in Equations 19 and 20 respectively.

$$T_R = 1 - e^{(-t/T)} \quad \text{Eq. 19}$$

$$T_F = e^{(-t/T)} \quad \text{Eq. 20}$$

In these equations, T represents the time constant for rise and fall time. For the waveforms generated here, a T of 400 ps was used. *Figure 51* illustrates the signal generated with these equations for a 150-MHz clock rate (300-Mbit/second bit-rate). This is equivalent to the data pattern generated by a D21.5 character.




Figure 51. Synthesized D21.5 Waveform

Running a 4096 point FFT on this waveform yields the spectral components shown in *Figure 52*. The vertical axis here is plotted on a log scale to match up with the spectrum analyzer outputs. This plot illustrates that the energy of a square wave having a symmetrical rise and fall is located at the odd harmonics.




Figure 52. FFT Spectrum of Synthesized D21.5 Pattern

An FFT is based on numeric analysis rather than a physical measurement and will calculate signal components with an amplitude of zero. Because $\text{Log}(0)$ is equal to $-\infty$, a calculated FFT does not have a noise floor. To plot its results in a usable form requires the addition of an artificial noise floor to present the points of interest on a reasonable scale. To allow a better comparison, *Figures 52 and 54* use a noise floor similar to that measured in the spectrum analyzer charts.

Unlike a spectrum analyzer, which only displays the magnitude of the spectral components, an FFT of a waveform yields both magnitude and phase in rectangular form as a complex number. To plot this information for comparison with a spectrum analyzer plot requires conversion to polar notation of magnitude and phase. This calculation of the magnitude portion is done using Equation 21 (Reference 24).

$$\text{Magnitude} = \sqrt{\text{Re}^2 + \text{Im}^2} \quad \text{Eq. 21}$$

This same FFT analysis was performed on the synthesized K28.7 pattern shown in *Figure 53*. This waveform uses the same 400-ps time constant as *Figure 51*. The FFT based spectral plot for this waveform is shown in *Figure 54*. Since it

uses the same time constant, this waveform has the same rise and fall times as the D21.5 pattern in *Figure 51*. Just as in the plot for the D21.5 pattern, all of the energy is contained in the odd harmonics.




Figure 53. Synthesized K28.7 Waveform




Figure 54. FFT Spectrum of Synthesized K28.7 Pattern

The spectral plots for both the D21.5 and K28.7 synthesized patterns contain slightly more energy in the higher frequency harmonics than the actual measured signals. This is primarily due to the sharp knee present when the synthesized waveform changes between rising and falling. This knee is much rounder in the actual signal.

Components

The selection of support components for a HOTLink communications environment should not be taken lightly. The correct parts allow construction of a high-bandwidth, low error-rate system.

Several parts can be considered key in a HOTLink system. These parts are

- Clock Oscillator
- Bypass/Coupling Capacitors
- Fiber-Optic Emitters
- Fiber-Optic Detectors
- Pulse Transformers
- Fiber-Optic Cable
- Copper Cables

- Circuit Board

Clock Oscillators

The HOTLink Transmitter and Receiver are designed to operate from a very stable clock source. To achieve the necessary frequency accuracy and stability it is necessary for this clock source to be based on a quartz crystal.

The current ANSI Fibre Channel standard (ANSI X3.230-1994) calls out a frequency accuracy of ± 100 ppm for both source and destination to allow reliable communications. Clock oscillators with this initial accuracy are available from multiple sources (Reference 3).

What must also be considered is lifetime stability. Most oscillator manufacturers can easily deliver product that meets the ± 100 -ppm rating right out of the box, but this limit must be met over the life of the product, and is affected by the operating environment. The two most critical parameters are referred to as *aging* and *temperature* stability.

Aging refers to how an oscillator's output frequency varies over time (assuming other environmental factors remain constant). This is usually expressed in ppm/year. For most common "AT" cut crystals, the typical aging is 5 ppm/year for the first year and 3 ppm/year thereafter (5 ppm = 0.0005%).

A crystal's resonant frequency also varies with temperature. How much it varies is based both on how the crystal is cut, and over how wide a temperature range it is used. The stability over temperature is a non-linear function and is usually expressed as some peak-to-peak frequency change over a temperature range. The process for measuring and specifying temperature stability is called out in MIL-O-55310. Temperature stability may easily exceed the initial accuracy specification. Ratings of ± 100 ppm for temperature alone are not uncommon. *Figure 55* shows a typical transfer curve of crystal frequency vs. temperature.

This curve can be rotated on the $+25^{\circ}\text{C}$ axis point by cutting the crystal differently. This can be used to create an oscillator that is more stable over a narrow temperature range (say 0°C to $+50^{\circ}\text{C}$), yet is much more unstable outside of this range.




Figure 55. Oscillator Temperature Stability

Temperature stability and initial accuracy are often combined in a vendor's specification; i.e., ± 100 ppm at 0°C to 70°C . These numbers do not take into account the aging characteristic of stability.

Modified oscillators that allow for a wider operating environment while maintaining a high stability are available. These are referred to as either TCXO (temperature compensated

crystal oscillator) or OCXO (oven controlled crystal oscillator).

The TCXO is usually built by adding a varactor diode in series with the crystal. A special thermistor network across the diode causes the oscillator to maintain a very stable operating frequency. Because of the desired stability of a TCXO (± 2 ppm), a better grade of crystal is used to provide better aging characteristics (± 1 ppm/year). Oscillators of this type are usually larger in size (and higher in cost) than the standard 4/14-pin DIP footprint of standard clock oscillators.

The OCXO provides the highest-accuracy oscillators. These are built by placing a standard oscillator into a temperature-controlled environment. Rather than having to both heat and cool the crystal, the operating temperature is set to the upper end of the oscillator's range. Crystals are also cut such that a nearly flat area of temperature response is located at the operating temperature of the oven. The normal operating temperature of crystal ovens is in the 60°C to 100°C range.

Oven-controlled oscillators are generally quite large, expensive, and dissipate large amounts of power. They also have a significant warm-up period, requiring from 15 to 30 minutes after power-on to achieve their specified stability (Reference 23).

HOTLink Oscillator Requirements

Unlike the ANSI requirement for ± 100 -ppm stability for end-to-end communication, the HOTLink family of parts operate with a substantially wider range of reference frequencies between the Transmitter and Receiver. The specification of 0.1% end-to-end frequency tolerance allows operation with oscillator sources having up to ± 500 -ppm tolerance. This allows even the lowest-cost oscillators to be used with HOTLink.

Bypass Capacitors

At the frequencies that the HOTLink Transmitter and Receiver operate, the proper usage of power-supply bypassing becomes quite critical. Strategically sized and placed capacitors are used both to provide an AC path between V_{CC} and ground (V_{EE}), and to source current when the power supply cannot respond quickly enough due to the parasitics of the power distribution system.

The base of any power distributing system is the circuit board. Due to the very high frequencies developed in a HOTLink-based communications link, it is strongly advised to use full power and ground planes, rather than attempting to distribute power and ground on the same layers used for signal distribution. These power layers should be made with a minimum of 1-ounce copper.

To properly bypass the HOTLink Transmitter and Receiver it is necessary to know which V_{CC} pins are assigned to which portions of the logic inside the part.

HOTLink Transmitter Power Pins

The pin configuration for the HOTLink Transmitter is shown in *Figure 56*. The transmitter has three pins assigned as V_{CC} and two assigned as ground. All three of these V_{CC} power pins are connected internally and *must* be connected externally to the same power rail. If different external V_{CC} supplies were used, the current flow from the slight voltage variations that would exist could damage the part.




Figure 56. CY7B923 HOTLink Transmitter Pin Configuration

Pin 4 of the HOTLink Transmitter is named V_{CCN} or Noisy V_{CC} . This pin provides power to the ECL emitter-follower output transistors. This pin is not usually a noise source if the ECL outputs are loaded in a balanced fashion. If these same outputs are operated single-ended with unbalanced loads, then a varying amount of current will flow through this pin as the outputs switch. To keep board noise to a minimum it is advised that, if an output is used, both outputs of the differential driver have similar loads.

Pin 9 of the transmitter is named V_{CCQ} or Quiet V_{CC} . This pin provides power to the CMOS logic core of the part and the TTL compatible input buffers. This includes the 8B/10B encoder and the counters and state machines used to control the flow of data through the part. Because the dynamic current draw through this pin should not be very large, the primary bypassing concern should be for higher frequency signal components present in the internal logic.

Pin 22 of the transmitter is also named V_{CCQ} or Quiet V_{CC} . This pin is probably the most critical of all the pins on the transmitter since it provides power to the analog core. This includes the charge pumps and comparators used with the PLL clock multiplier.

HOTLink Receiver Power Pins

The pin configuration for the HOTLink Receiver is shown in Figure 57. The receiver has three pins assigned as V_{CC} and three assigned as ground. All three of these power pins are connected internally and *must* be connected externally to the same power rail. The current flow from the slight voltage variations that would exist if different external V_{CC} supplies were used could damage the part.

Pin 9 of the HOTLink Receiver is named V_{CCN} or Noisy V_{CC} . This pin provides power to the TTL-compatible output buffers. Because there is no way to maintain a constant current load on these outputs (as can be done with the HOTLink Transmitter ECL outputs) there will always be significant dynamic current flow through this pin as the part operates.

Pin 21 of the receiver is named V_{CCQ} or Quiet V_{CC} . This pin provides power to the core CMOS logic in the receiver. This includes the 10B/8B decoder and the counters and state machines used to control the flow of data through the part. Be-

cause the dynamic current draw through this pin should not be very large, the primary bypassing concern should be for higher frequency signal components present in the internal logic.

Pin 24 of the receiver is also named V_{CCQ} or Quiet V_{CC} . This pin is probably the most critical of all the pins on the receiver since it provides power to the analog core. This includes the charge pumps and comparators used with the PLL and the input differential amplifiers for the high-speed serial data streams.




Figure 57. CY7B933 HOTLink Receiver Pin Configuration

Bypass Capacitor Types

For the purposes of power supply bypassing, capacitors are used to store charge, and deliver that charge to a nearby device when necessary. While many still believe that charge is stored on the plates of a capacitor, it is not. Charge is stored in the dielectric (Reference 21).

There are two primary types of chip capacitors used for power supply bypassing; they are identified as either High-K or Low-K capacitors. These capacitor types differ primarily in their dielectric material.

The K referred to here is the dielectric constant for the material used as a dielectric in the capacitor. High-K dielectrics for bypass-type capacitors are usually based on titanates of barium, calcium, strontium or magnesium. This material provides dielectric constants in the range of 1200 to 12,000. These High-K dielectrics allow construction of physically small capacitors that provide a large amount of capacitance per unit area. The generally available range of High-K capacitors is from 200 pF to 0.22 μ F. These High-K capacitors have temperature characteristics of type X7R, Z5U, or Y5V.

Both High-K and Low-K dielectrics are used for power supply bypassing. High-K dielectrics are usually not used for temperature-critical or high-frequency operations because of their thermal and frequency dependent characteristics.

One of the biggest problems with using these High-K dielectric capacitors is sensitivity to temperature. Per the graphs in Figures 58 and 59, these types of parts can change their capacitance values by over 80% over the operating temperature range of most commercial or industrial applications. (The

temperature characteristics for Y5V are similar to Z5U except that the peak capacitance occurs around 20°C lower in temperature.)




Figure 58. Capacitance vs. Temperature for Y5V and Z5U Dielectrics




Figure 59. Capacitance vs. Temperature for X7R Dielectrics

A second problem is that these titanate-based dielectrics exhibit ferroelectric properties; i.e., they do not respond linearly to an AC signal. The effect is similar to a hysteresis loop in magnetics. This makes these dielectrics a poor choice when a distortion-free analog response is required.

When used for high-frequency (RF) or communications-link type applications, High-K dielectrics have other drawbacks. Capacitors based on these dielectric types are also very sensitive to operating voltage and frequency.

Figure 60 shows the voltage sensitivity of High-K dielectrics. Here the capacitance loss can exceed 70% with as little as 25V applied to the part. This parameter may become critical if the capacitors are used as part of a DC-block in a communications link.

Figure 61 shows one of the effects of operating frequency on capacitance. As the operating frequency increases, the

High-K dielectrics exhibit less and less capacitance. If these High-K dielectrics are to be used at an RF frequency, a capacitance correction factor must be applied to determine the actual capacitance present in the circuit (Reference 15).




Figure 60. Capacitance vs. DC Voltage




Figure 61. Capacitance vs. Frequency

Low-K dielectrics are generally based on either titanium-dioxide ceramic, alumina, or porcelain. These materials provide dielectric constants in the range of 9 to 30. Because of the Low-K, these materials are only used for making small-valued capacitors in the range of 1 pF to 10,000 pF. These Low-K capacitors are usually identified as having NP0 or C0G type temperature characteristics, and are often referred to as RF-grade capacitors because of their high Q and low dissipation factors.

Low-K dielectric capacitors are very stable over temperature. Per Figure 62, these parts change in capacitance less than 0.5% over the full military temperature range of -55°C to 125°C. Because of this temperature stability, Low-K capacitors are preferred for many analog applications where fixed time constants and resonant frequencies are necessary.




Figure 62. Capacitance vs. Temperature for NP0/C0G Dielectrics

No capacitor provides a *pure* capacitance; i.e., there are other parasitic resistive and inductive elements present in the complex impedance of a capacitor over frequency, as shown in *Figure 63* (References 15, 16, 17, 21). These parasitic elements of a capacitor are due to the materials used in, and mechanical construction of, the physical capacitor. Because of these parasitics, a capacitor cannot be treated as having ever-decreasing impedance with increasing frequency. At some frequency the capacitor passes through its series resonant point and must then be treated as an inductor.




Figure 63. Capacitor Equivalent Model

A general rule of thumb is that as the capacitance decreases, the series resonant frequency increases. This relationship is shown in *Figure 64*. At this series resonant point, the capacitive and inductive reactance components cancel each other out, leaving only the Effective Series Resistance (ESR). For most common bypass capacitors, the ESR is well under 1Ω. When selecting parts for high-frequency operation, the smaller case sizes (0805 or 0603) are preferred because they have smaller inductive parasitics.

Resistors

Figure 65 shows a first order model of a real world resistor. Because of the parasitic L and C present, a resistor does not have a constant impedance over frequency. The actual amount of change in impedance from a pure resistance is based primarily on the construction of, materials in, and DC resistance value of the component.




Figure 64. Capacitor Impedance vs. Frequency




Figure 65. First Order Resistor Model

For high-frequency or RF designs, most low-value (<1 kΩ) composite (non wire-wound) resistors may be assumed to operate at or near their DC resistance. As the DC resistance of the part increases, its impedance at higher frequencies decreases. *Figure 66* shows this relationship for typical carbon film resistors. This change in impedance is referred to as the Boella Effect and is caused by the distributed shunt capacitance present in the conducting carbon particles (Reference 21).




Figure 66. Carbon Film Resistor Frequency Characteristics

This shows that low-value carbon film resistors have reasonable impedance characteristics for RF applications, but for higher values a different type of resistor must be used.

For higher resistance values at RF frequencies, metal film resistors should be used. Because these types of resistors are not formed from particulate material, the distributed capacitance is reduced. These types of resistors are manufactured by vacuum sputtering of thin films of mixed metals onto

a ceramic substrate. Because there are no individual particles of metal, the capacitance is much lower.

Care must also be used when selecting metal film resistors since some of these have significant inductive parasitics. These inductive parasitics are often caused by the method of laser trim used to adjust the value of the resistor. Those resistors created using two straight cuts, one from either side, are generally more inductive than those trimmed using a single straight or L-shaped cut.

Metal film resistors should be used for resistors in the analog data path. This includes the transmission line termination and line bias resistors at both the source and destination ends of the serial link.

Fiber-Optic Emitters (Drivers)

A fiber-optic emitter is an electro-optical converter that changes an electrical stimulus into light. A simplified block diagram of a fiber-optic emitter is shown in *Figure 67*. The input buffer is an ECL differential line receiver. While some emitters do provide a V_{BB} output to allow single-ended operation, its use is strongly discouraged. The ECL receiver controls a high-current amplifier. The amplifier drives its current through an LED or semiconductor laser to generate a shaped optical output in response to the ECL signal input. A micro-lens assembly (usually a small sphere of glass) is used to couple and direct the light into a port for an optical fiber. Because of the small core size of the optical fiber, the lens and fiber receptacle are aligned by the fiber-optic emitter manufacturer (Reference 26).




Figure 67. Fiber-Optic Emitter Module Block Diagram

Fiber-optic emitters are available in many different case styles, wavelengths, launch modes, data rates, etc. When selecting an emitter, the main concerns are:

- Optical Receiver characteristics
- Operating data rate
- Cable plant characteristics

Most of these areas deal with interoperability of data communications links. If a shortwave laser is used as an emitter, the optical receiver must be designed to operate with the specific data rates and spectral properties of that shortwave laser. While it would be nice if a more mix-and-match combination of LED, shortwave laser, and longwave laser emitters could be used, existing receivers do not allow this. If a 1300-nm LED-driver is used, an optical receiver designed for 1300-nm LED reception must be used to properly detect the signals. In addition the optical receiver must be designed to support the data rate used in the link.

Optical emitter assemblies are available from multiple sources, including AMP/Lytel, Siemens Optical, Hewlett-Packard, Sumitomo Electric, Methode Electronics, and others.

ANSI Fibre Channel Requirements

The current ANSI Fibre Channel standard calls out four optical interface technology options for use at the 25-MByte/second data rate supported by HOTLink. The ANSI designators for these technology options are (Reference 3)

- 25-SM-LL-L
- 25-SM-LL-I
- 25-M5-SL-I
- 25-M6-LE-I

These designators are interpreted as four fields. The first field identifies the data rate used (25 MBytes/second).

The second field identifies the media used. SM specifies single-mode fiber, M5 specifies 50- μ m core multimode fiber, and M6 specifies 62.5- μ m core multimode fiber.

The third field identifies the transmitter type. LL specifies a 1300-nm longwave laser, SL specifies a 780-nm shortwave laser, and LE specifies a 1300-nm LED-driver.

The last field identifies the distance class of the link. L specifies long distance (2m–10 km), and I specifies intermediate distance (2m–1.5 km).

HOTLink correctly operates with all of these link types. However, it is up to the user to select the proper combination of emitter and detector for each class.

For those users intending to implement laser-based optical links, there are a number of federal and international safety certifications required before any such link can be put into public use. These safety requirements (ANSI Z136.1 and Z136.2, F.D.A. regulation 21 CFR subchapter J, and IEC 825) are called out in the ANSI Fibre Channel standard (References 9, 10, 11, 12, 13). No such certification requirements are necessary for LED based links.

Power Distribution Requirements for Optical Drivers

The LED or laser used to drive the optical link is probably the largest noise generating item in an optical link. When the optical driver is turned on (sending 1s), currents of up to 50 mA are forced through the LED or laser. While current steering is often used to minimize dynamic current requirements, significant high-frequency noise is still generated. Most optical modules attempt to remedy part of this situation by providing multiple V_{CC} and V_{EE} pins on their package and including some power supply bypass capacitance inside the optical module. This does take care of some of the problem, but does not correct all of it.

While bypass capacitors are still necessary to provide dynamic current, additional power isolation and filtering is required to separate the high noise of the optical transmitter from the highly sensitive optical receiver, and from the serializer/deserializer operations of the HOTLink Transmitter and Receiver. Vendor's recommendations for this include a 10- μ F solid Tantalum capacitor located near the optical transmitter, and a 0.1- μ F decoupling capacitor directly connected to the optical transmitter V_{CC} pins (Reference 26).

Isolation is provided by separating the V_{CC} or power plane for the transmitter from the rest of the surrounding power plane, through an inductive path. This is done by placing a gap in the

V_{CC} plane around most of the transmitter V_{CC} pins with a single limited connecting point. If the transmitter package only has one or two V_{CC} pins, these may be treated individually by providing power through a small inductor or surface trace. For a low-noise environment this inductor may be constructed as part of the circuit board using a 15-mil-wide trace approximately 10 mm in length (approximately 5 nH). The specified bypassing should occur after this inductive trace, right next to the optical transmitter. The net result is to implement a π -filter using the circuit board and capacitors for the different filter elements. This is shown schematically in *Figure 68*.




Figure 68. Optical Module Power π -Filter

An example slotted power plane used to implement the inductive element in the π -filter is shown in *Figure 69*. This illustration details an actual power plane layout for an optical module. The black areas indicate the absence of copper. The slot in the center of the figure is used to separate the power for the optical transmitter from the optical receiver. The shaded line on the right hand side indicates a surface layer trace (inductor) used to separate power for the optical module from the remainder of the design.




Figure 69. Fiber-Optic Module Slotted Power Plane

Fiber-Optic Detectors (Receivers)

A fiber-optic detector is an opto-electric converter that changes a light stimulus into an electrical signal. A simplified block diagram of a fiber-optic detector is shown in *Figure 70*. Light enters the module through an optical fiber and is guided by the connector housing. A coupling lens focuses the delivered optical energy onto the active region of a light sensitive diode. The presence or absence of light affects the amount of current flow through the diode. This small current flow is then amplified by a transimpedance amplifier, which then feeds an

ECL differential driver. Many fiber-optic detectors also contain additional circuitry such as signal-detect (Reference 26).

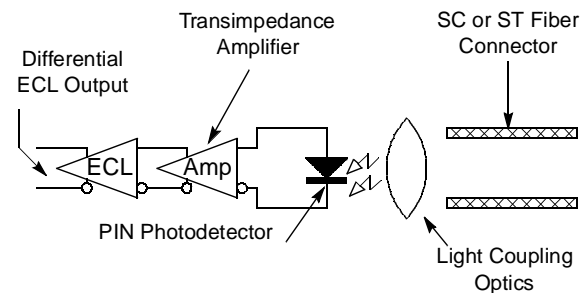


Figure 70. Fiber-Optic Detector Module Block Diagram

Fiber-optic receivers are generally available from the same vendors as fiber-optic emitters. As with fiber-optic emitters, the optical receiver must match the characteristics of the light driven into the optical fiber.

Unlike the optical emitter where there are multiple technologies used for light generation, all optical receivers are based on the response of a PIN (positive-intrinsic-negative) photo diode. These photodiodes are based on either silicon or gallium arsenide technology. The output of the PIN photodiode is a small (<1 μ A) change in current in response to received light. A fiber-optic detector module feeds the output of this PIN photodiode into a transimpedance amplifier. The function of this amplifier is to convert this small change in current into a large change (ECL 100K-level) in voltage.

For many optical receivers, it is possible to operate them above their stated maximum data rate. What is given up is receiver sensitivity and error rate margin; i.e., many 200-Mbit/second optical modules will operate at the ANSI Fibre Channel data 266-Mbit/second data rate, but with a 3-dB or greater loss of sensitivity. This loss may be converted directly into a shorter usable distance on the fiber-optic media.

Because the optical receiver has ECL outputs, care should be taken to maintain a balanced load on any differential outputs to minimize current transients. While some optical receiver outputs (i.e., signal-detect on endfire modules) may be single-ended, they usually do not change very often and should not affect data integrity when they do.

Power Distribution Requirements for Optical Receivers

The power filtering of the optical receiver is quite critical as the transimpedance amplifier must be responding to very low current variations. This filtering problem is usually compounded by the placement of the high-noise generating optical transmitter, directly adjacent to the optical receiver.

Depending on the type of receiver, it may be implemented with one or many V_{CC} pins. For those made with a single V_{CC} connection, this pin should be isolated through a π -network or other network that implements an inductive leg to block RF on the power lead.

For those optical receiver modules that use multiple V_{CC} pins, these pins are usually kept separate internal to the module, and feed different sections of the logic. For those V_{CC} pins that supply power to the ECL output emitter-followers and the ECL differential amplifiers, all that is necessary is a good 0.1- μ F decoupling capacitor next to the V_{CC} power pins. An

inductive-based filter is recommended for the V_{CC} pin that provides power to bias the PIN photodiode and the transimpedance amplifier to limit the external noise input from the system supply.

Just as with the transmitter, this inductive filter can be implemented either as a notched or slotted power plane, or by using a surface trace to act as an inductor. When implemented in this fashion the capacitor placed at the optical receiver end of the inductor should be $0.1 \mu F$.

Optical Modules

Thanks to the efforts of a group of optical component manufacturers (AMP/Lytel, Siemens Optical, Hewlett-Packard, and Sumitomo Electric), a *de facto* standard footprint has been developed for optical modules. While originally developed for the FDDI market, optical modules with speeds suitable for Fibre Channel and ATM are also available. This footprint specifies the mechanical dimensions and signal names of two different package styles, yet allows a common board layout to accept both. The dimensions and pin numbering of this footprint are shown in *Figure 71*.




Figure 71. Standard Optical Module Footprint

The two module types supported by this footprint are called DIP and endfire. The DIP modules utilize pins 1–32, while the endfire modules only use pins 33–41 (for signals) and pins 1 and 32 for package mounting. These two mounting pins are also larger in diameter than the other pins on the package.

These optical modules (DIP and endfire) share several signals. For compatibility with both module types, only the smaller set of signals present on the endfire module type should be used. A complete listing of the signals present in the standard footprint is found in *Table 5*. The signals present on the optical module are:

- SD – Signal Detect

- TD – Transmit Data
- RD – Receive Data
- Case – Outer Case of Module
- V_{CC} – Positive Supply Voltage
- V_{EE} – Negative Supply Voltage
- V_{BB} – ECL Base Threshold Voltage

The V_{BB} and SD– signals are only present on the DIP footprint package and thus should not be used in designs that wish to support interchangeable module types.

Table 5. Standard Optical Module Pinout

DIP Pin Assignments			
Pin	Signal	Pin	Signal
1	Case	2	No Pin
3	Case	4	V_{EE}
5	V_{EE}	6	+SD
7	–SD	8	Case
9	Case	10	RD
11	+RD	12	V_{CC}
13	V_{CC}	14	V_{CC}
15	Case	16	Case
17	Case	18	Case
19	V_{CC}	20	V_{CC}
21	Case	22	+TD
23	–TD	24	Case
25	Case	26	V_{BB}
27	Case	28	Case
29	V_{EE}	30	V_{EE}
31	No Pin	32	Case
Endfire Pin Assignments			
Pin	Signal	Pin	Signal
33	V_{EE}	34	+RD
35	–RD	36	+SD
37	V_{CC}	38	V_{CC}
39	–TD	40	+TD
41	V_{EE}		

Care must be used when connecting to the pins marked Case. These pins are not specified as being isolated, tied to V_{EE} , or tied to V_{CC} . As such, each manufacturer is allowed to connect them as they wish.

Isolated Case pins may be connected either to V_{CC} or V_{EE} . Usually this connection is made to whichever power rail is identified as ground in the system. When used with the HOTLink Transmitter, these types of modules are usually operated in PECL mode with the Case pins connected to V_{EE} .

When the case is connected to the V_{CC} pins, the part is designed for operation in a standard ECL (negative-referenced) system. Modules of this type may still be used with HOTLink, but some care must be taken in how they are interfaced.

Pulse Transformers

A pulse transformer is a magnetic device used to couple electrical energy from one stage to another with minimal distortion. This coupling occurs through magnetic induction. How well this coupling occurs is based on the construction of the transformer and the materials used for the core and windings.

Core Materials

There are three basic types of core materials used for transformers: metal, powdered iron, and ferrites. Metal cores consist of pieces of low conductivity metal having some magnetic properties; usually soft iron or steel. This metal core is usually made from multiple strips or laminations of material to limit eddy currents in the core. Metal cores have a practical upper frequency limit of about 50 kHz.

Powdered iron cores use metal powder fused together with an insulating binder. Because of the smaller size of the magnetic particles, the upper frequency for powdered iron cores extends to near 1 MHz.

Ferrites are a magnetic form of ceramic. Depending on the type of ferrite and construction of the core, transformers with ferrite-based cores are available with operating frequencies of near 1 GHz. This is the core material that must be used for transformers used with HOTLink.

ANSI Fibre Channel Specifications

The current ANSI Fibre Channel standard permits coupling to copper media with either capacitors or transformers. Because of their unbalanced nature, transformers must be used when interfacing to coaxial cables. The primary benefits of transformer coupling are ground isolation, common-mode rejection, and the ability to drive both balanced and unbalanced transmission lines with the same interface (Reference 3).

Just as with optical interfaces, the ANSI standard calls out multiple copper technology options for use at the 25-MByte/second data rate supported by HOTLink. The ANSI designators for these technology options are:

- 25-LV-EL-S
- 25-TV-EL-S
- 25-MI-EL-S
- 25-TP-EL-S
- 25-TW-EL-S

These designators are interpreted as four fields. The first field identifies the data rate used (25 MBytes/second).

The second field identifies the media used. LV, TV, and MI specify various grades of 75Ω coaxial cable, and TP and TW specify different types of shielded balanced-pair media.

The third field identifies the transmitter type. The EL identifier is used for all electrical classes.

The last field identifies the distance class of the link. S specifies short distances (<100 m).

While these are the only electrical classes that ANSI supports for Fibre Channel, the HOTLink Transmitter and Receiver functions with many other impedances and distances.

The typical transformer electrical characteristics to support these interface combinations are called out in the ANSI Fibre Channel standard in Section 7, Table 10 (Reference 3).

Pulse transformers suitable for coupling HOTLink to copper based cables are available from Technitrol/Pulse Engineering, Mini-Circuits, Premier Magnetics Inc., Valor, and others.

Fiber-Optic Cables

Optical media generally falls into two categories: multimode and single-mode. The usage of each type is dictated by the spectral characteristics and launch mode of the light into the fiber.

Single-Mode Fiber

Single-mode fiber is most often used with optical drivers that are both spectrally pure (i.e., a laser) and coherent in their output (well collimated, longwave laser). Fibers of this type have a very small core section to limit the modes of propagation of the transmitted light, and an index of refraction profile designed to only allow light to remain in the core that strikes the cladding at a very low critical angle. Its main propagation of light is by refraction (bending) of light that travels down the center of the core. In addition, a small number of turns of the fiber are usually placed near the optical transmitter to act as a filter for any of the higher-order modes of propagation that may be launched into the fiber. These turns change the incidence angle of the higher-order modes between the core and the cladding of the fiber, causing light at these modes to leave the core. A diagram of a single-mode fiber is shown in Figure 72 (Reference 18).




Figure 72. Single-Mode Fiber Propagation

Single-mode fibers are available in different core diameters for use with different optical source wavelengths. The fiber type called out for single-mode propagation in the ANSI Fibre Channel standard is 125-μm fiber diameter with a 9-μm core. With this core diameter, the fiber is limited to use with 1300-nm sources (Reference 3).

Multimode Fiber

Multimode fiber is usually used with optical drivers that are not spectrally pure (i.e., LED) or not coherent in their output (i.e., shortwave lasers). The lensing system used to couple the optical driver's light output to the fiber is not designed for collimation, but to couple the maximum amount of light energy. This type of fiber allows propagation of light both by refraction and by reflection.

Two distinct classes of multimode fiber are in use today: step-index and graded-index. In a step-index fiber, the primary mode of light propagation is through total internal reflection. Light that enters the core at one end is continuously reflected at the core/cladding interface until it exits the cable at the other end. A diagram of multimode step-index fiber is shown in Figure 73.




Figure 73. Multimode Step-Index Fiber Propagation

In a graded-index fiber, light is propagated through refraction rather than reflection. The fiber core is constructed of multiple concentric layers of glass. The index of refraction in each layer is slightly different, getting lower as you move out from the center of the core. Because light travels faster in a lower index media, the higher-order propagation modes that travel the farthest arrive in phase with the low-order modes that remain near the center of the core. A diagram of a multimode graded-index fiber is shown in *Figure 74* (Reference 18).




Figure 74. Multimode Graded-Index Fiber Propagation

The step-index form of multimode fiber is not normally used for high-speed data communications because its propagation characteristics limit the usable distance of a link. The ANSI Fibre Channel standard currently only supports graded-index fibers with core diameters of 50 μm or 62.5 μm , both with a cladding diameter of 125 μm (Reference 3).

Optical Pulse Dispersion

In a step-index fiber, light that travels straight through the core covers a shorter distance and arrives at the end of the fiber before light that repeatedly bounces off the core/cladding interface. This difference in delay through the fiber causes a narrow pulse launched into the fiber to widen as it travels down the fiber. Because this pulse widening or dispersion is caused by the different modes of propagation, this phenomena is known as modal dispersion.

When used with an LED driver, an additional source of dispersion comes into play. Unlike free space where all wavelengths of light propagate at the same rate, an optical fiber propagates different wavelengths at different rates. This causes any light pulse that is not spectrally pure (i.e., all the same wavelength) to widen as it travels down the fiber. Pulse widening caused by wavelength is called chromatic dispersion.

With multimode fiber one of the main limits to usable distance is the pulse spreading caused by light dispersion within the fiber. As the transmitted 1s (pulses of light) get wider through dispersion, they interact with adjacent transmitted 0s (absence of light). The effect of dispersion is illustrated in *Figure 75* (Reference 18).

With single-mode fiber, dispersion is usually not a limiting factor. Here the amount of attenuation over distance is the main limiting factor.




Figure 75. Pulse Dispersion

ANSI Fibre Channel Optical Fiber Requirements

Fiber-optic cables are available with many different optical and mechanical characteristics. International organizations have set standards for optical cable plants to allow manufacturers to standardize on some cable types.

The standards body that created the standards used for optical cable plants is called EIA/TIA (Electronic Industry Association/Telecommunications Industry Association). The governing document for all optical fiber types is EIA/TIA 492BAAA. This includes single-mode and both core diameters of multimode fiber.

The ANSI Fibre Channel standard has also selected a common fiber-optic connector type for use with all types of optical fiber media. This connector type was developed by NTT in Japan and is known as an SC-type optical fiber connector. A diagram of a simplex SC connector is shown in *Figure 76*.

These simplex connectors may be joined together using a plastic clip to form a duplex connector. In the duplex configuration the center-line spacing of the optical fibers is 0.5 inch.

Simplex and duplex cable assemblies are available from AMP, FOCS Inc., Alcoa Fujikura Ltd., Belden, and many others.

Copper Cables

There are three primary types of copper media available for distance data transmission: shielded twisted-pair (STP), twinaxial cable, and coaxial cable. Each of these cable types has specific advantages and characteristics.




Figure 76. SC Simplex Fiber-Optic Connector

Shielded Twisted Pair

Shielded twisted-pair (STP) cables are used for many low-cost LAN installations. One of the most common is the IBM Type-1 and Type-6 cables used for IEEE 802.5 token ring networks. For use with the ANSI Fibre Channel, the standard calls out Type-1 and Type-2 150Ω STP cables as defined in EIA/TIA 568 (References 3, 11, 20).

STP cables are constructed of two insulated conductors twisted together at a specific number of twists per unit length, with an overall shield and jacket. They are available with characteristic impedances from 78Ω to 200Ω. With this type of cable the transmission remains fully differential from source to destination. The shield is only used to prevent radiation and control susceptibility. Cables of this type are effective for long distances at low data rates, and short distances for high data rates. The main limiting factor for cables of this type is their attenuation at high frequencies. In many cases, cables of this type are so poor above 50 MHz that attenuation is not even specified at these frequencies. In some vendors' data, shielded twisted-pair cables are also referred to as twinax (Reference 20).

Twinaxial Cable

Twinaxial cable is a shielded form of twin-lead. Twinaxial cables consist of two parallel insulated conductors, maintained at a fixed spacing with an overall shield. Cables of this construction are often used for television reception lead-in cable. As with STP cables, twinaxial cables maintain a fully differential transmission system from transmitter to receiver. Twinaxial cables can have lower attenuation of high frequency signals than STP cables and can be used for longer distances.

Unshielded twin-lead, while having excellent high-frequency characteristics, is not generally usable for data communications due to the radiated emissions of the cable, and the impedance changes that occur as the unshielded cable is routed near metallic objects.

Twinax cables are available in impedances from 125Ω to 300Ω and velocities of 70% to 80% (Reference 20).

Coaxial Cable

Coaxial cable is used for the longest distances. They consist of a single center conductor surrounded by a dielectric spacer, surrounded by a concentric shield. Unlike either STP or twinax, coaxial cables are an unbalanced transmission line; i.e., the signal is transmitted and received as a signal relative

to a ground or shield, rather than a signal relative to another signal.

In a coaxial cable the outer conductor acts both as part of the transmission line to propagate the signal, and as a shield to prevent radiation of the transmitted signal and susceptibility from outside signals.

Coaxial cables are available in impedances from 50Ω to 125Ω and velocities of 66% to 90%. The main element that affects the velocity of propagation is the dielectric type used between the center conductor and the shield. Solid polyethylene is a common dielectric at the 66% velocity. The fastest speeds usually resort to foamed Teflon or partial air core. *Table 6* lists some common coaxial cable types and characteristics (Reference 20).

One thing that cannot be seen from this table are the cable's attenuation characteristics versus frequency. This is one of the characteristics that determines just how far a usable signal can be sent. The cables listed in *Table 6* are plotted for attenuation in *Figure 77*.

Table 6. Common Coaxial Cable Types

RG/U Type	Belden Type	Z ₀	Nominal O.D.	V _P
RG58A/U	8259	50	0.193"	66%
RG179B/U	83264	75	0.1"	70%
RG6/U	1223A	75	0.290"	83%
RG59/U	9259	75	0.242"	78%
RG11/U	87292	75	0.348"	82%
RG62A/U	9268	93	0.242"	84%
RG63	9857	125	0.405"	84%

ANSI Fibre Channel Copper Cable Requirements

The ANSI cable plant requires copper cables with specific operating characteristics. These characteristics are called out in Section 9 and Annex F of the Fibre Channel PC-PH standard (Reference 3).

Realizing these requirements means that the cable must be made with specific construction. For coaxial cables the V_P of 70% to 82% requires a foam dielectric.

The minimum necessary shield coverage for braid is 95%. This is necessary because of the high frequencies carried by the cables. With shield coverage lower than this, the signal leakage through the braid can allow not only significant signal radiation, but an impedance mismatch due to signal propagation down the outer surface of the braid. For best effectiveness, a 100% foil shield should be used in addition to the braid shield.

To meet flammability requirements, the National Electrical Code now requires that almost all installations use either CL2 or CL2P (plenum rated) jacket material (Reference 25).

Cables meeting all of these requirements are available from multiple vendors.

The ANSI standard also allows use of shielded twisted pair or twinaxial type cables. These cables all require a shield to meet EMI/EMC requirements. Unshielded twisted pair (used for many networks) should not be used. This is primarily due to radiated emissions rather than susceptibility.




Figure 77. Coaxial Cable Attenuation Characteristics

Copper Cable Connectors

There are three primary connector types called out for use with copper cables: BNC and TNC for coaxial cables (illustrated in *Figure 78*) and a 9-pin D-sub (illustrated in *Figure 79*) for twisted-pair/twinax cables.

For coaxial cables, the BNC connectors are used on the transmit end of the cable, while the TNC connectors are used on the receiver end of the cable. This dual connector configuration allows a duplex cable to be connected without having to identify one cable from the other. With these connectors the male end is always on the cable while the female end is used at the board bulkhead.

For twisted-pair or twinaxial type cables a 9-pin D-sub connector is used. This connector is required to have a metal shell because the shields of both the transmit and receive pairs are terminated to the shell of the connector. As with the coaxial connectors the cable gets the male connector while the board or bulkhead gets the female connector.

The STP cable is wired in a crossover fashion where the transmit pins at one end of the cable (as illustrated in *Figure 80*) are connected to the receive pins at the other end of the cable. The cable shields for both pairs are tied together and connected to the D-sub connector shell at each end.




Figure 78. TNC/BNC Cable Connectors




Figure 79. STP Cable Connector and Connector Pinout




Figure 80. STP Cable Connections

Because of the low current used in these cables, the connections are considered to be dry circuits. To prevent contact

oxidation from degrading the link over time, the contacts are required to be gold or palladium plated (Reference 27).

Conclusion

The HOTLink family of communications products provide designers with a simple yet elegant method of reliably moving large quantities of data at very high speeds from one place to another. These parts are capable of communicating over copper or optical media at distances well in excess of industry standards. Their BiCMOS implementation, along with their integrated power saving features, combine to offer one of the lowest-power, high-speed serial communications link standards available.

HOTLink is a trademark of Cypress Semiconductor.
 IBM is a registered trademark of International Business Machines Corporation.
 ESCON is a trademark of International Business Machines Corporation.

References

1. A. X. Widmer and P. A. Franazek, *A DC-Balanced, Partitioned-Block, 8B/10B Transmission Code*, IBM Journal of Research and Development, 27, No. 5: 440-451, September 1983
2. U.S. Patent 4,486,739, Peter A. Franaszek and Albert X. Widmer, *Byte Oriented DC Balanced (0,4) 8B/10B Partitioned Block Transmission Code*, December 4, 1984
3. *Fibre Channel Physical Standard, ANSI X3.230-1994*, American National Standards Institute, 1994
4. *Enterprise System Architecture/390 ESCON I/O*, SA22-7202, IBM Corporation, 1990
5. *F100K ECL Logic Databook and Design Guide*, National Semiconductor, 1990
6. Lawrence B. Levit and Marco L. Vincelli, *Characterize High-speed Digital Circuits: A Job For Wideband Scopes*, Lecroy Corp., EDN June 10, 1993
7. *The ABC's of Probes*, Tektronix Pub 60W-6053-3
8. *Product Overview*, Cascade Microtech, Product Overview, 1992
9. *Safe Use of Lasers ANSI Z136.1-1993*, American National Standards Institute, 1993
10. *Laser Safety in Optical Communication Systems ANSI Z136.2*, American National Standards Institute
11. *Commercial Building Telecommunications Wiring Standard EIA/TIA-568*, Electronics Industries Association/Telecommunications Industries Association
12. F.D.A. Regulation 21, Code of Federal Regulations
13. *IEC825*, International Electrotechnical Committee
14. Scott, Paul, Cypress Semiconductor, *Draft Paper - HOT-Link On Wire*, May 1992
15. *1990-91 Resistor/Capacitor Data Book*, Philips Components, 1990
16. *System Design Considerations When Using Cypress CMOS Circuits*, Cypress Semiconductor Applications Handbook, 1993
17. White, Donald R.J., *Electrical Filters, Synthesis, Design and Applications*, Second Edition, 1980
18. Sterling, Donald J.Jr., *Technician's Guide To Fiber Optics*, Second Edition, 1993
19. Blood Jr., William R., *MECL System Design Handbook*, Fourth Edition, 1988
20. *Belden Wire and Cable*, Cooper Industries, 1990
21. Botos, Bob, Hewlett Packard, *Designers Guide to RCL Measurements*, June 1979
22. *Hewlett-Packard Test and Measurement Catalog*, Hewlett-Packard Corp., 1992
23. *Crystal Oscillator Handbook and Catalog*, Vectron Laboratories, Inc., 1992
24. Ramierez, Robert W., *The FFT, Fundamentals and Concepts*, Tektronix, Inc. 1985
25. *National Electrical Code*, National Fire Protection Association
26. *Application Note 65074, Fiber-Optic Transmitter and Receiver*, AMP Incorporated, 1992
27. J. H. Whitley, AMP Inc., *Contacts and Dry Circuits*, AMP Symposium paper, October 1963