



HOTLink™ Jitter Characteristics

Abstract

This application note describes the basics of jitter in transmission systems and, using HOTLink™ as the example, shows how it can be analyzed and measured. Specific characterization data is presented to allow system integrators to understand the parameters needed to improve the reliability of their systems.

Introduction

This application note examines jitter from three different perspectives. First, as a background overview, it describes a few basic "jitter" concepts that affect digital systems. Second, it describes the jitter performance and characterization of the HOTLink Transmitter (CY7B923). Third, it describes the jitter tolerance and feed-through characteristics of the HOTLink Receiver (CY7B933).

Numerical characterization data is supported by descriptions of the various testing techniques and equipment that are required to obtain this information. Commercial, custom, and "home-brew" test equipment are described along with the connections used to gather data that illustrates the levels of performance attainable by HOTLink products.

The data contained in this application note will help users to understand the various characteristics of link components and HOTLink characteristics and capabilities. This data is offered to assist in the design of robust serial interconnect links.

Jitter

Jitter is a high-frequency semi-random displacement of a signal from its ideal location. These displacements can occur in amplitude, phase, and pulse width, and are generally categorized as either deterministic or random. For data communications links based on (or similar to) HOTLink, measurement and specification of jitter is usually restricted to timing displacements.

Deterministic jitter consists of those timing variations that are repeatable within a system and whose cause can generally be directly attributable to specific physical components or events. An example of this would be the jitter caused by the frequency-selective attenuation and phase delay of a signal in a transmission line.

Random jitter deals with those timing variations that are much more probabilistic in nature. While still observable and measurable in a system, this jitter is not directly predictable. Common sources for random jitter are thermal and electrical noise, both internal to and injected into a system or component.

Jitter in logic circuits is often characterized by its transfer function. This function, known as jitter feedthrough, is a measure of jitter output relative to jitter input of a system or component. Most circuits, when presented with jitter, tend to amplify that jitter in a few or many areas. Fortunately for data communications systems (which are plagued by high jitter creation elements), application of properly designed PLLs (phase-locked loops) can actually reduce or remove large amounts of jitter from a clock or data stream.

Background—Jitter in Logic Systems

The timing of logic signals flowing through a logic system is often assumed to be a series of simple voltage transitions that occur after some fixed delay. While this is a convenient and usually sufficient assumption for the logical function of a device, it is insufficient to analyze the limits of the timing or the reliability of the design.

The delay through logic devices (i.e., gates, flip-flops and other common building blocks) is defined to a first order by the time it takes for the inputs, the internal circuit nodes, and the outputs to change from one voltage to another. Since there is always some uncertainty about the exact voltage present at any node in the circuit, various logic families have been de-




Figure 1. Link Jitter Budget Depends on Link Components

vised with specific ways to assure reliable logic functions. Thresholds are well defined and inter-gate links have sufficient voltage margins to assure reliability. Typical components have output levels (e.g., V_{OH} , V_{OL} , etc.) that assure a significant voltage margin above and below the input thresholds (e.g., V_{IH} , V_{IL} , V_{TH} , etc.).

Most logic model libraries document a fairly wide range of possible delays through a logic element. This range includes the effects of many internal characteristics such as differences in output resting voltage, threshold voltage, signal ramp rates, and (to some extent) the speed the signals travel along the interconnecting wire, metalization, and leadframes. These delays, while supposedly covering the minimum to maximum range for the part, assume specific external operating and signal conditions. By presenting the logic element with input, output, or power conditions beyond those assumptions, it is possible for these logic elements to exhibit apparent delays both faster and slower than the specified minimum and maximum.

The noise carried on the V_{CC} or Ground rails (both internal and external) affects the actual timing of the I/O transition by causing changes in the starting levels of the active transition. *Figure 2* shows only the timing variation caused by ground bounce, but the influence of V_{CC} noise has a similar effect. If the signal begins its transition at some arbitrary but fixed time, and has a transition rate (i.e., rise time or fall time) that is mostly controlled by slew-rate limiting effects not related to the power supply glitch, the effective timing is determined by the placement of the glitch. If the transition begins on a glitch-peak, it arrives at the threshold voltage a little early. If the transition starts in a glitch-valley, it arrives a little late. This change in timing is usually invisible to the external examiner (except as power supply induced timing variation) because much, if not all, of the glitch is contained within the IC package, and is not externally observable.

The effect of this variation in starting voltage can cause significant variations in timing. A signal that has a 1 V/ns ramp rate (TTL edges are usually between 0.5–1 V/ns, and can be much slower), has an effective change in delay of about 1 picosecond per millivolt of disturbance. This equates to ± 100 picoseconds of delay variation for ± 100 millivolts of ground or V_{CC} noise, an amplitude which is normally deemed “quiet.” When noise spikes approach 1 volt, delay variations




Figure 2. Power Supply Glitches Affect I/O Actual Timing

can exceed 1 nanosecond. With 1 volt of power supply variation, other delay effects are sure to appear.

Additional timing variation can be caused by noise coupled into the external or internal logic through cross-coupled logic paths (including package-pin crosstalk), or by power supply noise injection. These “minor” variations in delay are typically ignored in the analysis of the logical function, since there is sufficient overdrive (voltage noise margin) to assure that the logical function is achieved. However this assurance is not transferred to the timing margins of a logic design.

Most of the delay of today's high performance logic is caused by an output “ramping” from its resting voltage to the actual threshold voltage (the voltage at which the gate begins to make its logical decision and subsequently change its own outputs). Any disturbance in either the internal threshold or the ramping input or output will cause a change in the apparent delay through the gate (see *Figure 3*). All single-ended logic gates suffer from this variable-delay characteristic.

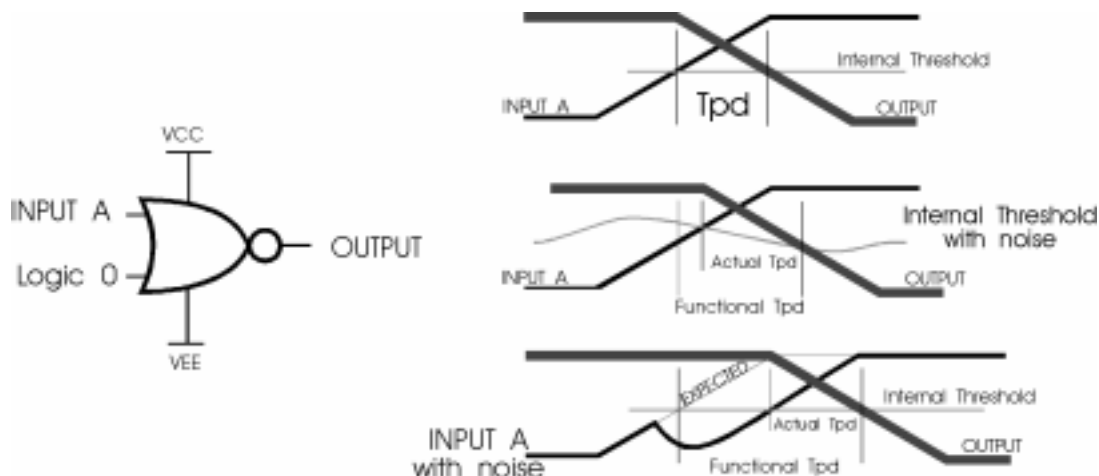


Figure 3. Delay Through a Logic Gate Changes with Injected Noise

Single-ended circuits include all TTL, CMOS, and any ECL logic that uses an internal or external threshold reference.

Differential circuitry can be used to partially mitigate the effects of injected noise, since the threshold of the gate is determined by a complementary input, hopefully carrying the same injected noise, but ramping in the opposite direction. The common mode range of such a differential gate helps to reduce many noise-induced delay characteristics. All of the critical timing paths in HOTLink products are implemented with differential CML (Current-Mode Logic) signals to minimize crosstalk and V_{CC} -coupled noise-jitter effects.

Various design techniques have been developed that maximize timing margins in logic, but in most of these techniques the timing of any particular logic element is considered a constant (or a range of constants). Except for the well known metastability characteristic of storage elements, the design tools assume that each element has a fixed delay, and the only accommodation to metastability is to attempt to avoid the conditions that provoke the unpredictable behavior.

Traditional design practices work on the simple assumption that if the logic path (delay) between storage elements is less than the time between clocking edges by some comfortable margin, then the logic will behave exactly as the designer intended. As clock speeds and product complexity increase, this comfortable simplifying-fantasy becomes more difficult to maintain. As is well documented in other literature, if the transition on the DATA input changes later than the required set-up time prior to the active transition on the CLOCK input, the delay of FF-1 (Figure 4) may increase or it may refuse to store the expected data. If the path length to FF-2 is running near its maximum limit, this increased delay could propagate through the logic causing unexpected and undesirable results.

Designs that meet all manufacturers specified set-up and hold times can also experience variable delays through the flip-flop. As the input transition approaches the "actual" set-up time of the internal latches, their clock-to-out delay begins to change (Figure 5). Typically, T_{setup} is specified at the point where delay has changed by less than some arbitrary amount (usually about 10%) of the cell's "nominal" delay. Inside of that point, delay increases radically until the flip-flop goes metastable. Similar effects occur as the hold time approaches zero.

Even if the nominal delays of the intervening logic are within design margins, voltage-noise effects can change the delays of the combinational logic devices. If that happens, metastable effects may be observed in the system. Normally in digital-logic systems, great care is taken to assure adequate




Figure 5. Propagation Delay Changes as Actual T_{setup} is Approached

timing margin, whereupon the error rate is "assumed to be zero," and ignored.

Jitter in PLL Systems

Phase-Locked Loops are typically used in high-speed clock multipliers and precision clock-recovery circuits. In their role as clock-source generators, PLLs are characterized for their timing precision. This is usually because any jitter that appears on the clock line must be compensated by an equivalent reduction in the timing margin allowed between flip-flops.

Jitter can enter a multiplier PLL in several ways (see Figure 6). The Clock input (1) can contain voltage-coupled noise or phase-noise that can affect the multiplied Bit Clock. The UP and DOWN outputs (2) of the PHASE FREQ DET are the digital-to-analog interface with the analog control circuits of the PLL and can suffer from the same voltage-coupled noise effects described earlier for logic. These digital signals carry the picosecond analog-timing information that controls the VCO. Any crosstalk or noise injection at this point will corrupt the "error" information that the PLL uses to maintain phase-lock with the input clock.

The output of the analog filter (3) contains both the gross center-frequency control, and the precision phase-control. Typically the input sensitivity of the VCO is hundreds of megahertz per volt, and microvolts of crosstalk or power-supply




Figure 4. Typical Logic Path Delay Limited by Minimum Clock Period

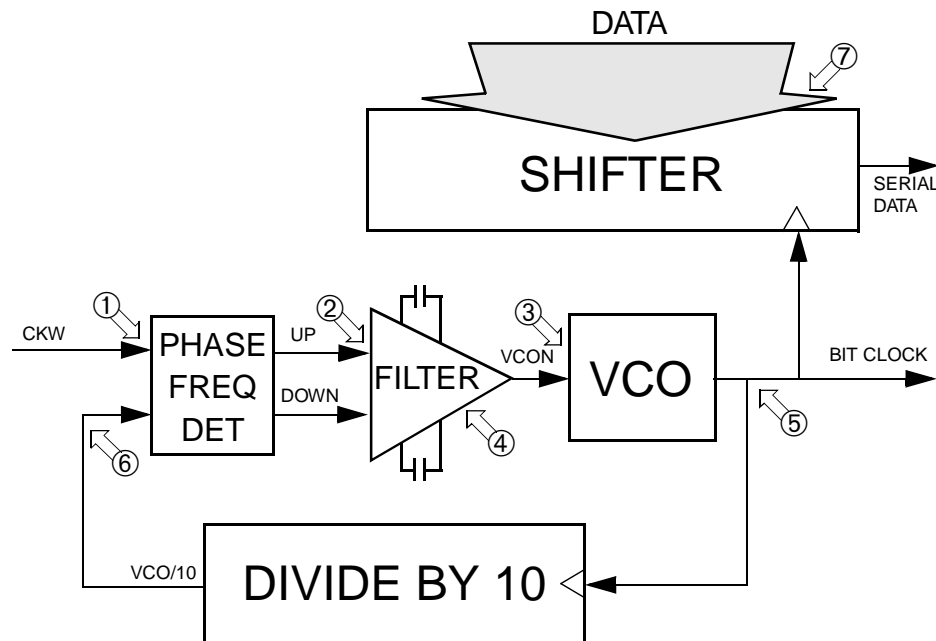


Figure 6. Clock Multiplier PLL Noise Injection Points

noise injection can add nanoseconds of jitter to the PLL output. Similarly, the capacitors (4) used in the Filter (either internal or external) can be susceptible to noise injection which cannot be eliminated by traditional circuit techniques.

HOTLink products use carefully designed, fully internal, Metal/Oxide/Silicon (MOS) capacitors. These huge, matched devices minimize external noise coupling. For noise sources that cannot be avoided, the capacitors and all of the other analog circuitry are designed to make coupled noise more rejectable by using fully differential, common-mode noise reduction methods. Older PLLs often use external capacitors which are notorious for noise injection through the external pins and circuit board traces required to connect these capacitors.

Noise injected at (2), (3), or (4), and to a lesser extent at the other points, can be only partially compensated by the normal filtering actions of the PLL. Noise at (2) or (6) exhibits different effects than noise injected at (1) and affects the Bit Clock (5) in different ways. These differences are illustrated in *Figure 11* and are discussed later.

Since the multiplier PLL only receives its correction information once every N VCO cycles (where N is the multiplication factor of the PLL, the VCO frequency divided by ten in this case), many specific errors do not cause a correction. Only the “average” of noise-induced errors result in compensatable disturbances. “Instantaneous” errors are not compensated by the PLL at all, especially if there are other errors of similar magnitude and opposite sign between reference updates.

Logic noise, as described earlier, can be injected into the Recovered Bit Clock (5) or at the feedback reference (6). These can be avoided by careful differential circuit and logic design. The parallel data moved through the Shifter (7) can cause transmitted output jitter, which is a function of the data pattern being sent (i.e., DDJ) as the set-up and hold times of the output flip-flop vary.

The operation of the PLL can cause jitter just by its normal operation (*Figure 7*). Whenever the phase detector adjusts the frequency of the VCO, it causes an instantaneous change in phase as part of the adjustment operation. This instantaneous phase change, followed by a drift until the time of the next correction, is the normal operation of the loop. Ideally, the correction is small, and entirely contained within one clock cycle, but if it is larger or lasts longer than one cycle of the VCO, it can cause bit-to-bit phase differences (i.e., jitter).

Clock Recovery, Data Separator PLL

The PLL used for clock synchronization and data recovery (shown in *Figure 8*) is different from the one used as a clock multiplier (shown in *Figure 6*). The phase correction information comes from comparisons between an arbitrary input pulse stream and an internal bit-rate VCO. In contrast to the Phase-Frequency Detector (PFD) used in the clock multiplier, this PLL uses a detector that is sensitive only to phase errors. Missing data transitions are ignored, and corrections follow each and every data transition. In contrast to the predictable correction rate of the PFD, the Phase Detector makes corrections at the rate of the incoming data. It can vary from one correction per VCO cycle (when data contains alternating 10101...) to once per byte (or less) for some serial protocols. This variation in correction density can cause some forms of jitter and, by affecting the loop stability and bandwidth characteristics, affects jitter feed through.

Jitter can enter a synchronizing PLL in several ways. The input data (1) contains significant jitter which accumulates on the serial transmission link. This is the jitter that the receive PLL is intended to remove.

The noise injection points at (2), (3), (4), and (5) are the same as those in the multiplier PLL, and affect the receiver PLL in similar ways. The main difference is that this PLL gets a phase-error update on each input data transition. This allows noise events to be corrected more often than those in the




Figure 7. Phase/Frequency Corrections in Multiplier PLL

multiplier PLL, but the noise induced corrections can be affected by the corrections already required by the jittered data. Conversely, these noise-induced jitter components reduce the data-recovery circuit's tolerance to input-data jitter.

The Phase Detector (or PFD in clock multiplier PLLs) in clock synchronizer PLLs is intended to give a "unit" of phase correction information for a "unit" of error. This correction should be directly proportional to the error, regardless of error magnitude. A poorly designed (or poorly implemented) phase detector in a PLL (either a multiplier or clock synchronizer loop) can exhibit what is typically called a "dead-zone" if the error/correction relationship does not hold for miniscule errors. This effect is shown in *Figure 9* as the less-than-ideal transfer function which effectively removes the phase correction con-

trol in the neighborhood of "zero error." This "hole" in the transfer function causes an otherwise perfectly locked loop to exhibit jitter because the loop is unable to maintain linear control and wanders between the two inflection points. HOTLink Transmitter and Receiver PLLs have been designed to eliminate this undesirable behavior.

The closed-loop PLL acts like a Low-Pass Filter to incoming noise (*Figure 10*). All frequency components that fall below the roll-off frequency of this filter are passed unattenuated. Frequencies above the roll-off frequency of the filter are attenuated, and frequencies around this point might be amplified to some extent. Some forms of jitter have low-frequency characteristics that pass through the PLL and appear on the re-

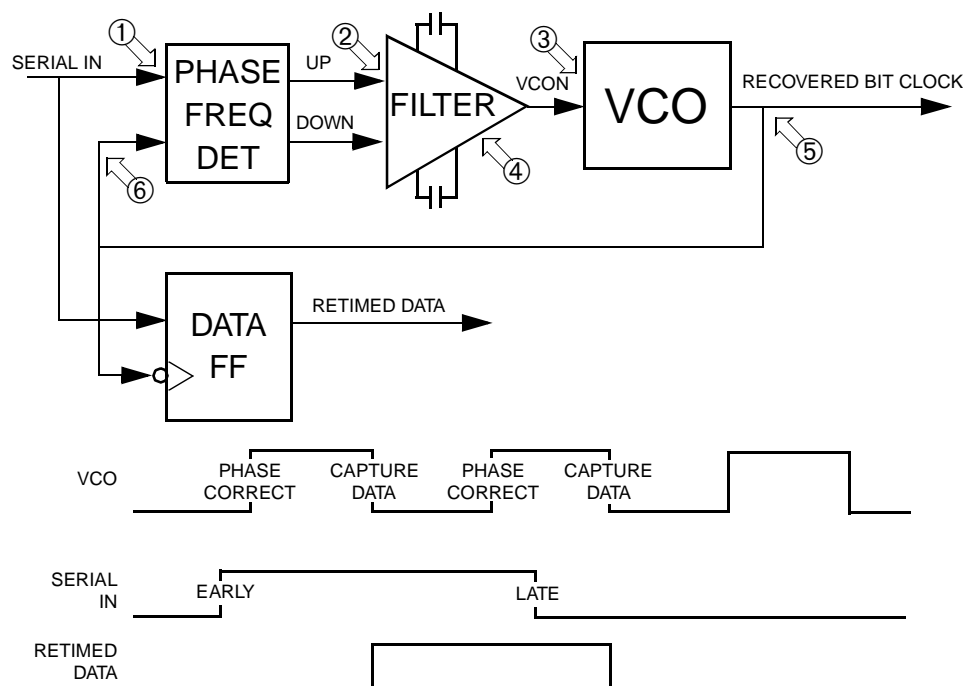


Figure 8. Receive-PLL Block Diagram




Figure 9. Phase Corrections Should be Linear with Error Magnitude

sulting high-frequency clock output (e.g., low-frequency wander passes unattenuated through the Receive PLL).

The PLL low-pass filter model is valid for jitter that enters the system at the PLL input. However, jitter that is injected (or is present) *inside* the loop “sees” the loop as a high-pass filter. The dynamics of the closed-loop system allow it to compensate for low-frequency injected jitter with an automatic (and opposite) low-frequency phase adjustment. As the frequency of the injected jitter rises toward the roll-off frequency, the loop becomes incapable of fully compensating for the injected jitter. Above the roll-off frequency, the loop passes injected jitter without attenuation (see Figure 11).

Since V_{CC} noise is injected (and can result in undesirable delay variations) at multiple points inside the loop, the resulting component of jitter attributable to V_{CC} noise usually shows a peak at approximately the roll-off frequency, and less output

jitter at both higher and lower frequencies. The ratio of jitter magnitude for V_{CC} noise frequencies above and below the PLL roll-off frequency might give a clue to the probable noise injection point(s). Larger jitter magnitude for V_{CC} noise frequencies below roll-off is probably attributable to logic delays in the input path of the PLL (e.g., logic delay changes in gates prior to the loop, passing through the PLL low-pass filter). Larger jitter magnitude for V_{CC} noise frequencies above roll-off is probably attributable to noise injection into some circuit inside the loop (e.g., jitter inside the loop being affected by the PLL as a high-pass filter). Jitter magnitudes that are approximately the same above and below roll-off are probably attributable to delay changes being introduced into circuits that follow the loop.

Jitter on a Multiplier-PLL clock can affect the operation of a system in several ways. The most obvious effect is cycle-to-cycle jitter (high-frequency or random effects) that causes a reduction in clock-to-clock spacing, and affects logic timing as was discussed previously. Another undesirable effect is the longer term clock wander that can cause problems in a system when two parts of the system are clocked by different clocks of the same frequency but with a variable phase relationship.

Jitter can affect a Clock-Synchronizer PLL in at least two ways. First, the jitter in the PLL clock directly reduces the jitter tolerance of the Clock/Data separator. Jitter tolerance of the receiver is a measure of design margin in a serial communication link. Second, the jitter of the PLL clock is transferred directly to the receiving host system clock. This jitter-reduced timing margin affects the system in exactly the same way that Multiplier PLL jitter affects logic timing.

Bit Error Rate

Jitter is often directly equated with Bit Error Rate (BER). While it is true that jitter accumulation on a serial link is the primary determinant of the link's reliability, its statistical character




Figure 10. PLL Closed Loop Response is a Low Pass Filter




Figure 11. V_{CC} Noise Injection Transfer Function

makes it difficult to understand and impossible to accurately predict. Since PLLs also exhibit a statistical nature which must be added to the jitter of the link, some designers incorrectly assume that the PLL jitter is what causes the errors. In fact, the PLLs used to create the high-frequency clocks in the transmitter and receiver serve to increase reliability rather than decrease it.

BER is a term which is common to both serial communication devices and communication systems, and could also be applied to any system where data moves from one storage location to another. BER is the ratio of “corrupted data” received to “good data” sent, and is calculated using Equation 1. Sometimes BER can mean “Byte Error Rate” in systems that transmit multi-bit-wide information. Errors that affect the fill-bits (the “non-information” bits that occupy the time between actual data transfers) are not usually counted toward BER, but might be used to predict overall system margins.

$$BER = \frac{\text{Bits in error}}{\text{Bits transmitted}} \quad \text{Eq. 1}$$

Usually BER is expressed as a large, negative exponent (e.g., 2.5×10^{-12} or 1×10^{-9}) because acceptable systems perform almost flawlessly. For example, a 250 MBaud system operating with a BER of 1×10^{-12} would only experience about one error per hour while sending continuous information (see Figure 12).

While one error per hour may seem excessive, it comes naturally from the technology of the serial interconnect used. Even when adequate margins are designed into the link, there are physical and electrical effects which can cause occasional errors. Communication systems have been engineered to accept small error rates by including some level of error checking in the physical layer hardware, and extensive error recovery is built into most of the communication protocols.

The PLL(s) used in the communication link don’t cause the failures. Their ability to recover timing information from a severely distorted serial interface is one reason that these types of links are possible at all. The PLL supplies the logical clock necessary to correctly process the data. Without it the difficul-

$$\text{Example: } \left(\text{Bits transmitted} = 250 \times 10^6 \frac{\text{bits}}{\text{sec}} \times 3600 \frac{\text{sec}}{\text{hour}} = 9 \times 10^{11} \frac{\text{bits}}{\text{hour}} \right)$$

$$BER = \frac{N}{9 \times 10^{11}} = N \times 1.11 \times 10^{-12}$$

MTBF of a link running at 250 Mbaud @ BER of $1 \times 10^{-12} = 1$ hour between errors

MTBF of a link running at 250 Mbaud @ BER of $1 \times 10^{-15} = 46$ days between errors

MTBF of a link running at 250 Mbaud @ BER of $1 \times 10^{-18} = 127$ years between errors

Figure 12. BER Example Calculations

ty and cost of aligning data and clock across any distance would be immense if not impossible. But because of the designer's inability to predict to the exact statistical nature of the interconnect link, and because PLL jitter contributes to this uncertainty, PLLs are often wrongly equated with BER.

HOTLink CY9266-C Evaluation Boards operating at 250 MBaud with a short coax link have been tested continuously for over 4000 hours without any errors. During this time, over 3.6×10^{15} bits were sent, received, and checked by the HOTLinks using the BIST (Built-In Self-Test) function. This error-free time yields an estimate of BER less than 8×10^{-16} with greater than 95% confidence. Link error rate is not impaired by the addition of a deterministic interconnect link, such as a long coaxial cable. HOTLink CY9266-C Evaluation Boards running the BIST test at 250 Mbaud, and interconnected with 300 feet of RG-59 coaxial cable (90% of the maximum uncompensated distance) have operated with no errors during a 1500 hour test. This shorter test still yields an estimated BER of less than 2×10^{-15} with greater than 95% confidence. No error has *ever* been recorded in all the time spent testing HOTLinks with deterministic links (sensitivity to environmental-noise injection was not included in this test). Actual BER rates have never been determined because after an uneventful six month test, the tests were terminated to free the equipment for other uses.

HOTLink Transmitter Jitter

The PLL used in a Transmitter application (clock multiplier) is intended to provide a high-speed, stable clock that tracks a low-speed reference (CKW in Figure 13). This clock (Bit Clock) is used to run the parallel-to-serial converter and all of the internal logic in the HOTLink Transmitter.

Jitter is an undesirable and often unpredictable misplacement of any particular transition from its ideal position. Transmitter output jitter can be characterized as Random or Deterministic Jitter, RJ and DJ respectively. It can further be subdivided into Intrinsic Jitter, Transferred Jitter, or Injected Jitter.

To separate the various types of jitter, carefully designed tests were performed on HOTLink parts selected from the full spec-

trum of manufacturing tolerances. These tests were designed to separate the effects of Power Supply, Clock Sources, and various PLL characteristics. Manufacturing tolerances include variations in all types of resistors used in the design, characteristics of Bipolar and CMOS transistors, and other normal process variations. Environmental effects include V_{CC} variation over at least the full specified range, and ambient temperature variation over the full military and commercial ranges.

Unless otherwise noted in the following text, static variations in power supply levels (4.5V to 5.5V), ambient temperature (-55°C to 125°C), and process variations (within manufacturing tolerance limits) cause virtually no change (within the accuracy of the measurement system) to any jitter tolerance or PLL characteristic. This should be true for any *well-designed* PLL, though it is often not true for all products in the marketplace.

Transmitter Random Jitter

Random Jitter is an undesirable and unpredictable misplacement of any particular transition from its ideal position that cannot be correlated to either data-stream content, or parameters of the hardware. To separate Random Jitter from the other effects, the HOTLink is configured to send various square-wave patterns. This minimizes any possible deterministic (DJ) effects caused by loading or variations in internal circuit delays. The set-up used to measure HOTLink Random Jitter is shown in Figure 14. The clock source was a combination of an HP 8656B and HP 8131 generator chosen to give the lowest possible source jitter. Other generators, described later in this application note, are also satisfactory for this clock source. The output was measured on both an HP 54720D and a Tektronix 11801 Digital Sampling Scope, each of which have sufficient sampling bandwidth to accurately show the performance characteristics of the device under test.

Random jitter, measured in this way, is shown in Figure 15. In virtual time, the "random" characteristic can be seen in these histograms taken on a Tektronix 11801 Digital Sampling Oscilloscope.

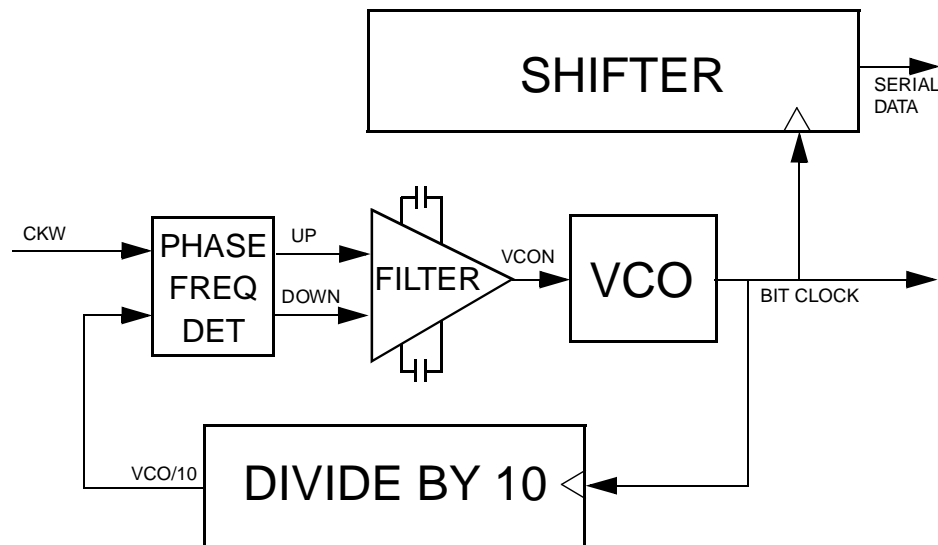


Figure 13. HOTLink Transmitter PLL Block Diagram




Figure 14. HOTLink Transmitter Random Jitter Set-Up

The left histogram in *Figure 15* shows the quality of the reference-clock signal used to gather the data that follows. The reference clock, coming from an HP8656B frequency synthesizer and buffered through an HP8131 Pulse Generator, has minimal jitter. This TTL input clock-source adds a negligible amount of jitter to the output jitter measurements that follow. Other generator/buffer combinations are possible, but if any appreciable jitter is present at this point, it will be difficult if not impossible to separate the input jitter from the jitter accumulated in the part under test. Other tests (described later) show the effect of noise at the CKW input.

Cycle-to-cycle jitter out of the HOTLink while sending a “perfect” bit-rate square wave, measured one bit away from the output edge used to trigger the scope, is very small. This is a measure of the jitter that can accumulate between adjacent VCO clocks in the multiplier PLL. Since this measurement

(right histogram in *Figure 15*) was triggered and measured on the same output, and the scope was not constrained to sample at any particular rate, this photo shows the superposition of all possible bit positions. It also shows the (small) magnitude of Deterministic Jitter built into the output circuitry.

The serial outputs of HOTLink are PECL differential signals that must be combined differentially in the front-end of the sampling scope to provide an accurate measurement of both the signal transition and any jitter present on those transitions.

Figure 16 displays the differential measurement of the accumulated jitter on the OUTA± outputs of the HOTLink Transmitter. This is a measure of the total jitter accumulation through the entire PLL and output circuit (combined with any jitter present on the reference clock) while sending a “perfect” byte-rate square wave.





Figure 15. Histograms of CKW Source and OUTA± One Cycle Away




Figure 16. HOTLink Transmitter OUTA± Rj vs. CKW

The wide vertical bar shows the accumulated jitter measured in 100,000 samples of the 0-to-1 transition while being triggered by the CKW reference. The different shades of gray in the vertical bar represent different concentrations of signal samples that occurred at that specific time/amplitude coordinate. The darker the sample point, the more samples that occurred at that point. The very center of Figure 16 contains a narrow rectangle centered around the HOTLink Receiver threshold region. All samples that occur within this rectangle are plotted in the black histogram at the bottom of the figure. The Gaussian shape of this curve confirms that the primary component of this jitter is truly random in nature. While there




Figure 17. HOTLink Transmitter Random Jitter as a Function of Frequency

is a slight increase in output jitter as the operating frequency decreases (see Figure 17), there is no appreciable change in HOTLink jitter due to V_{CC} , temperature, or process variation.

In contrast to the virtual-time measurements shown in the previous figures, real-time measurements allow an insight to the behavior of the HOTLink Transmitter in terms of sequential events. Figure 18 shows the edge displacement from the ideal location of all sequential rising edges of a continuous data stream. In this sequential, real-time measurement, it can be seen that there is no obvious or repetitive pattern to the jitter, confirming the validity of the virtual-time measurement. A minor pattern is visible in the running-average histogram, that shows a small amplitude, continuous oscillation in the sign of the edge misplacements. Peak-to-peak deviation in the real-time illustration is smaller than that indicated in the virtual-time measurement. This is consistent with the large difference in the number of samples in each, and the fact that many of the extreme excursions that occupy the tails of the distribution may not be PLL variations, but are probably caused by pulse-noise injected into some logical or measurement function.




Figure 18. Real-Time HOTLink Transmitter-Output Byte-Rate Jitter




Figure 19. Real-Time HOTLink Transmitter Bit-Rate Jitter Output

The pattern of jitter does not change appreciably when the output pattern changes from one cycle-per-byte to one cycle-per-bit (see Figure 19). The peak excursion remains about the same, and the distribution is similar.

Transmitter Deterministic Jitter

Deterministic jitter is an undesirable and often difficult-to-predict misplacement of any particular transition from its ideal position that can be correlated to the content of the data stream or some characteristic of the circuit or hardware. To measure deterministic jitter attributable to the internal circuitry of the HOTLink Transmitter, the parts were measured in several ways to separate the different possible DJ sources. (For all of the following DJ measurements, repetitive output patterns are averaged to remove RJ effects.)

The basic Deterministic Jitter tests fall into two categories:

1. There could be DJ that causes edges to be misplaced because of variations of internal rising and falling delay, or variations of internal delays caused by the spacing between adjacent transitions (e.g., internal logic swing limitations or flip-flop metastable-delay effects).
2. There could be DJ that is caused by some deterministic PLL-multiplier effect which misplaces the internal clock edges from their ideal transition time (e.g., PLL phase corrections might cause bit 0 to be always early, bit 2 always late, and the others on time).

Deterministic Jitter as a Function of Data Pattern

To check for internal delay effects and sensitivity to data content, the HOTLink Transmitter was configured to send various data patterns. Figure 20 shows how the bit position is affected by the pulse width and pattern content of the serial bit-stream. This test is performed by measuring the exact timing of a single data transition as other transitions occur at various other bit positions and pulse spacings. By plotting the exact transition time against the expected transition-time, variations in internal delay become apparent. This is a test of the output

buffer and circuits in the serial output logic, because it checks the position of a single transition as the preceding pulse width changes from nine bits to one bit LOW.

HOTLink manages to place any particular edge within ± 10 ps of the ideal location regardless of data pattern. In this test it is impossible to separate pulse distortion from slight propagation delay changes. The 50% duty cycle pattern was used as the reference location for the "ideal" pulse position. The absolute position of the rising edge was measured as the ten-bit data pattern was stepped through all nine of the possible pulse widths. Edge placement is not significantly affected by temperature or process variation. The change in edge displacement in the figure above is mostly caused by V_{CC} variation and may be related to output load currents variation.

Deterministic Jitter Caused By PLL Corrections

By causing the Transmitter to send an alternating 1-0-1-0... bit pattern, and synchronizing the measurement system to the byte-rate clock, it is possible to detect misplacement of internal clock transitions. In this test, performed by measuring the exact time of each voltage transition (averaged to remove random jitter effects), and plotting it against its expected tran-




Figure 20. Data Dependent Edge Displacement

sition-time, PLL phase corrections and clock-synchronous noise appear as fixed, repetitive, displacements from their ideal position. The square-wave bit pattern minimizes delay effects which might otherwise be present.

The results of this evaluation show that there are no deviations observable within the measurement accuracy of the test equipment. Extensive testing showed less than 2 ps deviation from the ideal position of all ten of the output transition locations, regardless of frequency.

Total Transmitter Jitter While Sending BIST

As a measure of total transmitter jitter, the tests were repeated while sending the Built-In Self-Test (BIST) sequence. In contrast to the special patterns used to measure the various components of overall jitter, this test is a more comprehensive measure of real HOTLink performance, since it highlights any effect that might have been obscured by the individual jitter-component tests.

The scope photo in *Figure 21* shows the jitter characteristics of the HOTLink output while sending the BIST pattern. This 511-character pseudo-random pattern includes all of the legal data patterns in the 8B/10B code, and is a good representation of a typical data transmission. The resulting jitter includes all of the random and deterministic jitter accumulated by the clock source, the multiplying PLL, and the internal logic and output circuits.

Transmitter Jitter Transfer Function

PLL output jitter can be affected by the noise characteristics and stability of the clock source used as its reference. The closed-loop transfer function of this type of PLL is a low-pass filter. Noise components below the rolloff frequency (f_{-3dB}) of the PLL are passed unattenuated and those above f_{-3dB} are attenuated by varying amounts. By injecting a measurable and controlled amount of noise (jitter) into the clock reference,




Figure 22. Clock-Jitter Generated by Mixing Noise into CKW

as shown in *Figure 22*, the PLL transfer characteristic can be measured.

In this configuration (*Figure 22*), the noise source is added to the clock source with a resistive mixer. Because the clock source has a significant ramp rate, the addition of the noise causes a controlled variation in the effective threshold crossing, thus causing jitter. The noise source can be any controlled source, but for this test, it was a good quality sine-wave generated by a stable generator. The amplitude was adjusted to create the desired CKW jitter amplitude (ns Pk-Pk), and the frequency was varied over a wide range while the output jitter was monitored on OUTA±. The graph in *Figure 23* shows the relationship between input and output jitter at various input jitter frequencies as the jitter frequency is increased from about 10 kHz to over 70 MHz.

While the vertical axis of this type of chart is usually expressed as a “gain” term, and uses units of dB (i.e., 20 log out/in), this data is presented as a pure ratio of input to output jitter. This allows a clearer visualization of jitter magnitude, and shows that for all frequencies of operation and noise, the output jitter falls to approximately the same “noise-floor.” By maintaining the vertical log scale, it is obvious that the effect being illustrated is the same as the closed-loop PLL transfer function described earlier.

As expected, low-frequency jitter passes through the PLL unattenuated. Higher frequencies are attenuated until the jitter frequency approaches the Character-Clock (CKW) frequency. This jitter-feed-through peak at about the reference frequency is the result of the sum-and-difference frequencies that naturally result from mixing. A significant frequency component is generated at the “difference” between CLOCK (f_{REF}) and NOISE (f_N) sources. When this “difference” frequency falls within the PLL filter bandwidth, it passes unattenuated to the output of the PLL and appears as jitter exactly as if it was caused by an equivalent low-frequency input-noise source. This effect is enhanced by using a single frequency noise source. The energy at any particular frequency of a wide-band noise source is relatively small, but feeds through in exactly the same way. Narrow-band noise sources that operate synchronously with the HOTLink CKW input-rate might cause more of these “mixed-down” frequency components that would also feed through and emerge as output jitter.

The bandwidth of the HOTLink Transmitter PLL varies slightly as a function of the operating frequency as shown in *Figure 23* ($t_B = 1/\text{baud}$; $t_B = 6 \text{ ns} = 160 \text{ Mbaud}$, etc.). This variation is attributable to variations in VCO gain that are a function of operating rate.




Figure 21. HOTLink Transmitter OUTA± Total Jitter in BIST vs. BIT Rate Reference




Figure 23. HOTLink Transmitter Jitter Transfer Function

The scope traces in *Figure 24* graphically show the jitter feed through from the TTL-CKW input to the PECL outputs of the HOTLink Transmitter. Low jitter frequencies pass through unattenuated, and high frequencies are significantly attenuated. For high-frequency input jitter, the only jitter remaining on the output is roughly equivalent to the intrinsic jitter of an undisturbed PLL.

V_{CC} Jitter Transfer Function

To characterize HOTLink output jitter characteristics in the presence of noise carried on the power supply, the Random Jitter set-up was modified as shown in *Figure 25*.

In this test the power supply was intentionally disturbed. By injecting a measured amount of noise into the V_{CC} pins of the HOTLink Transmitter (using an external driver), the jitter effects of power supply noise could be observed.

As expected, when the power supply is disturbed, the output contains some additional jitter. Increasing amplitude disturbances cause increased jitter amplitudes. It is difficult to create this much noise on normal system boards because of the normal power supply bypassing that is usually applied to this type of component. Large amplitude V_{CC} spikes are removed by the bypass capacitors.

As the V_{CC} noise frequency is varied, the jitter out also varies. At low frequencies of noise, the jitter is small, and at high noise frequencies the jitter is also small. Around the PLL roll-off frequency (as measured in the previous analysis), the jitter output increases (see *Figure 26*).

Transmitter PLL Lock Time

Multiplying PLL lock characteristics are mostly a function of the internal loop dynamics. While the loop is changing the clock period from one frequency to another, the transition is

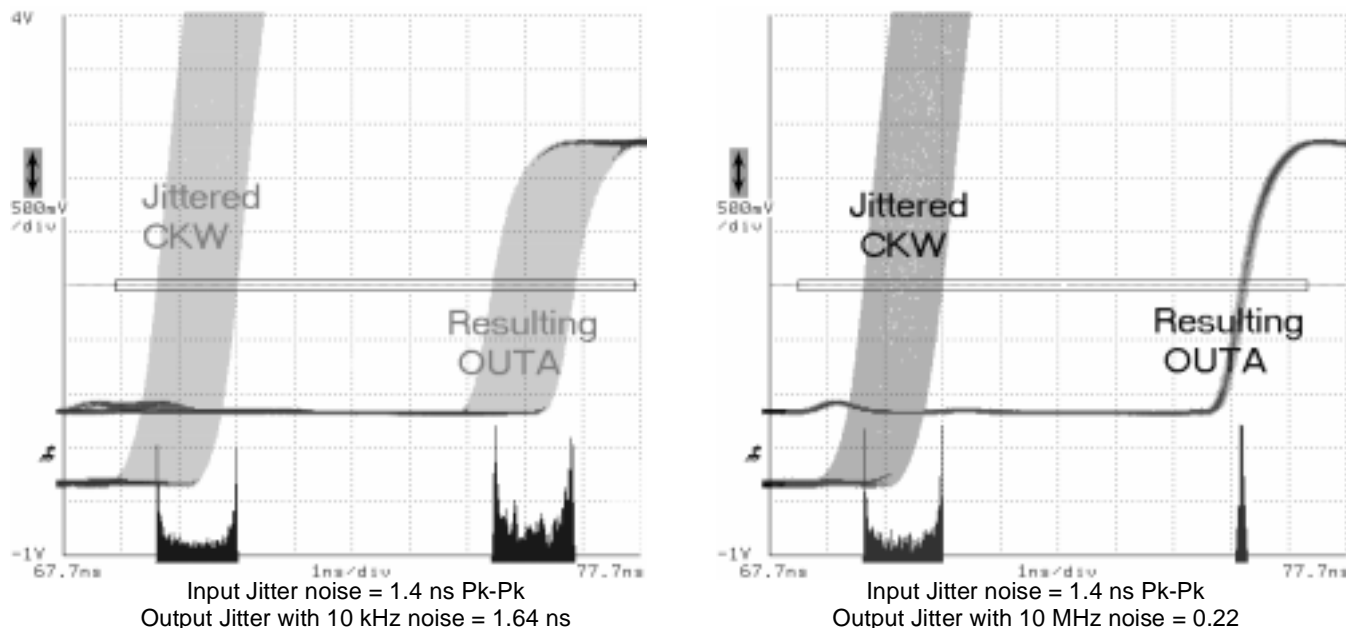


Figure 24. Serial Output Jitter Varies as a Function of Input Noise Frequency




Figure 25. HOTLink Transmitter V_{CC} Coupled Jitter Set-Up

not monotonic as might be expected, but appears to oscillate around the acquisition trend-line until the frequency falls within the PLL "Lock Range." This effect is a result of the normal characteristics of the Phase Frequency Detector (PFD) and the Loop filter. The PFD output is a continuously varying series of pulses that cause the VCO to change frequency in the desired direction, but the resulting pulse widths are not constant. The graph in *Figure 27* shows each of the Bit-Clock periods during the PLL "re-lock" progression while the loop locks to a higher and a lower frequency.

The Bit-Clock period listed in *Figures 27* and *28* is the period of the internal VCO inside the HOTLink Transmitter. While this signal is not directly observable at a HOTLink Transmitter pin, its period may be directly calculated from measurements

made at the serial output (OUTA±) pins. In this test the serial data pattern was set to the equivalent of a K28.7 special code (0000011111), fixing the bit-clock period at one tenth the interval between output rising transitions.

The two locked frequencies were selected to be the minimum and maximum actual operational limits of the part under test. Cycle-by-cycle times were recorded in a continuous stream using an HP54720D. The frequencies used for these illustrations are well outside the data sheet operational limits of HOTLink, but were the actual functional limits of this particular (typical) part. While all parts behave in a similar manner, some have slightly higher or lower operational frequency limits. The acquisition rate varies slightly with temperature, being slowest at higher temperatures.




Figure 26. HOTLink Transmitter V_{CC} Coupled Jitter Transfer Characteristic




Figure 27. Transmitter PLL Acquisition Characteristic (from Locked to Locked)




Figure 28. Transmitter PLL Time to Lock (Quiet to Locked)

The PLL behavior is slightly different when the character clock (CKW) is removed for some time (see *Figure 28*). Instead of immediately beginning to acquire the new frequency, there is a time after CKW begins but before there is any change in the VCO frequency. This time is required for the internal control nodes to move from their “ranged-out” levels (resulting from the PLL trying to track to zero or infinity Hz) to within the compliance-limits of the amplifier and VCO. After the change begins, it moves at the same rate as for the previously described cases.

When CKW is removed, the loop immediately begins to slew toward the lowest possible speed. The transition does not have the “jaggies” typical of a loop tracking to another clock rate, because there are no reference clock edges to modulate the PFD output. The frequencies shown in this illustration are well outside the data sheet operational limits of HOTLink, but were the actual functional behavior of this particular (typical) part. While all parts behave in a similar manner, some have slightly higher or lower frequency limits. The acquisition and slew rates were similar but vary slightly with temperature, being slower at higher temperatures.

HOTLink Receiver Jitter

The PLL used to synchronize an internal clock to a received bit stream (i.e., in the HOTLink Receiver) has different requirements than those for a multiplying PLL. This loop is effectively a one-to-one loop where the bit clock (Received Bit Clock, an internal signal) runs at the same rate as the incoming data stream (Serial IN, an external signal). The Received Bit Clock is used to sample the Serial input at regular intervals, thus extracting the serial data (Retimed Data, in *Figure 29*). This same signal runs all of the internal logic for deserializing, framing, and decoding the serial data. Any disturbance that can affect the PLL and the Recovered Bit Clock will affect both the quality of the data recovery and the quality of the byte-rate, data-synchronous clock that is provided to the receiving system.

Receiver jitter affects systems in at least two ways. Jitter tolerance is a major determinant of system margin, and Jitter feed-through can reduce timing margins in the receiving host system.

Jitter feed-through is a function of the PLL filter characteristics, and can be directly measured at the CKR output of the HOTLink Receiver in much the same way used to test Transmitter jitter feed-through.

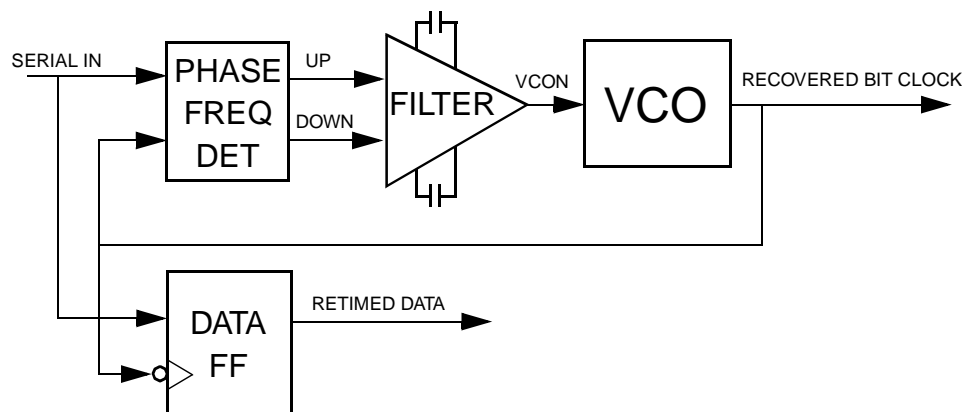


Figure 29. HOTLink Receiver PLL Block Diagram




Figure 30. Technique to Measure Static Alignment

Jitter tolerance is more complicated, since it is a measure of the Receiver's ability to correctly capture and interpret incoming data, and must be measured indirectly. Jitter tolerance is both a function of the intrinsic jitter in the receive-clock synchronization PLL and the effects of received data upon it. Tolerance is also a function of the precision timing and alignment of internal clock edges (i.e., the clock edge used in the PLL to synchronize the data, and the clock edge used to sample the incoming data stream). The data-sampling flip-flop set-up/hold timing characteristics and their variation contribute to further jitter tolerance degradation.

To isolate the effects and tolerance limits to various types of jitter, carefully designed tests were performed on HOTLink parts selected from the full spectrum of manufacturing variation. These tests were designed to separate the effects of power supply, data characteristics, external clock sources, and various PLL characteristics. Unless otherwise noted, static variations in power supply levels (4.5V to 5.5V), ambient temperature (-55°C to 125°C), and process variations (within manufacturing tolerance limits) cause virtually no change (within the accuracy of the measurement system) to any of the following jitter tolerance or PLL characteristics.

Static Alignment and Error-Free Window

To maximize jitter tolerance, the receive circuit is designed to sample the incoming data at a point *exactly* halfway between the ideal transition times of uncorrupted data. This requires that the PLL track the incoming data and align itself with the "average timing" of the received edges. The precision of this alignment is often called "Static Alignment" and should have a magnitude of zero, indicating perfect alignment of VCO and the data and perfect 50% sampling alignment. Using this recovered clock, the incoming data is sampled at the point that gives maximum tolerance to misplaced edges and maximizes the error-free window. Any misplacement of this sampling point will reduce jitter tolerance.

Static alignment of the HOTLink Receiver was evaluated using the technique shown in Figure 30. The HOTLink Transmitter and the Receiver under test were configured to send and receive the BIST pattern. Then, by inserting a BIST-synchronous pulse on the FOTO pin (using a generator triggered on the $\overline{\text{RP}}$ output of the HOTLink Transmitter), one transition in

the transmitted data pattern was varied to find the maximum "misalignment" possible before the onset of an RVS error indication. This configuration allows the receive PLL to have about 3000 "ideal" transitions (i.e., the total number of transitions in the 511 byte BIST loop) and only one misplaced edge. Shorter patterns modified in this way (e.g., a single data byte with byte-synchronized FOTO pulses having a single misplaced transition) give an erroneous result. The very large phase error which occurs in one of the ten bit positions is averaged out by small-compensating phase-adjustments during the other nine bit-times. The BIST pattern test allows the PLL phase-correction response from the single-edge error to settle out before the next error appears so that the averaging effect does not color the data-capture results.

Data transitions can be misplaced from their ideal position by almost half of a bit-time without erroneous sampling by the data recovery flip-flop. The data characterization summary in Figure 31 indicates that the HOTLink Receiver accepts misplaced edges to within about 250 ps of the half-bit point. The center of the small error region where data is not sampled correctly (at approximately 180 ps after the ideal mid-bit point) is the actual PLL static alignment position. The width of the error region (about 150 ps) is attributable to both the sampling flip-flop metastable region, and the internal PLL clock jitter.

This data alone implies that any data edge could fall anywhere within a bit time (minus about 500 ps) and still be de-




Figure 31. HOTLink Receiver Static Alignment as a Function of Frequency

coded correctly. This is almost correct, except for the effect of receiver clock jitter caused by the various types of incoming jitter.

Duty Cycle Distortion Jitter Tolerance

The characteristics of some types of interconnect circuits cause Duty Cycle Distortion which the receive system must tolerate. DCD jitter alters the placement of all transitions in the data stream by about the same amount (in alternating directions) regardless of the bit pattern being sent. For small amounts of jitter, this alternating error tends to cancel out, and the loop behaves normally while recovering data without error.

As the magnitude of jitter increases, phase correction pulses from adjacent misplaced edges start to interact. Each correction pulse has some finite duration, usually a significant percentage of the expected bit time, and is proportional to the magnitude of the edge misplacement. Since jitter is also expressed as a percentage of a bit (usually a large percentage) the interaction between jitter magnitude and phase correction pulse width determines DCD jitter tolerance. When adjacent phase corrections interact, they sum in unexpected ways which affect the resulting correction response. When these interactions are rare or small, there is no apparent effect. If the interactions affect most of the phase correction events, the PLL stability, predictability, and output jitter are affected and data is not captured correctly.

Figure 32 shows HOTLink Receiver DCD jitter tolerance. This test was performed by carefully corrupting the link between a HOTLink Transmitter and Receiver with increasing magnitudes of DCD (see Jitter Generator circuit and description in Figure 44). Using the BIST test capability included in the HOTLink devices, DCD tolerance limits are declared as being exceeded when the RVS output of the Receiver indicates approximately one error every ten seconds (i.e., $BER \leq 4 \times 10^{-10}$ at 250 Mbaud). Slight differences in jitter tolerance exist between parts from different process corners, but no appreciable variation was found for V_{CC} or temperature variation. The DCD tolerance characterization data shown in Figure 32 varies by less than 5% across the full process spread (e.g., from 1.42 ns to 1.39 ns out of a 3.0 ns bit-time). The threshold of failure is very abrupt. At these levels of DCD jitter, changes in jitter amplitude of less than ± 100 ps make the difference between almost-perfect data reception, and almost-total corruption.

In contrast to the predicted jitter tolerance that comes from the Static Alignment test, and the DDJ tolerance (see following text), DCD tolerance at first appears to be much smaller. This apparent reduction in jitter tolerance is entirely due to PLL and Phase-Detector effects, and does not result from any anomaly in the data recovery path. Data can be recovered correctly at the levels of edge misplacement that are found at the limits of DCD tolerance, but not above. By carefully approaching the limit, it can be seen that the PLL loses lock at the jitter magnitudes shown in Figure 32, and then regains it at slightly higher jitter levels, but with a massive clock jitter, often slipping bits as the jitter goes through the "magic point," destroying any possibility of data recovery. The recovered clock shows almost no jitter feed-through when DCD is present, and remains below the "data-corruption" threshold (as shown in Figure 46).

Fortunately, most transmission links don't include large amounts of DCD. The most common contributors are mis-




Figure 32. Duty-Cycle-Distortion Jitter Tolerance as a Function of Data Rate

matched output loads on differential or single-ended PECL outputs, and improperly designed or operated optical interface modules. Single-ended PECL outputs can change the effective delay of the driver by about ± 0.5 ns. Differential outputs are typically more symmetrical. Optical-to-electrical (receiver) interface modules running with extremely high or low light levels can have non-linear and asymmetrical delay characteristics that affect the pulse symmetry of the receiver output, used by the PLL data recovery circuits. The optical emitter in an electrical-to-optical interface module also has non-symmetrical turn-on and turn-off characteristics which are normally compensated by careful design of the drive electronics. At the limits of performance, optical modules can add more than ± 1 ns of DCD.

Data Dependent Jitter Tolerance

The characteristics of some types of interconnect circuits cause Data Dependent Jitter which the receive system must tolerate. The same "correction-pulse" interaction that limits DCD tolerance also affects DDJ tolerance. Since the collisions between adjacent correction pulses occur at a much less frequent and regular rate, the effect is smaller. The "clock-jitter" that results from these corrupted corrections reduces the jitter tolerance to less than the ideal maximum that the Static Alignment test might predict.

Figure 33 shows HOTLink Receiver DDJ jitter tolerance where the DDJ was generated by an artificial generator. This test was performed by carefully corrupting the link between a HOTLink Transmitter and Receiver with increasing magnitudes of DDJ (see Jitter Generator circuit and description in Figure 44) while sending a continuous BIST pattern. Errors were most typically associated with the long running bit pattern included in a K28.5 character, and the same tolerance was observed while receiving only corrupted K28.5s. The worst DDJ peak always follows the 1111101 and the 0000010 contained in the special characters. Using the BIST test capability included in the HOTLinks, DDJ tolerance limits are




Figure 33. Data-Dependent-Jitter Tolerance as a Function of Data Rate

declared as being exceeded when the RVS output of the receiver indicates approximately one error every ten seconds (i.e., BER 4×10^{-10} at 250 MBaud). Slight differences in jitter tolerance were found between parts from different process corners, but no appreciable variation was found for V_{CC} or temperature variation. The DDJ tolerance characterization data shown in *Figure 33* varies by less than 5% across the full process spread (e.g., from 2.04 ns to 1.86 ns out of a 3.0 ns bit-time). The threshold of failure is very abrupt. At these levels of DDJ jitter, changes in jitter magnitude of less than ± 100 ps make the difference between almost-perfect data reception, and almost-total corruption.

Interconnect Link Jitter Tolerance

The tolerance to synthetic-DDJ, as shown in *Figure 33*, is slightly worse than that found when the jitter is natural-DDJ. The variation is caused by unintentional DCD introduced by the test system used to create a stable and repeatable test pattern at all frequencies over which HOTLink might operate.

Wire transmission-line jitter is dominated by DDJ caused by the variation in attenuation as a function of frequency. Higher frequencies are attenuated more than lower ones. This rising attenuation-with-frequency characteristic of wire links causes the wider pulses (i.e., multi-bit strings of ones or zeros) to have a higher amplitude than the shorter pulses, since the higher frequencies (those attenuated the most) are required to make the fast edges and narrow pulses, while the wider pulses contain more low-frequency components. This variation in amplitude results in variations in pulse placement, since the edge rate is almost constant and the variation in amplitude causes variations in the time at which a transition crosses the receiver threshold.

This effect is most visible when a single, worst-case character is measured. *Figure 34* shows the edge misplacement caused by the different-length pulses in a continuous K28.5 pattern (i.e., 11000001010 00111110101...). When the data




Figure 34. DDJ Characteristic of K28.5 at 250 MBaud after 250 ft. RG-59

is more normally distributed, it becomes more difficult to see the distinct pulse positions, and the jitter just merges into a continuous “uncertainty-zone” (see *Figure 35*).

Using actual data and real transmission lines, the HOTLink tolerance to DDJ appears to be a more constant function of bit-rate than shown in *Figure 33*. If about 500 ps of eye-opening can be maintained, the data is recovered correctly, regardless of the data rate. However, recovered clock jitter increases with increased DDJ (see *Figure 47*).

In wire transmission links, the accumulation of DDJ determines the maximum distance over which data can be reliably communicated. The characteristics of the chosen media




Figure 35. BIST data at 370 MBaud after 250 ft. of RG-59 coax (BER $< 4.5 \times 10^{-11}$ with < 700 ps eye opening)




Figure 36. Maximum Data Rate vs. Uncompensated Wire Length (BER < 3 x 10⁻¹²)

determine the usable distance. The total attenuation of the line is rarely sufficient to limit the maximum usable distance, even though the data bits that are incorrectly interpreted have minimal amplitude at the time the error. This loss of amplitude is a result of the variation in peak voltage attained during any particular pulse.

HOTLinks have been designed to offer more than 20 dB of attenuation margin between the transmitter output and the receiver input. Typical maximum-distance, uncompensated links have less than 10 dB of high frequency attenuation due to the transmission line and interconnect components. The remainder of the interconnect budget is used to compensate for the difference between high- and low-frequency attenuation of the wire transmission line. Compensated wire links have been built that operate reliably over more than double the distances shown in Figure 36.

Fiber-optic links, in contrast to the wire links described previously, are limited by optical attenuation, chromatic and modal dispersion, and the resulting Random Jitter in the optical-electrical converter. At the limit of operational optical margins, the low light levels into the receiver and the dispersion from the fiber combine to create misplaced data transitions. These displacements are usually random, but in the case of some optical modules, can also include significant Duty Cycle Distortion.

Peak random jitter tolerance *should* be approximately the same as the Static-Alignment limits described previously (Figures 30 and 31). The simplest way to generate random jitter involves a long piece of fiber-optic cable, and appropriate fiber-optic interface modules. As fiber length is increased, both dispersion (i.e., pulse distortion caused by the variations in propagation delay through the fiber, as a function of optical wave-length and propagation modes) and attenuation also increase, causing the jitter out of the optical-to-electrical converter to increase. There is, however, a limit to the attenuation, beyond which the fiber-optic receiver cannot recover the

data correctly. Attenuation alone, without the effects of long fiber-optic cable, often causes significant DCD in the link. This DCD obscures the real random jitter behavior of the receiving PLL.

The random jitter output of a 5-km piece of 62.5-μm core multi-mode fiber is shown in Figures 37 and 38, and illustrates a typical problem that occurs when trying to measure random jitter and jitter tolerance. These pictures were taken at the limit of frequency vs. length, as indicated by BIST errors appearing on RVS. The “eye-diagram” in Figure 37 was taken using the traditional infinite-persistence scope measurement, where the scope is triggered by a pristine bit-clock. The trigger-clock, shown below the eye-diagram for reference, is arbitrarily placed with respect to the jittered data trace. This is the resulting display of an HP54720D at 8 Gs/s after about four hours of jitter accumulation (approximately 30,000 traces). It would appear that the jitter tolerance of the receiver is only about 45% (i.e., 4.0 ns – 2.19 ns) at the measured BER. *This conclusion is incorrect.*

Figure 38 offers another view of the same link and error rate, when triggered by the error event, and shows the *actual* eye opening. This view, triggered by the pristine bit-clock qualified by RVS (ANDed), shows that when the HOTLink indicates an error event, the “eye” is actually fully closed. This picture displays only those traces that contained an error event, about one every four seconds at 250 MBaud. It is impossible to determine from these pictures exactly where the PLL and the data sampling flip-flop have placed the bit boundaries, but it is obvious that if the transition doesn’t cross the threshold, the data is lost. (The “ghost” traces that appear in the picture are parts of other error-traces where the eye-closure occurred at some other bit position beyond the limits of the screen.) The discrepancy between these two figures is caused by the triggering and display characteristics of the scope. Even though there are over 30,000 patterns displayed on the first one (Figure 37), it just happened that none of the error bits were cap-




Figure 37. Random Jitter out of Fiber-Optic Link Triggered by Bit-Clock

tured. This could have been because of the relative rarity of the events, and the trigger hold-off caused by the scope processing that occurs between measurements.

Receiver Data-Phase Acquisition Time

To measure the HOTLink Receiver response to phase-hops in the incoming data stream, it is necessary to produce a data stream that has a controlled phase change. It is possible to use the two selectable inputs of the HOTLink Receiver to switch between two identical, but skewed, data streams. The data stream used for these tests comes from a HOTLink Transmitter using a good quality clock source. The HOTLink BIST function provides a convenient source of repeatable data and is accompanied by a convenient trigger pulse in the RP output that occurs once per BIST loop. The Receiver BIST comparator can be used to determine whether the receiving PLL has maintained phase lock without slipping by monitoring its RVS output. This output pulses only if there is an error in the received data pattern.

In the test set-up shown in *Figure 39*, the input to the INB+ pin of the Receiver is skewed with respect to the INA± input using the precision skew capability of the Colby delay generator, which can add delay up to 10 ns in 1 ps increments. A carefully placed control pulse (i.e., inputs are changed only when both inputs remain at the same logic level for a few bit times to ensure that the change does not affect the serial data stream), which is a BIST-synchronous control signal (i.e., the pulse is triggered by RP which occurs once in per BIST loop), switches the receiver input between the two data streams. As expected, when the A/B input switches between these two streams, no errors are indicated if the skew is small. When the skew is increased, and approaches almost half of a bit time (i.e., 135 to 150 degrees as seen by the PLL Phase Detector) errors are indicated by pulses on RVS. These errors are caused by "bit-slip" in the PLL as it reacquires the new data stream.




Figure 38. Random Jitter out of Fiber-Optic Link Triggered by RVS & Bit-Clock

By triggering the HP54120D on the signal that changes data streams, it is possible to observe the real-time behavior of the receiving PLL. The scope can be programmed to measure either clock period or propagation delay between two channels. The former shows each clock period as the loop acquires the new data stream. The latter set-up shows the more traditional phase-alignment measurement that defines Phase-Locked-Loop acquisition characteristics.

Measurements were taken with various amounts of phase difference between the two input channels. *Figure 40* shows the characteristics of the HOTLink Receiver with phase errors less than 180 degrees and with phase errors at as close to 180 degrees as possible. This illustrates typical link performance. *Figure 41* shows the worst-case phase acquisition characteristic.

In the test set-up shown in *Figure 39*, the HOTLink Receiver input is switched to one of the inputs, allowed to stabilize there for a few byte times, and then switched back. The second switch is an equal magnitude phase offset, but opposite sign. During the time when the PLL is trying to regain phase alignment with the incoming data stream, it adjusts the period of the VCO, and thus the output clock of the HOTLink Receiver. As illustrated in *Figure 40*, the phase correction begins immediately after the change in data stream.

Since the phase error is less than 180 degrees, the correction is always in the expected direction. When the new data stream "lags" the current PLL position, the clock is stretched for a few cycles until it realigns with the incoming data. Likewise, when the new data stream "leads" the current PLL phase, the clock is shortened for a few cycles until it realigns with the incoming data.

The change between any pair of clock (CKR) periods is small, and the maximum deviation usually varies by less than ± 1 ns midway through the seven to ten byte-times required to re-




Figure 39. Set-Up to Measure HOTLink Phase Acquisition Characteristics




Figure 40. Phase Hop of less than 180 degrees without Data Corruption

align the clock. The magnitude of change that can be accommodated without error varies slightly with frequency, and the time needed to resume normal clock periods varies by one or two character times. There is little or no correlation between settling time and the sign of the phase-change, data-speed,

process-corner, V_{CC} -level, or ambient-temperature. For all frequencies, it seems that any phase change that is less than a half-bit time (less about 500 ps) is accommodated without data corruption. The Character-Clock adjustment shown in *Figure 40* is the accumulated sum of the ten Bit-Clock periods




Figure 41. Phase Hop Timing with Exactly 180-Degree Phase Difference

that combine to make up the Character-Clock adjustment, each of which is probably much smaller.

When the phase change is carefully adjusted to the 180-degree position, the correction behavior changes. The correction can now occur in either direction, since both have an equal capability to realign the PLL clock phase. One direction will cause a bit slip since the decoding logic will find the data appearing one bit earlier or later than expected. The other direction might not slip, but will probably still indicate a corrupted byte because of a metastable response from the data sampling flip-flop.

Additionally, the phase correction does not start immediately after the change in incoming data phase (see *Figure 41*). The time it might take cannot be calculated, because the loop is operating outside its linear response region, and will assume some metastable behavior that could theoretically take forever to clear. It takes several byte times before the PLL accumulates enough error information to cause it to realign itself. When the data has exactly 180 degrees phase offset to the PLL VCO, the Phase Detector may have either no phase-correction effect or a small reverse phase-correction effect, in contrast to its normal, increasing-correction with increasing-error, linear-phase-correction response to smaller phase errors. Once it begins to change, the PLL completes the phase hop in about the same way as the earlier example showed, although over a slightly longer duration. Perhaps counter to intuition, the more quiet the received data stream, and the cleaner the VCO clock, the longer this "hang time" becomes. (Products with "jitter problems" will never exhibit this "hang phenomenon.") Any jitter or frequency deviation between the incoming data and the VCO provides a tie-breaker and gives enough error information to allow the Phase Detector to begin its change. Once the relative-phase has moved only a little bit, it becomes obvious to the Phase Detector that the error is large and requires a large correction. Complete phase alignment is not achieved until several byte times after the CKR output has resumed its normal period.

This final alignment time is immaterial for most data-communications systems, since the receiving system will have long since resumed correct data recovery because of the wide jitter tolerance of the receiver. As shown in the limited phase-step experiment, the HOTLink Receiver recovers the bit-stream correctly when the input transitions are more than

about 500 ps away from the mid-bit point. However, the *data* must be "framed" to be interpreted correctly, and the time necessary to accomplish framing (HOTLink requires one or two K28.5 characters to frame, depending on current framer mode) depends on the protocol being used, and how often SYNC characters occur.

Receiver Data-Frequency Acquisition Time

Two serial data streams rarely operate at exactly the same frequency, so the PLL must first acquire the new frequency before it can align the clock to the phase of the new data. The frequency offset that must be accommodated is different for each standard system, but is usually a few hundred parts per million (PPM) variation from a specific frequency. Fibre Channel, for instance, specifies a maximum frequency offset of ± 100 PPM ($\pm 0.01\%$). HOTLink is specified to accommodate frequencies of $\pm 0.1\%$, but typically accommodates more. The chart in *Figure 42* illustrates the behavior of a HOTLink receiver as it switches between two data streams that are offset from the local REFCLK frequency by $\pm 0.2\%$. The acquisition time and the effective clock-period transient time is equivalent to that seen when the receiver is only adjusting for phase differences.

This is typical of the HOTLink acquisition behavior at all operating frequencies, and doesn't vary significantly across process, V_{CC} , or temperature variations. The exact transient size and duration varies from event-to-event because of the statistical nature of the "first-change." The excursion could be either to a shorter or a longer clock period depending upon the perceived phase of the new data when the change occurs. Likewise, the time to achieve phase alignment varies slightly depending on the probability of having a perfect 180 degree phase alignment after the change. None of the clock period excursions measured during this test exceeded ± 1.0 ns.

When the data stream is not within the frequency tolerance limits of the receive PLL, the HOTLink automatic frequency-range-control mechanism modifies the transient behavior of CKR. This function continuously monitors the frequency of the VCO and compares it to the frequency of REFCLK. When they are different by a sufficient amount, and for a sufficient time, internal logic forces the VCO to align to the REFCLK frequency. This automatic mechanism ensures that the absolute frequency of CKR is never far from its ideal period, and




Figure 42. Frequency Hop Within $\pm 0.2\%$

that when “good” data returns, the PLL can rapidly align to it, and begin valid data recovery.

This test is similar to the previous one, except that the data is not switched externally, and the recorded transient is only the result of the internal re-lock to REFCLK behavior. In this case (shown in *Figure 43*), the applied data-stream was offset from REFCLK by about 3.0% (well beyond the datasheet limit of 0.1%). The HP54720 was “glitch triggered” when RVS was HIGH for >60 byte times. (RVS-HIGH for 64 byte times is the PLL out-of-lock indication, since normal data does not yield continuous error indications.)

For the first few bytes (out to about byte-time 45 in *Figure 43*), the average period of CKR is about 3% faster than the expected 30 ns which indicates that HOTLink has been successful in acquiring the data frequency. When the built-in automatic range control is asserted, there may be a momentary transient in the CKR period caused by the phase and frequency of the PLL relative to the instantaneous bit-stream phase.

Next, the VCO is pulled to the REFCLK frequency by the internal range-control logic (from about Byte 45 to about Byte 110 in *Figure 43*). Finally the PLL is released to track the incoming data, whereupon it might immediately return to the previous frequency (the frequency of the incoming bit-stream, if any), or (as in this illustration) hunt around for an indeterminate time (possibly an indefinite time) until it again finds a signal within its acquisition and tracking range. The exact PLL behavior depends on the frequency, transition density, timing characteristics, and stability of the applied data stream.

CKR period excursions are slightly larger when this range control mechanism is applied, but still under about ± 1.2 ns. The period of CKR is the sum of all Bit-Clock periods that occur between CKR transitions.

Receive PLL Jitter Transfer Function

PLL jitter, and consequently recovered clock jitter, can be affected by the noise characteristics and stability of the incoming data stream. The closed-loop transfer function of the PLL




Figure 43. Frequency Acquisition from/to $\pm 3.0\%$ Changes CKR Period




Figure 44. Data-Jitter is Generated by Mixing Noise into Serial-Data-In

is a low-pass filter. Noise components below the natural frequency (f_n) of the PLL are passed unattenuated, and those above f_n are attenuated. By injecting a measurable and controlled amount of noise (jitter) into an otherwise stable data stream (as shown in Figure 44), the PLL jitter-transfer characteristic can be measured.

In this configuration, the noise source is added to the data-clock source by a resistive mixer, similar to that used for transmitter jitter-transfer testing. The mixer output drives the external bit-rate clock input of a high speed data generator. The Microwave Logic GigaBERT 1400 can run with clock rates above 1 GHz, and can send serial data from an internal memory using this clock. By jittering the external clock, it is possible to create a controlled serial data stream with single frequency jitter noise. The amplitude of input jitter was adjusted to create the desired data jitter amplitude (ns Pk-Pk), and the frequency was varied over a wide range while the jitter was monitored on the CKR output.

Direct jitter generation is difficult to manage because of the need for a single frequency noise source superimposed on an otherwise perfect data stream. Most jitter generators seem to generate either multiple frequency noise sources or have significant DCD and DDJ. The method described for creating

jitter suitable for Transmitter jitter-testing creates significant DCD which is ignored by the transmitter PLL, since it only responds to the rising edges of its reference input. Because the receiver responds to both edges of the pulse, this DCD affects the results in undesirable ways. The graph in Figure 45 shows the relationship between input and output jitter at various input jitter-noise frequencies

As expected, low frequency noise passes through the PLL filter unattenuated and higher frequencies are attenuated as theory would predict. Also as expected, the apparent bandwidth of the PLL filter varies as the transition density of the data stream varies. For the highest possible transition density (e.g., a 1010101... data stream) the natural frequency is highest, and for lower transition densities it is proportionally lower. The information shown here is characteristic of the HOTLink while receiving normal data. In this case the data was the BIST pattern.

Effective loop-bandwidth varies as a function of data rate, as shown in Figure 45. This variation is caused by various gain changes within and between the PLL component blocks. Some blocks have analog gain variations as a function of frequency, and others have a constant output response regardless of operating frequency. The behavior shown in Figure 45





Figure 45. HOTLink Receiver Jitter Transfer Function (BIST Data)



250 MBaud, Data = BIST,
DCD = 1.94 ns Pk-Pk


**Figure 46. CKR Output Jitter as DCD
Distorted Data is Being Received**

is unaffected by temperature, V_{CC} variation, and variations in manufacturing tolerance.

The Receive PLL jitter transfer function is not sufficient to determine what the actual jitter out of the HOTLink Receiver might be. Different types of jitter have different transfer characteristics. DCD-type jitter causes essentially no output jitter for input jitter magnitudes up to the point where the data is corrupted. The waveforms in *Figure 46* illustrate the jitter feed-through characteristics of the HOTLink Receiver. The input waveform is a continuous stream of 1-0-1-0-... bits that have been artificially distorted with the DCD Jitter generator described later in this application note and shown in *Figure 46*. The 4.0 ns bits have been narrowed by almost half (about 1.96 ns, see the twin-peak histogram in *Figure 46*), and the CKR output shows less than 100 ps of jitter as illustrated by the darker trace superimposed on the input jitter waveform (note that the two traces have different vertical scales, but the same time scale).

When DDJ is applied to the data input, CKR jitter increases. The eye diagrams in *Figure 47* shows that when DDJ approaches maximum tolerable levels, the CKR output jitter increases appreciably. The test results shown in *Figure 47* were performed using the same maximum tolerance jittered data shown in *Figure 35*. This signal was generated using the BIST sequence transmitted through 250 feet of RG-59 coaxial cable at 370 MBaud, while operating with a received BER of $<4.5 \times 10^{-11}$. (The measurement in *Figure 47* is triggered by the pristine-bit-clock, which results in copies of the byte-rate TTL clock displayed at bit-clock intervals.)

This jitter feedthrough is partly caused by the low-frequency characteristic of the jitter, which is determined by its data content, and partly because the actual PLL failure mode (as opposed to Data failure mode) is the same for DDJ as for DCD. In either case, when any data-pulse falls below the DCD



370 MBaud DDJ in = 2.18 ns Pk-Pk,
CKR Jitter = 1.31 ns Pk-Pk

Figure 47. CKR Jitter Output as a Function of DDJ Input

pulsewidth limit, the PLL drops some of its tracking and locking information. In a normal data stream this loss is not regular, and causes minimal disturbance. The main effect is to increase jitter on the CKR output.

Summary

The summary data in *Table 1* is representative of the devices tested and described in this report. This evaluation included parts from across the full manufacturing spread, which were tested over the full range of temperature, voltage, and frequency of operation. This data is representative of HOTLink in-system performance, but because of the small sample size tested, it cannot necessarily be assumed to be worst case.

Hints to Improve Measurement Accuracy

- Use differential scope inputs instead of single-ended measurement systems to remove common-mode amplitude variations from timing jitter. Minor variations in power supply levels that are passed through to the complementary PECL outputs are ignored by the differential receiver, and so should be removed from the measurement. Systems with only single-ended scope inputs should carefully monitor V_{CC} -coupled signals, since a few millivolts of vertical shift can result in several picoseconds of apparent delay variation. Faster edges and minimal loading can minimize the problem, but not eliminate it.

Table 1. Summary of HOTLink Jitter Characteristics

| Parameter | Characteristic | | Condition |
|--|---|--|---|
| Tx Cycle-Cycle Random Jitter | < 6 ps RMS | < 50 ps Pk-Pk | |
| Tx Input-Output Random Jitter | < 20 ps RMS < 22 ps RMS < 30 ps RMS | < 175 ps Pk-Pk < 190 ps Pk-Pk < 250 ps Pk-Pk | (330 MBaud) (250 MBaud) (160 MBaud) |
| Tx Data Dependent Edge Displacement | | < ± 10 ps Pk-Pk | |
| Tx PLL Deterministic Edge Displacement | | < ± 2 ps Pk-Pk | |
| Tx Total Transmitted-Data Jitter | < 26 ps RMS < 28 ps RMS < 36 ps RMS | < 230 ps Pk-Pk < 250 ps Pk-Pk < 300 ps Pk-Pk | (330 MBaud) (250 MBaud) (160 MBaud) |
| Tx Closed-Loop Bandwidth (3 dB) | | 1.5 MHz 0.6 MHz 0.3 MHz | (330 MBaud) (250 MBaud) (160 MBaud) |
| Tx Re-Lock Rate (Locked to Locked) | | > 11 MHz/ μ s > 9 MHz/ μ s | Typical Hot |
| Tx Crash Rate (From CKW Stop) | | > (45 MHz + 19 MHz/ μ s) > (21 MHz + 16 MHz/ μ s) | Typical Hot |
| Tx Lock Time (Quiet to Locked) | | < 45 μ s < 60 μ s < 80 μ s | Typical (160 MBaud) Typical (330 MBaud) Hot (330 MBaud) |
| Rx Error-Free-Window (Static Alignment) | | > $t_B - 250$ ps | Note: $t_B = 1/\text{Baud rate (ns)}$ |
| Rx Random Jitter Tolerance (BER < 1×10^{-12}) | | > $t_B - 500$ ps | |
| Rx DCD Tolerance (BER < 1×10^{-12}) | | > $0.42 \times t_B$ | |
| Rx DDJ Tolerance (BER < 1×10^{-12}) | | > $0.62 \times t_B$ > $0.82 \times t_B$ > $0.95 \times t_B$ | (330 MBaud) (250 MBaud) (160 MBaud) |
| Rx Total Jitter Tolerance (BER < 1×10^{-12}) | | > $t_B - 500$ ps | |
| Rx Input-Output Random Jitter | < 39 ps RMS < 25 ps RMS < 24 ps RMS | < 224 ps Pk-Pk < 180 ps Pk-Pk < 148 ps Pk-Pk | (330 MBaud, no jitter BIST) (250 MBaud, no jitter BIST) (160 MBaud, no jitter BIST) |
| Rx CKR Cycle-Cycle Peak Jitter (does not include reframing CKR-stretch) | | < 100 ps < 300 ps < $0.7 \times t_B$ < 1.0 ns < 1.5 ns | (No input jitter, single data) (No input jitter, random data) (Worst-case input DDJ) (Data Phase Hop only) (Loss of Lock) |
| Rx CKR Maximum Instantaneous Offset Freq. | | < REFCLK $\pm 5\%$ | (Unstable, range control active) |
| Rx CKR Maximum Continuous Offset Freq. | | < REFCLK $\pm 0.25\%$ | (Stable, range control inactive) |
| Rx Run-Length Limit (without cycle slip) | | > $200 t_B$ > $200 t_B$ > $200 t_B$ | (330 MBaud) (250 MBaud) (160 MBaud) |
| Rx Phase Acquisition Time (BER < 1×10^{-12}) | | < $60 t_B$ < $250 t_B$ | (typical, <180 degree hop) (includes 180 degree hop) |
| Rx Frequency Acquisition Time (BER < 1×10^{-12}) | | < $50 t_B$ < $700 t_B$ | (delta-freq $\leq \pm 0.2\%$) (delta-freq > $\pm 0.2\%$) |
| Rx Closed-Loop Bandwidth (3 dB) | | 9.0 MHz 4.5 MHz 2.5 MHz | (330 MBaud) (250 MBaud) (160 MBaud) |

Table 1. Summary of HOTLink Jitter Characteristics (continued)

| Parameter | Characteristic | Condition |
|--|---|-----------|
| Rx REFCLK Re-Lock Rate (Locked to Locked) | $> 2 \text{ MHz}/\mu\text{s}$ | |
| Rx Lock Time (REFCLK Quiet to Locked) | $< 200 \mu\text{s} + 2 \text{ MHz}/\mu\text{s}$ | |
| Rx Crash Rate (from REFCLK & DATA stop) | $> 80 \text{ ps}/\mu\text{s}$ | |

- AC-coupling of input, output, and measurement signals cause unexpected problems if the wave form is non-repetitive, not DC-balanced, or if the signaling rate changes. The components used for blocking the DC voltage in the signal exhibit impedance variations because of their reactive nature. They almost always have non-monotonic transfer functions, and often have self-resonant characteristics that are not well documented. High quality DC-blocking modules from HP and other sources are typically specified to be effective over a very wide frequency range (e.g., HP 11742 Blocking Capacitor is useful at 0.01 through 26.5 GHz), but the more common “capacitor soldered on a board” is usually unsuitable for critical measurements.
- Random jitter measurements should be taken at the approximate center of the differential swing to minimize “scope arithmetic” and round-off errors that obscure actual performance.
- Bypass PECL load circuits to remove “load-ringing” effects. Power supply and PC board impedance adds directly to the impedance of the termination circuit. Careful attention to power supply bypassing minimizes load related errors.
- Simplified, high quality PECL measurements are possible using the connection shown in *Figure 48*. This is a derivative of the standard 80Ω/130Ω Thévenin termination for PECL in which the lower 50Ω of the 130Ω is provided by the scope input impedance. By using low-impedance, passive probes to maintain the full input bandwidth of the scope, and by separating the scope probes from the loads, a more representative measurement is possible. This connection yields a probe with an approximate attenuation of 2.6:1, instead of the more usual 10:1 probes. For critical voltage measurements, each such connection must be calibrated, because the actual attenuation factor will depend on the actual values of resistor used for the PECL termination. Since most AC measurements are differential and use only relative voltage levels, this connection is preferred to more expensive probe configurations. Of course, good low-capacitance layout and good quality 50Ω cables and connec-

tors are required to maintain the bandwidth of the measurement system. When the scope is not connected to the test points, a substitute 50Ω resistor should be connected to allow the PECL outputs to operate correctly.

Test Equipment

Relevant Characteristics of Measurement Equipment

Good quality, high-bandwidth, measurement equipment is mandatory to determine the actual performance of the HOTLink and the systems used with it. To gain an accurate insight into 300-MHz transmission lines, and the picosecond variations which characterize the components that define the limits of operation, it is necessary to use test systems capable of making accurate measurements up to multiple Gigahertz. The list that follows (and the short listing of their relevant attributes) are not the only applicable measurement systems, just the ones used by this design team.

HP54720D High-Speed, Real-Time, Digital Sampling Scope

Sample Rate = 8 Gigasamples/second

Trigger Jitter < 10 ps

Bandwidth = 2 GHz

1 GHz on each channel with 54721A Input module

2 GHz on single channel with 54722A Input module

This high-performance scope offers the opportunity to observe the actual wave shape with its “real-time” capability. In contrast with the more traditional sampling scope, this instrument records the signal on its inputs at 125 ps intervals until its input buffers are full. The 54720D has the ability to place the triggering event at the beginning, middle, or end of the stored waveform, which allows it to capture random and non-repetitive events.

Tek 11801A Digital Storage Oscilloscope

with SD-22, 12.5-GHz Sampling heads for precision low-impedance measurements

with SD-14, 3.0-GHz Sampling heads for low-load, high-impedance measurements

and DL-11, 5-GHz Delay Line for measurements at the time of the trigger event

Trigger Jitter < 3 ps

Bandwidth > 20 GHz, bandwidth on each channel, limited by the sampling head (SD-22 or SD-14).

This high-performance scope has sufficient bandwidth to observe the actual performance of the PECL outputs of HOTLink. Lower bandwidth scopes and probes often




Figure 48. PECL Scope Probe

give an erroneous impression of the voltage waveform being measured. The 11801A is best used for measuring repetitive waveforms, since it only accumulates a single "dot" for each trigger. Accumulated over time, this is sufficient for observing repetitive waveforms, and its color-graded histogram ability is very useful for capturing and displaying jitter performance.

HP 8560A Spectrum Analyzer

50 Hz to 2.9 GHz

Used for monitoring jitter transfer tests and various clock-source attributes to assure the accuracy of the bench set-up. The displays that appear on a spectrum analyzer are often ambiguous, since frequency, phase and amplitude variations all cause similar indications. This is a fun instrument to use, but must be interpreted with care. It usually gives more information than can be fully understood, but does offer another view of the system under test from the frequency domain.

HP 54610 500-MHz, 2-Channel Oscilloscope

This is a small, relatively portable bench scope (i.e., about one cubic foot and can be carried with one hand, in contrast to the other scopes which require a dedicated cart) used for monitoring the function of various bench set-ups and the functionality of the part under test. It has sufficient bandwidth to give an accurate picture of the circuit under test, but is too slow to give accurate results in the previously described precision tests. These scopes are typically used for setting up the various generators, clock sources, and data generators, and for cross-checking the validity of many of the measurements. They were not used to gather actual data, but offer sufficient performance to see that the set-up is working as expected.

NOTE: When using any type of sampling scope, care must be made to avoid misinterpreting the displayed information. If the data is sampled at too low a rate, aliasing effects can mask the true appearance of the measured signal.

Clock Sources

Crystal oscillators are typically used in operational systems because of their stable, predictable, low-noise characteristics (as well as their low cost). They were found to be unsuitable for the previously described tests, because of their low-frequency delay and wander characteristics. These unrepeatable effects obscure the jitter characteristic being measured. In operational systems, these effects cause *no* reduction in link performance, and are merged into the unmeasurable, insignificant, background characteristics of the system. To gather the precision information described in this application note, several clock and data sources were used. The list that follows (and the short listing of some relevant attributes) are not the only applicable clock sources, just the ones used for these evaluations.

RF Generators

HP 8656B Generator 1 kHz to 990 MHz

HP 8647 Signal Generator 250 kHz to 1000 MHz

Used as frequency reference generators because of their spectrally clean output and their high-frequency function. They generate small, ground-referenced sine waves with great accuracy and are easily programmed

from the panel or using a GPIB controller. These generators are typically used to trigger high-performance Pulse Generators, which produce the required levels and edge rates. The generators themselves have acceptable stability and jitter performance for most AC and functional evaluations, but are not sufficient for jitter related tests.

When triggered by a stable source, the jitter performance of the generator improves to almost that of the triggering reference.

Clock Generators

HP 8131A 500-MHz Pulse Generator

Pulse generators are used to generate the PECL and TTL clock and data sources for testing HOTLink products. The HP8131 can be used by itself or triggered by an RF source. It offers two independent channels, each with complementary outputs.

Wavetek 178 Function Generator

0- to 50-MHz Function Generator

The Wavetek 178 is convenient for generating low-frequency signals such as Receiver REFCLK and swept frequency-range tests. It has the capability to generate various wave shapes and can sweep its output frequency across a wide range. It has good stability and is relatively "clean," but exhibits about 200 ps of low-frequency jitter.

Colby Instruments Pulse Generator PG-1000A

The Colby pulse generator is a very stable oscillator that is mechanically tuned, and offers very good spectral purity and good control. It suffers from slight frequency drift until it is fully warmed-up. The design of the instrument is very modular, and offers many specialized controls and options to meet various voltage translation and buffering needs.

Pattern Generators

Microwave Logic GigaBERT-1400 TX

1.4-GHz maximum clock rate

< 2 ps RMS clock jitter, < 20 ps Pk-Pk

No jitter added to output when divided by N to create Bit or Byte Clock

This instrument is actually a very high quality clock generator, packaged with a bit-rate data generator. It can be used for generating bit-clock inputs without the need for an external oscillator trigger source. It was used for many of the bit-rate referenced tests described in this application note by programming it to the required pattern.

Translators and Delay Generators

Colby Instruments Custom Clock Buffer and Translator Box

This general-purpose translator box was used to convert between differential PECL and both true ECL (–5.2V referenced) and "zero-crossing" signals used in various tests. It can accept single-ended signals and return differential outputs with extremely fast edges and no appreciable increase in jitter noise. The inputs all include integrated high-quality transmission line terminators that simplify most bench configurations.




Figure 49. Duty Cycle Distortion Jitter Generator Schematic

Colby Instruments Programmable Delay Line PDL-30A

This general-purpose, mechanical delay generator is capable of generating a repeatable and stable delay up to about 10 ns in increments as small as 1 ps. It is most useful for adjusting mismatched delay lines, and for creating desired skews between various signals. It is essentially a 50Ω transmission line that can be mechanically adjusted in small increments to change the delay. It is programmable by an external keyboard with a digital readout of programmed delay.

Non-Commercial Test Equipment

Synthetic-DCD Jitter Generator

Duty Cycle Distortion (DCD) can be generated by the circuit shown in *Figure 49*. This circuit uses the stages in a 10H116 (ECL triple-differential amplifier) to perform

- Differential-PECL-input buffering
- Ramp generation
- Threshold shifting
- Level restoration
- Differential PECL output buffering

In this circuit the Transmitter data stream (or other data source) is fed through the Jitter Generator while the Receiver monitors and checks for correct operation. As the control voltage (V_j) input is varied between the 10KH V_{IL} and V_{IH} levels, the duty cycle of the data stream is impaired in a repeatable and measurable manner. Either of the V_j inputs can be independently adjusted, or they can be differentially driven to get different jitter effects.

The first differential stage of the 10H116 is used as a differential-ramp generator with controlled output impedance and symmetrical rise and fall times. The series resistor and capacitor to ground are adjusted to provide a relatively long voltage transition ramp that can be used to manipulate the edge transition timing. The ECL output termination resistors (shown at the outputs of each differential stage) are part of the normal PECL output loads, and can be either the parallel terminations shown at (a) or the single pull-down shown at (d).

The R-C ramp generator at (c) must be tuned to each data rate, to ensure that 100% voltage swing is maintained for the narrowest pulses expected. If the ramp is too long, it is possible to raise V_j above the level of some data bits, thus “losing” data.

The second differential stage of the 10H116 serves as a voltage comparator that translates the differential, artificially extended voltage-ramps back to PECL swings. The differential (or single-ended) control voltage (V_j) level modifies the restored DC levels of the AC coupled ramps. By adjusting the DC levels at the input of stage two, the average (DC voltage component) of each ramp can be independently adjusted. This adjustment moves the “crossing voltage” which the differential inputs of stage two convert to changes in the timing of the data bits. Additional DC filtering may be required between the V_j input and its input to (d) to ensure that high-frequency, single-ended noise does not corrupt the data flow.

The third differential stage of the 10H116 is used to restore crisp-edged, full-swing levels to the serial data, and to drive the subsequent transmission line. In some cases, the PECL output terminations of this stage are provided by the transmission line terminations.

Synthetic-DDJ Jitter Generator

Data Dependent Jitter (DDJ) that approximates the natural effect of long wire-transmission lines, can be generated by the circuit shown in *Figure 50*. This circuit uses the stages in a 10H116 (ECL triple-differential amplifier) to perform

- Differential-PECL-input buffering
- Ramp generation
- Threshold shifting
- Level restoration
- Differential PECL output buffering

In this circuit the Transmitter data stream is fed through the Jitter Generator while the Receiver monitors and checks for correct operation. As the control voltage (V_j) input is varied to cause variations in the “data-corruption” ramps, the data stream is impaired in a repeatable and measurable manner.




Figure 50. Data Dependent Jitter Generator Schematic

The first differential stage of the 10H116 is used as a differential-ramp generator with controlled output impedance and symmetrical rise and fall times. The series resistor and voltage-variable capacitor (c) are adjusted to provide a relatively long voltage transition ramp that is used to manipulate the edge transition timing. The ECL output termination resistors (shown at the outputs of each differential stage) are part of the normal PECL output loads, and can be either the parallel terminations shown at (a) or the single pull-down shown at (d).

The R-C ramp generator at (c) must be tuned to each data rate, to ensure that the ramp covers the same number of bits for each data rate. If the ramp is too short, the full spread of pulse-width-dependent jitter is not generated.

The second differential stage of the 10H116 serves as a voltage comparator that translates the differential, artificially extended voltage-ramps back to PECL swings. The differential restoration resistors put the degenerated waveforms at the optimal voltage so that the inputs of the receiver gate can make a proper logical translation.

The third differential stage of the 10H116 is used to restore crisp-edged, full-swing levels to the serial data, and to drive the subsequent transmission line. In some cases, the PECL output terminations of this stage are provided by the transmission line terminations.

Fiber-Optic Test Bed

The set-up used for testing fiber-optic interface capabilities of HOTLink is shown in *Figure 51*. It consists of a HOTLink Evaluation card, several lengths of fiber-optic cable, and appropriate measurement equipment.

A 3-km piece of fiber-optic cable, containing a single splice, was used to generate chromatic and modal dispersion. The shorter pieces of fiber, with two connectors between every 500 meters, and an optical attenuator, were used to add connector attenuation. The optical splitter and power meter were used to ensure repeatability of the measurements. The limits of distance and data rate were mostly set by the optical interfaces used, and by the number of connectors in the link.

Coax Test Bed

The set-up used to test wire links is shown in *Figure 52*. It consists of a HOTLink Evaluation Board with suitable connectors, and a length of the cable to be used for testing. Various cable types have been tested for data rate and distance characteristics. The HOTLink BIST function and the Evaluation Board error-count indicators combine to offer a clear and unambiguous system to determine the quality of an interconnect link, and its suitability to perform at a specified data and error rate.

HOTLink Evaluation Board CY9266-C, CY9266-T, and CY9266-F

The HOTLink Evaluation Card was designed to facilitate early HOTLink system evaluation without expensive or hard to find test equipment. These cards (shown in *Figure 53*) have convenient interfaces for user data and control signals, using either the 48-pin connector used on the IBM OLC-266 card, or a 60-pin card edge connector.

The CY7B923 and CY7B933 include an exhaustive Built-In Self-Test function that can be used to effectively test link performance. It can also be used as a controlled and predictable data source, and as a grader for received data. The receive comparator assures correct functionality of the HOTLink Transmitter, the internal logic in the HOTLink Receiver, and the interconnect link that joins them. These are the essential components of a Bit-Error-Rate tester, except for the reporting mechanism. To fill this need, the Evaluation Cards include a PLD programmed to be a two-digit accumulator and display driver. The Error Display shows the number of Error Bytes received during the BIST sequence, by counting the HOTLink RVS outputs.

BIST

HOTLink Transmitter and Receiver include a comprehensive link test function, as part of the functionality of the basic chips. When the HOTLink Transmitter $\overline{\text{BISTEN}}$ is enabled, the part creates a continuous 511 byte ($2^9 - 1$ bytes) pseudo-random stream of 8B/10B-encoded data patterns which the HOTLink Receiver checks byte-by-byte. The 256 possible data patterns are sent once each, and the 12 Special Characters and the 4 specified error codes are sent sixteen times each (ex-




Figure 51. Fiber-Optic Test Bed Facilitates Random-Jitter Testing




Figure 52. Coax Test Bed to Test for Deterministic Jitter

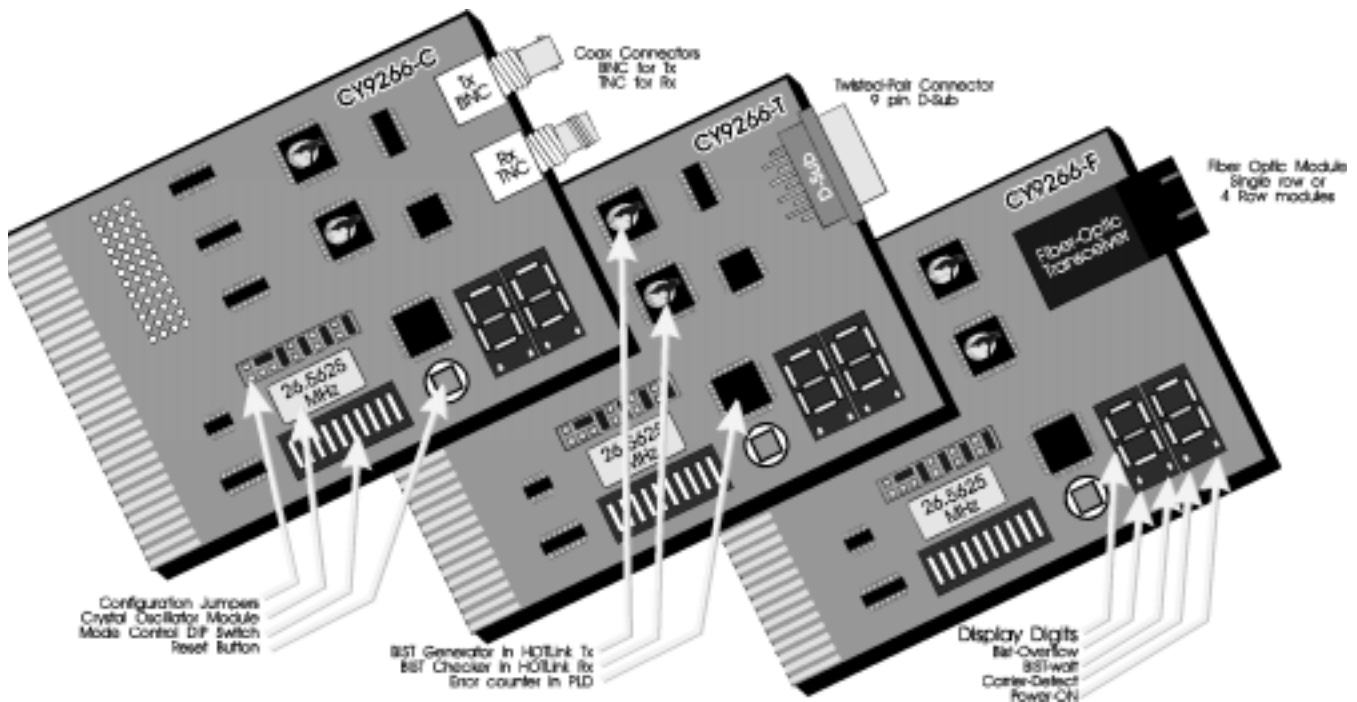


Figure 53. HOTLink Evaluation Boards Form the Core of a Comprehensive Evaluation System

cept C0.0 which is sent only 15 times) for a total of another 255 data patterns. For a complete list of codes used in the 8B/10B encoder and the Special Character and Error Codes, see the CY7B923/933 HOTLink Tx/Rx Data Sheet.

If errors are discovered in the received sequence, received running disparity, or received transmission codes, they are flagged by the RVS output of the HOTLink Receiver. A full discussion of the BIST function of HOTLink is contained in the "HOTLink Built-In Self-Test (BIST)" application note.

For Further Information

HOTLink User's Guide

Hewlett-Packard Catalog

Hewlett-Packard Test & Measurement Division
Mail Station 51LSJ P.O. Box 58199
Santa Clara, CA 95052-9943
(800) 452-4844 or (408) 553-7271

Tektronix Catalog

Tektronix
26600 Southwest Parkway P.O. Box 1000
Wilsonville, OR 97070-1000
(800) 426-2200 or (503) 627-1916

Microwave Logic
285 Mill Rd
Chelmsford, MA 01824
(508) 256-6800

Colby Instruments, Inc.
1810 14th St.
Santa Monica, CA 90404
(310) 450-0261

HOTLink is a trademark of Cypress Semiconductor Corporation.