

# An Automated, Complete, Structural Test Solution for SERDES

Stephen Sunter, Aubin Roy, J-F Côté    LogicVision, Inc.  
{sunter,aubin,jfcote}@logicvision.com,    Ottawa, Canada

## Abstract

*Gigahertz serialization and deserialization (SERDES) has become a dominant inter-chip and inter-board data transmission technique. Signal integrity is the primary factor determining its bit error rate, typically less than  $10^{-12}$ , so the primary production test challenges are testing picosecond jitter and the signal eye opening. Off-chip jitter and rise/fall time measurements are limited by hardware complexity, access, bandwidth, and noise. Published on-chip measurement techniques are limited by delay line jitter. This paper presents a new jitter test technique that has been demonstrated on an FPGA to achieve less than 1 ps RMS self-jitter, and a new signal eye test that has unlimited bandwidth; neither test uses high speed circuitry. The all-digital technique uses the receiver itself to demodulate the signal jitter to a low-speed bit stream that is analyzed by a single-clock domain, synthesizable circuit. This is combined with logic BIST and 1149.6 boundary scan to completely test an IC.*

## Introduction

High speed, serialization of digital data has proven to be a very effective way to solve important integrated circuit (IC) problems:

- High pin counts - 4~16 input/output (I/O) signals can be compressed into a differential pin pair, and fewer power pins are needed because the glitches induced on the power rail during switching are greatly decreased;
- High power at high data rates - Differential signals can have smaller amplitudes than single-ended signals, e.g. 500 mV peak-to-peak instead of 3 V, which greatly reduces switching power ( $P = f CV^2$ );
- Low  $V_{DD}$  - Lower signal swing can be tolerated for differential signals because noise is canceled when it is common to both signals of a differential pair.

Serial data communication is the most efficient way to convey data long distances by wire, fiber optics, or radio waves.


There has been dramatic progress recently in increasing the speed and quality of serial I/O signals. The maximum speed is about 6 Gb/s for two-level signals, and over 10 Gb/s for multi-level signals, on wire (and much higher on optical fiber), for normal circuit board materials (FR-4) and short distances (<1 metre). The signal integrity is best measured by bit error rate (BER). BER of commercial chips must typically be better than  $10^{-12}$  but some companies are claiming better than  $10^{-15}$ . For a BER of  $10^{-12}$ , it takes 7 minutes, on

average, to detect a single bit error at 2.4 Gb/s; a repeatable test would need to detect at least 100 bit errors. Higher data rates don't really help: for a BER of  $10^{-15}$  and a transmission rate of 10 Gb/s, it takes 2.5 hours, on average, to detect one bit error.

Clearly, the IC signal quality cannot be tested by measuring BER in production. Fortunately, the BER is primarily determined by the amount of jitter in the signal edges, and the BER can be predicted by measuring the RMS jitter, which is the sigma of its probability distribution, for the random and deterministic portions. The peak-to-peak random jitter will be less than 6 sigma for about 99.74% or  $1-10^{-3}$  of the time. Thus, a BER of  $10^{-3}$  corresponds to 6 sigma. A BER of  $10^{-12}$  corresponds to 14 sigma. At a given data rate, and no deterministic jitter, 14 times the jitter's sigma must be less than about half of the bit duration. For example, at 2.5 Gb/s, the bit duration, or unit interval (UI), is 400 ps, and so sigma must be less than  $200/14 = 14$  ps. In fact, this is the budget for jitter from all sources; for just the receiver or just the transmitter, the jitter should be less than half this, and deterministic jitter must be accounted for. Thus, the RMS jitter is typically specified to be less than 10 ps RMS or 2% of the UI.

A commercial time interval analyzer (TIA) can typically measure to an accuracy of 1 or 2 ps RMS in about 200 ms for any of 8~16 data inputs. TIA units have an input bandwidth of several GHz, so can measure rise fall times down to about 150 ps. The very latest sampling oscilloscopes have a bandwidth in the tens of gigahertz and better than 0.5 ps jitter measurement accuracy, but require up to 10 seconds to achieve this for each of 1~4 inputs. The performance of TIAs and oscilloscopes is limited by their bulk, about 2 cubic feet, which prevents them from being very close to the circuit under test, and by the difficulty of conveying an on-chip signal off chip, across a board, through connectors, multiplexers and coaxial cable, to the equipment. Variations in the path's characteristic impedance can affect the measured jitter. A final consideration is that the jitter that really matters is the jitter perceived by the serial receiver after differential buffering, including jitter in the recovered clock, and this is deduced by measuring the receiver's jitter tolerance.

While jitter and jitter tolerance are the most important parameters to test, the next most important is the signal eye opening, as shown in Figure 1, because the eye opening determines the signal-to-noise ratio.




**Fig. 1 Signal eye opening**

The eye's width and height are determined by the signal's rise and fall times, its duty cycle distortion, and its minimum voltage swing. If we can measure the jitter, rise and fall times, and duty cycle, we can calculate the width of the eye opening. We can measure the signal voltage swing at DC to determine the eye height, but it is preferable to measure this height at-speed, especially if the receiver is AC-coupled.

Different signal types have different signal voltage swings. The two most common are LVDS and CML. Low voltage differential signaling (LVDS) has 100~450 mV swing, for DC-coupled signals, between 155 Mb/s and 1.25 Gb/s. Current mode logic (CML) may be DC or AC coupled, has a 400~1200 mV swing, and requires more power than LVDS, so it is typically only used above 2.5 Gb/s.

Above 2 Gb/s, transmitters usually add pre-emphasis. The first bit of each series of consecutive same-value bits has an increased voltage swing, as shown in Figure 2. After low pass filtering by the wire capacitance and inductance between the transmitter and receiver, the received waveform will have uniform amplitude for all bits, and hence will have less data-dependent jitter. Pre-emphasis does not affect the average voltage of a high-speed random data signal.




**Fig. 2 Waveform showing pre-emphasis**

The number of high speed differential pin pairs on individual ICs continues to grow. Chips have been reported with more than 64 such signal pairs. A scalable solution needs to be able to test many of these signals simultaneously, unless the test time for each channel is less than, say, 100 ms.

Although DC coupling is still common, AC coupling is becoming the favored way to connect high frequency (HF) signals because it provides voltage isolation to allow the driver and the receiver to each operate in their preferred voltage range. AC coupling requires the data to have a constant DC average (an equal number of 1's and 0's in any short interval) and this is typically accomplished with 8B/10B encoding. The

popularity of AC coupling, and the inability of 1149.1 boundary scan to test it [11], recently led to the development of IEEE 1149.6 boundary scan [12], also known as AC-EXTEST or AC-JTAG.


Figure 3 shows a common receiver termination and coupling scheme. The termination resistance is typically 100 Ω (possibly comprising two 50 Ω resistors connected to a bias voltage) because each leg sees a characteristic impedance of 50 Ω which is the most common impedance used in printed circuit board design. The coupling capacitor is typically 10 nF or larger if it is on the transmitter side of the termination resistor. If the capacitor is on the receiver side as shown, the biasing resistors can be 10 kΩ (and their value is non-critical) so the capacitor can be 50 pF or less. This allows the capacitors to be on-chip to considerably reduce cost (about \$0.50 per capacitor). These on-chip metal-metal capacitors occupy very little area (each less than a bond pad).



**Fig. 3 AC coupling, termination/bias resistors**

A phase-locked loop (PLL) is typically used for clock and data recovery (CDR), as shown in Figure 4. Typically, after power-up, the receiver is in Initialization mode and the PLL locks to a reference frequency,  $f_{RXREF}$ . The reference frequency is typically required to be within 500 ppm of the incoming data frequency. After phase or frequency lock is achieved, the receiver locks to incoming data. If the PLL ever loses phase-lock, the receiver returns to Initialization mode, re-locks to the reference frequency, and then re-locks to the data.

For frequencies higher than 2 Gb/s, the phase error of the receiver in Figure 4 can become excessive, and a different design is typically used – a bang-bang CDR [1]. However, some multi-channel, high speed applications require the CDR to phase-lock to a clock derived from another channel, so in this case, phase error can be tolerated if it is constant and less than a few bits.



**Fig. 4 Typical receiver design**

Of course, it would be best to measure jitter on-chip, but no silicon results have ever been reported for a built-in self-test (BIST) circuit that achieves better than 10 ps RMS accuracy. In our experience, corroborated by the hardware results in other researcher’s publications, this is an almost technology-independent noise floor, and is due to jitter inherent in the BIST circuit’s delay line [2] or Vernier oscillators [3][4], or both [6], in addition to coupling to the PLL via power rails and the substrate.

Very few on-chip techniques have been reported for accurately measuring the signal eye opening of gigahertz signals [14]. One paper [5] reported a technique for measuring rise and fall times, but its accuracy was limited by voltage-dependent delay through the signal buffer and by jitter in the signal – the peak-to-peak jitter can be comparable to rise/fall times in gigahertz signals.

In the rest of this paper, we will describe a new on-chip jitter measurement scheme whose measured noise floor is below 1 ps RMS and improves with technology. We will then describe a new unlimited-bandwidth technique for measuring the received signal eye opening. We’ll show how these techniques have been combined with LFSR-based testing of the mux/demux structure, logic BIST of the sub-gigahertz logic, and 1149.6-based testing of board-level interconnect, to form an automated, complete test solution for SERDES circuits.

### New on-chip jitter measurement circuit

The jitter test will first be described for a 1010 transmitted serial data stream. Parallel input data is sampled by a transmit clock,  $f_{TxCLK}$ , and transmitted serially at N times this rate. To generalize, the data is serially transmitted at a data rate  $f_D$ , as shown in Figure 5. This data is normally sampled by the receiver at an exactly equal rate using a clock recovered from the data by the receiver’s PLL. For the new jitter test, we under-sample the data at a rate slightly lower than the transmitted data rate or at a rate slightly different from an integer fraction of the transmitted data rate.

For example, if the data is sampled at the data rate minus 1%, then the resolution of the jitter test is 1% of the UI, regardless of the frequency. If the data is sampled at a frequency that is an integer fraction minus

1%, once every N bits, then the resolution of the jitter test will be  $N \times 1\%$ . To achieve 1% UI resolution for this case, the data needs to be sampled at the data rate divided by N, minus 1% divided by N. Typically, 0.2% resolution is sufficient to measure jitter that is 1~2% of the UI. Therefore, for the most common serialization rate, 10 bits per word ( $N=10$ ), the frequency difference needs to be about 0.02% (200 ppm).

The output of the sampling latch is conveyed to a jitter analysis circuit which calculates the jitter properties. The analysis circuit, which is a single-clock finite state machine, monitors the signal samples and, when a transition is detected after many same-value bits, captures the unstable bit regions that are caused by jitter. As can be seen in Figure 5, the sampling latch output is a mostly low frequency (LF) waveform, whose lowest fundamental frequency is a “beat” frequency equal to the difference between the sampling frequency and the nearest integer fraction of the data frequency.

It is important to note that the sampling clock can be generated by the receiver’s own PLL whose input is a frequency-offset reference clock from a stable off-chip oscillator. The sampling latch can be the receiver’s latch with no additional latch added.

For example, to test a 1 Gb/s, 10-to-1 SERDES, we supply a 100 MHz clock to the transmitter, whose PLL generates 1 GHz. The receiver’s PLL phase-locks to a 99.98 MHz external reference clock, and the received 1 Gb/s data is latched by the receiver using the 999.8 MHz generated by its PLL. The deserializer conveys every 10<sup>th</sup> sample to each of its ten parallel outputs (with any decoding bypassed), and the jitter analysis circuit is connected to one of these outputs. The receiver’s reference clock (99.98 MHz) is generated by digitally dividing the 100 MHz source clock frequency by 5,000 to

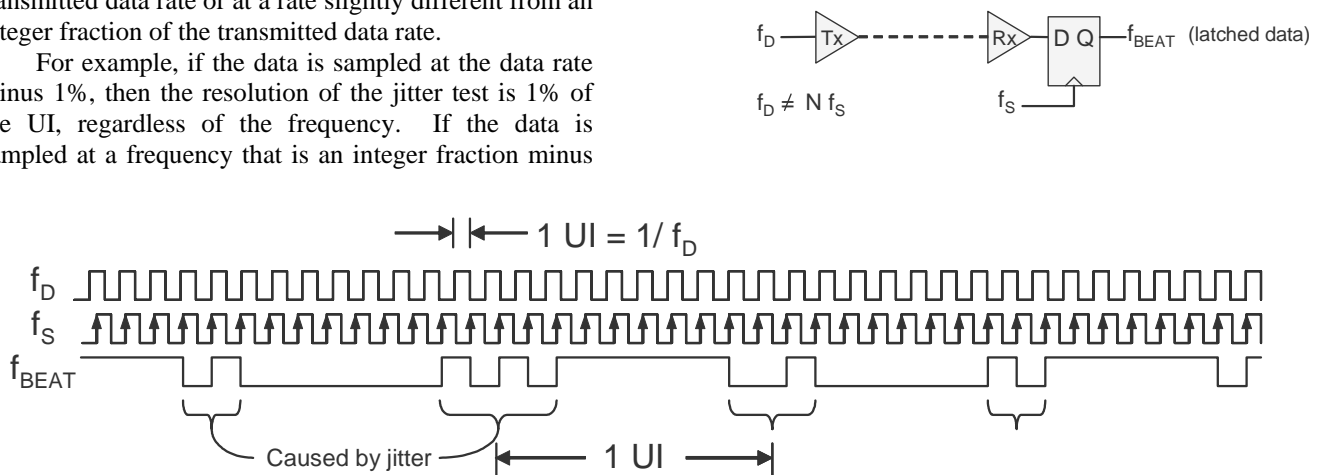


Fig. 5 Operating principle of new jitter test: circuit and simplified waveforms

obtain 20 kHz, and then multiplying the 20 kHz by 4999 using a voltage-controlled crystal oscillator (VCXO) -based PLL on the loadboard. As described in [7], this is both a low jitter and low cost way to generate a clock. The beat frequency is 20 kHz, so 1000 beat frequency edges are sampled in 50 ms with 0.2% UI measurement resolution before averaging. At 10 Gb/s, the measurement time reduces to 5 ms.

The jitter analysis circuit measures the RMS jitter using a cumulative distribution function (CDF)-based algorithm [2] which is accurate for jitter with an approximately Gaussian probability distribution function (PDF). It does this by super-imposing each group of unstable bits and adding the logic 1's in each time slot. The circuit measures either high or high-low frequency relative timing jitter, where the separation between and high and low frequency is equal to approximately half the beat frequency, as detailed in [13]. High frequency jitter is measured by aligning the medians of each captured group of unstable bits. Wide band jitter is measured by aligning the groups according to each group's position relative to the nearest edge of the average beat period. The circuit also simultaneously counts the number of unstable bit transitions to verify that the jitter appears random. As the RMS jitter value is shifted out, it is simultaneously compared on-chip with shifted-in test limits, so that pass/fail bits can be output too. The circuit can also output a histogram of the jitter to allow graphical display and off-chip calculation of separate RMS values for random and deterministic jitter.

The only process-sensitive element of this jitter measurement circuit is the sampling latch (or D-type flip-flop), which is preferably the receiver's latch as mentioned, so we want its noise included in the measurement. In other words, this jitter measurement scheme measures the jitter as perceived by the receiver, which off-chip measurements cannot do directly. Importantly, this also means that the entire circuit can be accurately simulated in a digital simulator. For 1000 beat periods, our logic simulations showed CDF-derived RMS measurements are within 3.2% of the true value for approximately Gaussian distributions.


Many off-the-shelf SERDES ICs permit the receiver's PLL to remain phase-locked to a reference clock while sampling data (for high-speed parallel data reception when the clock is derived from one of the channels) – this permits use of our jitter measurement method without any modifications to the SERDES high frequency circuitry. The jitter analysis circuit can process any one of the low-speed parallel data outputs from a channel – all of the outputs will contain the same data bits (stable and unstable bits) but at different phases of the beat frequency.

The data pattern sent by the transmitter for this jitter test can be any periodic sequence, preferably less than or equal to 2N bits (N is the width of the serializer's multiplexer). However, to obtain more-representative jitter, the data pattern must be a pseudo-random bit stream (PRBS). If we undersample a random pattern, we will not obtain a region of stable logic 1's (or logic 0's) to delineate the unstable regions caused by jitter. The solution is to use a slightly modified PRBS in which two adjacent bit positions in every transmitted word are kept constant at 01 (or 10) when measuring rising (or falling, respectively) edge jitter.

Transmitted jitter can also be measured independently by sampling with the  $f_{R_{xCLK}}$  directly – VCXOs with period jitter of 1 ps RMS are commonly available.

To verify the jitter floor of this technique, we measured the jitter of a crystal oscillator using a voltage-controlled crystal oscillator as the reference clock. Specifically, the data for a Xilinx Virtex™ FPGA D-type flip-flop was the output frequency of a 27.0000 MHz crystal (UI = 18.5 ns), and the clock was the output of a VCXO (Epson VG-1011JA) adjusted to 27.00073 MHz, to give a resolution of 1 ps. The high frequency relative timing jitter for 1000 beat frequency edges was less than 1 ps RMS when there was no other active circuitry.

We also verified the performance at 3.11 Gb/s for a SERDES transceiver of an Altera Stratix GX™ FPGA, and obtained the histograms shown in Figure 6. The transmitter's reference was a 155.52 MHz crystal, and the receiver's was a PLL connected to a 33.333 MHz crystal. The PLL multiplied its input frequency by 14/3 to obtain 155.554 MHz, which is 0.022% more than the transmitter's reference frequency.




**Fig. 6 Measured HF timing jitter histogram for 1000 beat periods; loopback via 30 inches of coax**  
a) Modified PRBS  
b) Modified PRBS, with 15% pre-emphasis  
c) Word = 1111100000, with 15% pre-emphasis (650 MHz Wavecrest DTS2070 measured 2.5±0.5 ps)

## New transition time measurement circuit

Jitter measurement alone does not reveal the signal integrity. The amplitude of the signal and its rise and fall times are a necessary part of the picture.

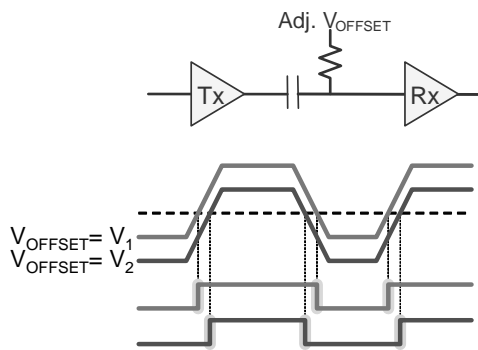
The technique described in [5] measured transition times by measuring the received duty cycle for two different reference voltages of input single-ended comparator, as shown in Figure 7.



**Fig. 7** A previous circuit [5] for rise/fall time measurement

The circuit of Figure 7 has several problems – the delay through the comparator depends on the reference voltage, it is single-ended (not differential), and the analysis technique is sensitive to jitter. The solutions to these problems are to exploit the AC-coupling that is so common in SERDES applications, and to use a jitter-insensitive analysis algorithm.

The circuit of Figure 8 illustrates the principle of our approach. The input signal is AC-coupled to an input buffer which has a constant switching point voltage. For an incoming periodic waveform, the bias voltage controls the voltage level, relative to the waveform, at which the limiter/buffer switches. Increasing the bias voltage raises the waveform relative to the constant switching point, and hence decreases the switching point voltage relative to the waveform. For a non-zero rise/fall time, the output duty cycle of the undersampled data bits will indicate the apparent width of the logic 1 pulse. By measuring the duty cycle for a bias voltage that shifts the waveform down by 30% of its amplitude, and then for a bias voltage that shifts it up by 30%, the 20~80%




**Fig. 8** Single-ended circuit for rise/fall time measurement, and waveforms

rise fall time (as percent of UI) is calculated as follows, for a 1010... input pattern:

$$t_{\text{RISE+FALL}} = \text{dutyCycle}_{80\%} - \text{dutyCycle}_{20\%}$$

The circuit of Figure 9 permits injection of a differential offset voltage that does not change the common-mode voltage. A tester's parametric measurement unit (PMU) sources a current and another PMU sinks an equal current. The injected offset is equal to the current multiplied by the termination resistance (e.g., 100  $\Omega$ ; its measurement is discussed later). For example, 1 mA will inject a 100 mV differential offset. A relatively high value resistor (e.g. 1~5 k $\Omega$ ) between each PMU and the high frequency pin will reduce the transmission line effects that connecting the PMU might cause.



**Fig. 9** Differential circuit for rise/fall time measurement

The duty cycle of the received waveform, with or without an injected offset voltage, is measured by counting the number of logic 1's in the undersampled data (of one of the parallel outputs) from the latching circuit while the frequency-offset clock is applied. The number of 1's is divided by the total number of sampling clock cycles in the measurement interval (preferably an integer number of beat periods). This measurement is jitter tolerant – jitter only introduces unstable bits (Figure 5) which comprise on average an equal number of 1's and 0's.

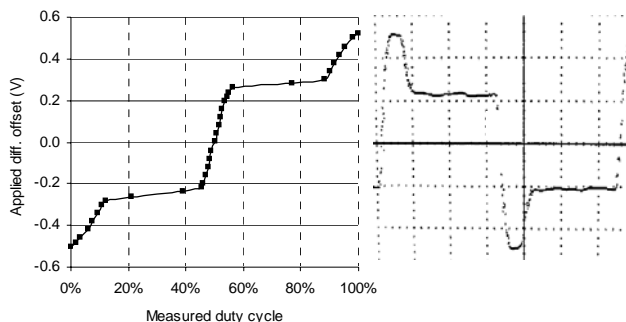
The measurement resolution of this duty cycle test, using the previous frequency offset calculations for jitter, will be 0.2% of the UI, for any frequency. The measurement has no bandwidth limits, unlike external equipment, other than limits imposed by the receiver itself.

For a 1 volt signal swing (500 mV single-ended), and a 100  $\Omega$  termination impedance, we must inject 3 mA to obtain a 30% or 300 mV differential offset voltage. If 1 k $\Omega$  resistors are used to isolate the PMU, the voltage drop across them will be 3 volts, which may be above the power rail voltage but it is almost impossible for it to cause damage to a defect-free device.

If the AC-coupling is on-chip, then PMU access to the receiver needs to be accomplished on-chip. There are several ways to implement this. An on-chip analog bus can deliver the PMU currents to each receiver input – a bias resistance of greater than 10 k $\Omega$

means that less than 300  $\mu\text{A}$  is needed to inject a 300 mV offset. To ensure linearity and minimal capacitive loading of the receiver inputs, the analog bus should only be connected to the receiver inputs via a resistance (for example, 10 k $\Omega$ ). An alternative, simpler, and faster way to inject a precise offset voltage is to use a separate bias voltage bus for non-inverting and inverting inputs of the receiver. This avoids delivering any current to the receiver input, permits all receivers to be offset simultaneously, and does not rely on knowing the exact value of the termination resistance. However, it might require an extra IC pin to allow a PMU to drive each reference bus.

We verified the method by comparing the measured duty cycles for various waveforms. In the waveform of Figure 10, the average rise/fall time was calculated to be 95 ps based on the duty cycles, and the measured times were 80 ps and 110 ps (average = 95 ps). Other patterns did not give such accurate results.



**Fig. 10 Duty-cycle measurements and 3.1 Gb/s waveform (word = 1111100000) from TDS8000 DSO with 20 GHz bandwidth; 500 ps/div., 200 mV/div.**


To reconstruct the signal eye opening, as perceived by the receiver, we need to measure the duty cycle distortion (DCD) while transmitting a 50% duty cycle bit pattern, and measure the duty cycle while injecting +30% and -30% voltage offsets. We also need to know the absolute value of the voltage swing – the height of the eye. We shall describe that measurement technique next.

### Measuring received signal amplitude

If we increase the injected offset voltage sufficiently, the offset will exceed the amplitude of the received signal. Thus, we can inject a positive differential offset that is equal to the minimum acceptable positive voltage swing and verify that the measured duty cycle for a 1010 transmitted pattern (measured with an offset sampling frequency) is, for example, greater than 25%. Similarly, we can inject a negative voltage offset and verify the received duty cycle is less than, for example, 75%.

The procedure described will measure the amplitude as perceived by the receiver. Any offset in the receiver will be detected as an inadequate logic 1 or logic 0 voltage swing, or as duty cycle distortion.

If there is pre-emphasis or equalization present in the received waveform, its amplitude can also be deduced by the measured duty cycle at a large enough offset voltage, as seen in Figure 11 (and Figure 10), measured for 11001100 pattern at 3.1 Gb/s.



**Fig. 11 Duty cycle measurements for different amounts of pre-emphasis and equalization**

A key aspect of testing a high speed receiver, that we have not mentioned yet, is delivering a gigahertz data rate signal to the receiver. The most commonly used technique to test a transceiver is external loopback: the output of an IC's serializer is connected via wiring on the device-tester interface board to the input of the IC's deserializer. Sometimes there is additional low pass filtering and/or attenuation. Using this connection, some companies simply verify that the SERDES circuit can transmit/receive a few million pseudo-random bits without any bit errors, to avoid using a high frequency tester and to save test time. Of course, a BER of  $10^{-12}$  can hardly be verified with  $10^7$  bits, but if the jitter and signal eye tests just described in this paper are also used (or used instead), then this is quite feasible.

Loopback of the high speed serial data is usually implemented also on-chip, connecting from the output driver's input to the post-receiver input. At least one commercially available IC (from Xilinx) also loops the serial signal from the output bond pad to the input bond pad, thus including the effects of output amplitude, pre-emphasis, and equalization.


A classic short-coming of any loopback test is that it does not test the transmitter and receiver separately, so compensatory defects might escape detection. For example, if the receiver had poor sensitivity (low gain), but the transmitter delivered an excessive voltage swing, the loopback test would pass (however, it is important to note that Rx+Tx RMS jitter will always exceed either individual value, so it is practical to just measure the sum). We shall next describe several ways to test the transmitted signal amplitude independently.

## Measuring transmitted signal amplitude

Obviously, the simplest way to test transmit logic levels is to measure them at DC, while the driver is delivering a steady-state logic 1 or logic 0 and driving a 100 Ω differential load. 1149.1 or 1149.6 boundary scan facilitates very simple control for this test. However, for DC-coupling the voltage swing that the driver achieves at 5 Gb/s will likely be less than the DC swing due to the load capacitance, and for AC-coupling the at-speed swing will always be less. Nevertheless, the DC test may be sufficient if the at-speed signal eye is verified using the duty cycle-based test described previously.

Another way to test the at-speed voltage swing is to transmit a periodic sequence having a known duty cycle, measure its average voltage, and then transmit a different periodic sequence having a different known duty cycle, measure its average voltage, and then compute the signal swing based on the two voltages. 8B/10B encoding must be bypassed for this test.

To minimize transmission line effects of PMU access to the transmitter, high-value isolation resistors should be used (as described for the receiver), as shown in Figure 12.



**Fig. 12** Differential circuit for transmit level measurement

For example, the 10-bit word 1100011000 ( $N=10$ , and  $M_1=4$  ones) is transmitted continuously, and the average voltage is measured:  $V_1 = 2.0$  volts. Next, the word 1110011100 ( $M_2=6$  ones) is transmitted continuously, and the average voltage is again measured:  $V_2 = 2.1$  volts. The voltage swing is:

$$\Delta V = N (V_2 - V_1) / (M_2 - M_1) = 500 \text{ mV}$$

Other pattern pairs, such as 1000010000 and 1100011000 can also be used to verify that isolated ones (or zeroes) achieve full amplitude.

The PMU end of the resistors can be connected to capacitors to provide low pass filtering, but some PMUs integrate over a sufficiently long time interval to achieve this filtering inherently.

The voltage measurements can be performed by measuring each of the single-ended signals and subtracting the two results to compute the differential results, or measuring the differential voltage directly – both cases gave the same results in our hardware experiments because all noise is inherently cancelled.

The test can be performed with/without pre-emphasis to ensure that the pre-emphasis does not affect either average voltage, nor the computed voltage swing.

A pair of on-chip analog buses could be used to measure these DC voltages, one pin-pair at a time. As also discussed for the receiver, isolation resistors are needed to ensure linearity (CMOS transmission gates have a notoriously non-linear resistance versus bias voltage) and isolation.

## Measuring input/output resistance

To ensure minimal transmission line reflections and jitter, the resistance of the transmitter and receiver must match the transmission line. Some ICs rely on off-chip termination resistors, however, for optimal GHz performance the resistors are typically on-chip, and their values should be measured. For LF input/output pins, measuring resistance is only a matter of connecting the tester's PMU to the pin and measuring current for an applied voltage. However, if PMU isolation resistors have been included so that the transmission line can go directly from the transmitter to the receiver, without any open circuit stub to the PMU, then these resistor values must be subtracted.

If the isolation resistors have a value of 1000 Ω  $\pm 0.1\%$ , then the maximum error introduced in measuring 100 Ω, is 2 Ω, or 2% (two of the 1 kΩ resistors are involved). Inductors with good GHz quality could be used for access: their series resistance will be less than an ohm so the inaccuracy will be less.

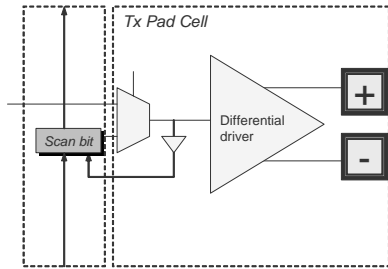
It is interesting that AC-coupling allows these DC measurements to be easily performed. DC-coupling would make it more difficult to measure the transmitter and receiver resistances independently.

## 1149.6 boundary scan


Defects in board-level differential signals are difficult to diagnose with 1149.1 boundary scan, and AC-coupling further decreases the testability. As mentioned earlier, 1149.6 was developed as an extension to 1149.1 to address these types of signals.

To make an 1149.6-compliant transmitter boundary scan cell, connected as shown in Figure 13, it is necessary to add an exclusive-or gate and multiplexer so that the TCK clock, divided by two, can be delivered to each pin. A multiplexer that selects mission mode or test mode data is also needed (as in 1149.1). Inserting a multiplexer in the high-speed bit-stream is practical for several gigabits per second, with custom layout, but at faster rates it has proven more practical to insert the multiplexer in the low speed parallel data path, implemented as N, 2-to-1 multiplexers as shown in Figure 14. This allows all of

the boundary scan logic for the transmitter to be synthesized and laid out automatically.




**Fig. 13** 1149.6 boundary scan for transmit



**Fig. 14** Modification to use low speed data

For a differential receiver, each pin (and optionally the mission receiver output) requires a boundary scan cell bit, as shown in Figure 15, and a special Test Receiver (TR) comprising a hysteretic comparator and possibly an RC low pass filter – if the signal is always AC-coupled, then the filter is not needed. The hysteresis of each comparator is initialized by its boundary scan bit. Aside from the Test Receivers, which are directly connected to the receiver pins, the rest of the receiver’s 1149.6 boundary scan logic can be synthesized and laid out automatically because it does not handle GHz serial data.

For many SERDES devices, the parallel “low speed” data is, in fact, a fairly high frequency. For example, 8-to-1 or 16-to-1 SERDES devices that transmit at 5 or 10 Gb/s have parallel inputs that each operate at 311 or 622 Mb/s. For 155 Mb/s and higher, the I/O pins are typically LVDS, and hence 1149.6 is also needed for these pins.



**Fig. 15** 1149.6 boundary scan for receive

## A total solution

Software has been developed to synthesize logic to implement the new tests described here, integrated with existing structural tests such as logic BIST, memory BIST, PLL BIST [2], and 1149.1 boundary scan, as shown in Figure 16. All tests are controlled via the 1149.1 test access port (TAP) controller.

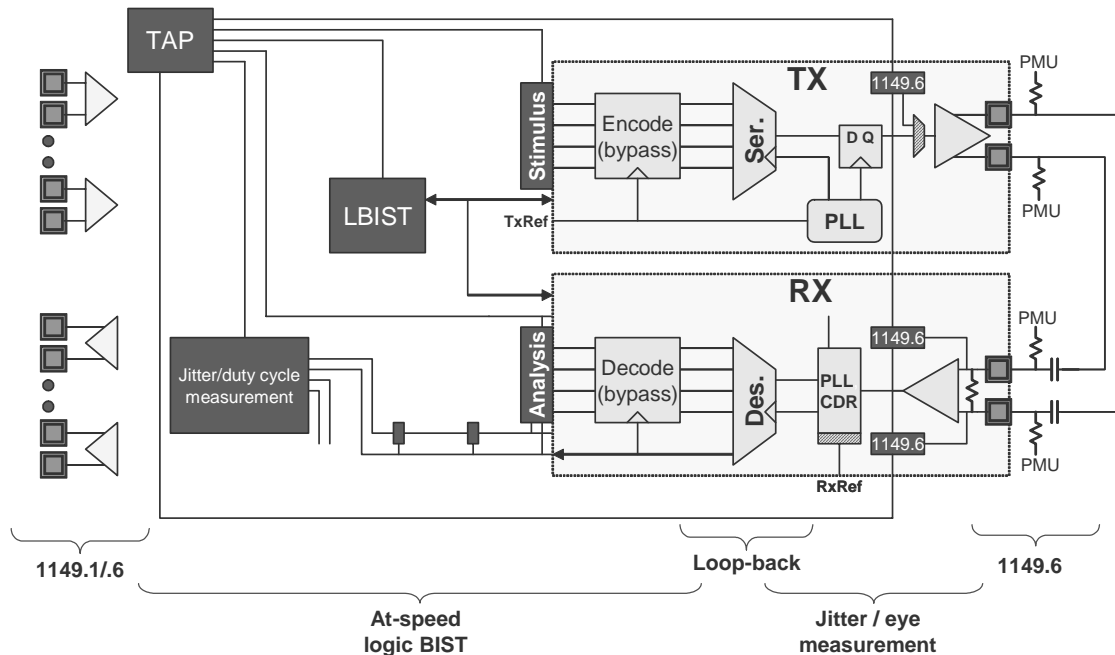
The logic that interfaces to the multiplexer of the serializer and the demultiplexer of the deserializer can be implemented within a “hard” macro. This low speed logic permits structural testing of the high-speed circuitry alone, without involving the low-speed functions, to improve diagnosability. For example, the transmitter logic (“Stimulus” in Figure 16) provides optimized stimulus patterns for each test. The receiver logic (“Analysis”) analyzes the response for some of these patterns, and the jitter/duty-cycle block analyzes the other responses. Loopback, as shown in Figure 16, is handled automatically, but interfacing to “golden” devices (e.g., FPGA) is also possible.

All other logic can be synthesized as part of soft macros. The jitter and duty cycle analysis circuitry, about 5K gates, may be shared among a group of SERDES transceivers with pipe-lining of the jitter data when the across-die transit time warrants it. There may be multiple analysis blocks on a single IC to reduce test time or to measure the impact of cross-talk. The jitter measurement circuit can also test jitter of the PLL alone, and any unrelated PLLs.

## Discussion

The jitter measurement principle described here is not completely new. A few prior publications (including patent applications) have used variations of this principle to measure peak-to-peak jitter using a latch or tester channel, or to measure the jitter modulation using analog circuitry that demodulates using conventional sine and cosine multiplication. In a recent ITC paper [8], the authors used the tester’s strobing frequency (100 MHz) to sample a periodic output from the DUT having a slightly different frequency (99 MHz). By recording the first fail and last fail times as a “sliding delta clock” moved by the sampled signal edge, the peak-to-peak jitter could be deduced. Of course, this is susceptible to jitter in the tester’s clock, and may not measure the jitter perceived by the IC. More importantly, it only measures peak-to-peak jitter. To estimate BER, it is the RMS jitter that is vitally important. The other approach, on-chip demodulation, requires analog circuitry, which introduces more circuitry to be tested, only tests LF jitter, and requires on-chip RMS jitter calculation using squaring and square roots.

Note that it is only HF jitter that causes bit errors because the receiver’s PLL can track LF jitter.



**Fig. 16** Components of the complete solution and their test coverage

Some ICs have unequal numbers of transmitters and receivers, so loopback testing is not feasible. In this case, a golden device is often used to provide or capture the high frequency serial data. A variety of manufacturers sell high speed FPGAs that are well-suited to this task, especially because they can include the necessary protocol handling and because they are highly programmable – the SERDES test scheme described in this paper is well-suited to this setup. The transmit level and jitter of the golden device may need to be re-measured for every DUT because they will be affected by the load and temperature. Many off-the-shelf SERDES transceivers permit the receiver to be phase-locked to an independent reference clock while sampling data, and can thus measure the jitter using the method we described.

The most important contributor to gigahertz IC performance is usually the circuitry connected to the gigahertz I/O pins – it contributes transmission line effects, loading, and noise. The techniques described in this paper enable all parameters to be tested with minimal external circuitry connected to the pins, and, for the previously mentioned IC that has on-chip loopback at the bond pads, no external connections are needed. Multi-chip module (MCM) or system in package (SiP) assemblies can be tested without connection to the high speed serial channels.

In one study [9], it was calculated that for RAMBUS® memories, it would cost \$1.1M per picosecond of measurement error when testing a hundred million 800 Mb/s memories per year. When

measuring 2 ps RMS jitter for 6 Gb/s SERDES, achieving sub-picosecond accuracy might become even more important. The time to market for these devices is imperiled because picosecond error can either greatly reduce yield or increase the number of test escapes, and neither is acceptable.

Presently, many companies do not test for jitter and/or jitter tolerance, due to its complexity. These companies rely on design margin, as verified in characterization, to obviate this testing. However, the design margin comes at a cost – extra power and IC area, and the possibility of delivering defective ICs.

To completely test a SERDES device, it is necessary to test the multiplexer/demultiplexer, the 8B/10B encoder and decoder, protocol logic, and the lower frequency I/O pins. Most SERDES devices already include pseudo-random pattern generation (PRPG) and BER detectors. When combined with multiple loopback paths, as done in many commercially available ICs, this circuitry will test much of the logic with block-level diagnosability. To provide gate-level diagnosability of the lower speed circuitry, structural test patterns and scan-based access are needed. A recent ITC paper [10] discussed why PRPG-based loopback testing is inadequate for much of the SERDES logic – it did not address the frequency drift recovery circuitry, for example. The offset frequencies needed for our new jitter test are well-suited for functionally testing this circuitry.

The jitter measurement is suitable for measuring jitter at any frequency, including low frequencies, for

any clock or data signal, and only requires a quiet (preferably differential) latching element. The latch's output can also be subjected to Fourier Transform analysis to measure the jitter's spectral content.

The jitter and duty cycle measurement techniques described here are reusable at the system level, including fiber optic transmission systems. For these systems, each transceiver board is typically required to have a high stability crystal oscillator having a frequency within 100 ppm of the target. The slight difference in frequency between a distant transmitter and the local receiver can be exploited by this jitter test technique, by phase-locking the receiver to the local crystal. The beat frequency, which may be only a few hundred hertz, is sufficient to measure the jitter with picosecond accuracy. In fact, this is very similar to the experiment that we described earlier in this paper (with the two crystal oscillators). Voltage offset can be injected by using two on-chip voltage references during test mode, as described earlier, one for the inverting and one for the non-inverting inputs.

### Future Work

Other parts of the complete solution, which we have explored successfully, include measuring the parameters for multi-level signals, direct measurement of receiver jitter tolerance, cancellation of sampling clock jitter, measurement of rise and fall times separately, and measurement of low frequency phase modulating waveforms (when measuring LF "jitter" caused by spread spectrum clocking, the Rx/Tx frequency offset must be greater than the modulation frequency). All of these measurements exploit the techniques presented in this paper, and have similar accuracy in theory, but have not yet been verified experimentally in hardware.

High frequency jitter tolerance can be determined more directly than using off-chip techniques: the median sampling edge time is measured relative to the median data edge time, with a duty cycle measurement. This technique requires access to a GHz signal, and so could not be verified on the FPGA – we plan to do this on an ASIC. The PLL's low frequency tracking can be measured using the PLL loop gain test described in [2].

### Acknowledgments

The authors would like to thank Pierre Gauthier and Denys Richard for automating 1149.6 boundary scan logic generation and test pattern generation, and Benoit Nadeau-Dostie for developing new structural tests for the multiplexer/demultiplexer.

### Conclusions

We have presented new techniques (patents pending) for measuring LF and especially HF jitter,

and the signal eye opening, using only low frequency digital and DC analog circuitry. The jitter floor can be zero when the receiver's sampling latch is used, or less than 1 ps RMS when a D-type flip-flop is added, which we verified in an FPGA. All complexity is transferred to a process-insensitive digital analysis circuit which is a low data rate, single clock, finite state machine (that can itself be tested thoroughly by scan or logic BIST). The signal's jitter, duty cycle distortion, and rise/fall times can be measured with almost unlimited time-resolution accuracy at any operating frequency, from kHz to GHz. When combined with BIST of the low speed circuitry and 1149.6 boundary scan of the I/O pins, these techniques facilitate an automated, complete test of a SERDES function with minimum pin-count access.

### References

1. J.Alexander, "Clock recovery from random binary signals," *Electronics Letters*, vol. 11, p. 541, 1975.
2. S.Sunter, A.Roy, "BIST for Phase-Locked Loops in Digital Applications", *Proc. of ITC*, 1999, pp. 532-40.
3. S.Sattler, D.Oberle, J.Eckmüller, "PLL Built-In Self-Test Jitter Measurement Integration into 0.18u CMOS Technology", *Proc. of Test Methods and Reliability of Circuits and Systems Workshop*, 2001
4. A.Chan, G.Roberts, "A Synthesizable, Fast and High-Resolution Timing Measurement Device Using a Component-Invariant Vernier Delay Line," *Proc. of ITC*, pp. 858-867, 2001
5. J-L.Huang, K-T.Cheng, "An On-Chip Short-Time Interval Measurement Technique for Testing High-Speed Communication Links", *Proc. of VLSI Test Symp.*, pp. 380-385, 2001
6. P.Dudek, S.Szczepanski, J.Hatfield, "A High Resolution CMOS Time-to-Digital Converter Utilizing a Vernier Delay Line", *IEEE J. Solid-State Circuits*, vol. 35, pp. 240-247, Feb. 2000
7. J.Ferrario, R.Wolf, S.Moss, "Architecting Millisecond Test Solutions for Wireless Phone RFIC's", *Proc. of ITC*, pp. 1151-1158, 2002
8. G.Bao, "Challenges in Low Cost Test Approach for ARM9™ Core Based Mixed-Signal SoC DragonBall™-MX1", *Proc. of ITC*, pp. 512-519, 2003
9. W. Dalal and S. Miao, "The Value of Tester Accuracy", *Proc. of ITC*, pp. 518-523, 1999
10. G.Hetherington and R.Simpson, "Circular BIST testing the Digital Logic with a High Speed SERDES", *Proc. of ITC*, pp. 1221-1228, 2003
11. S.Vandivier, M.Wahl, J.Rearick, "First IC Validation of IEEE Std. 1149.6", *Proc. of ITC*, pp. 632-639, 2003
12. IEEE Std 1149.6-2003, *IEEE Standard for Boundary-Scan Testing of Advanced Digital Networks*
13. S.Sunter, A.Roy, "On-Chip Digital Jitter Measurement, from Megahertz to Gigahertz", *IEEE Design & Test of Computers* [to be published July/August 2004]
14. B. Casper et al., "An 8-Gb/s Simultaneous Bidirectional Link with On-Die Waveform Capture," *IEEE J. Solid-State Circuits*, vol. 38, no. 12, Dec. 2003, pp. 2111-2120