JAZIO Enhancements To Achieve Very High Data Rates Per Pin For Parallel Interfaces

I. INTRODUCTION

JAZiO technology is a digital signal switching technique that uses differential sensing but requires only a single pin per signal. [1], [2], [3]. It is based on detecting change or no-change from the previous state instead of the traditional high or low voltage compared to a reference. JAZiO is suitable for low latency transfers for high-speed inter-chip or intra-chip signal transfers. Signal transfer rates of greater than two gigabits/second/pin are achievable at low cost using standard semiconductor, package, and printed circuit board (PCB) technologies. The technology is easily applied to DRAMs (Memory Bus), SRAMs (Back Side Bus), CPUs (Front Side Bus), Network Processors, ASICs, Intra-Chip buses, etc.

JAZiO is targeted for parallel transmissions at high data rates and offers superior performance per pin compared to other known solutions. Parallel transmission is a data transmission in which a group of bits are moved over several channels or wires simultaneously. Data transmission where the bits move over a single channel, one after the other, is referred to as serial transmission. Semiconductor devices that are located near one another normally use parallel transmission because it is much faster. As distances between devices increase beyond one meter parallel transmission becomes increasingly difficult and serial transmission is typically utilized for transmission over longer distances. This paper concentrates on the effects of high frequency; such as inter symbol interference (ISI) jitter, signal attenuation, and bit-to-bit skew of parallel interfaces. A previous white paper [3] covered the basic technology in detail. Simple enhancements to existing JAZiO technology that further reduce these effects are described. Attenuation and ISI affect both serial and parallel interfaces but bit-to-bit skew is only applicable to parallel interfaces.

II. HIGH FREQUENCY ISSUES IN LONG TRANSMISSION LINES

A. Attenuation

Signal attenuation occurs at high frequency on long transmission lines due to the skin effect because most of the current flows on the surface of the conductor (other attenuation effects lurk at even higher frequencies). In the past this effect could be ignored in most digital systems because the frequencies and distances were not so great. However, high performance systems are now experiencing this effect and future systems must provide a solution. The worst effect of attenuation is known as the “lone pulse problem” or the “first pulse problem” and is a pulse of a single bit time that follows several bit times of non-changing values.

A standard PCB trace represents a lossy transmission line. A pulse propagating down the lossy transmission line is attenuated by an amount determined by the distance and frequency. The attenuation manifests itself as high frequency roll-off that slows any fast transition. If the frequency is high enough then the slowed pulse is not able to make a complete transition during a single bit time period. However, whenever data values are constant for several bit times then the signal comes to rest at its maximum high or low level and the stage is set for the lone pulse problem. The leading edge of the lone pulse starts at the maximum high or low level but is too slow to make a full transition before the bit time ends and the trailing edge begins and, therefore, stops short of a full transition. The amount of shortfall as a percentage of a full transition is known as the percent attenuation. This attenuation results in a reduced “data eye opening” at the receiver which can cause malfunction. It is important to recognize that a constantly changing signal is also attenuated and cannot reach either the maximum high or low level but achieves a steady-state which is centered around the midpoint between the maximum high and low values and, thus, does not have a lone pulse problem. The significance of this fact will soon become apparent.

Most of the current methods used for high speed and/or relatively long distances in which the lone pulse problem is prevalent use differential signaling technology. The source transmits the differential data and its complement (D and /D) across the PCB trace to a receiving chip with a termination at the receiver. The PCB trace attenuates the transmitted signal by the time it reaches the receiver. If the attenuation is 50% then a lone pulse effects both sides of the differential pair and totally eliminates the eye opening at the receiver as shown in figure 1. Even at lower attenuation the eye opening is significantly reduced making reliable operation very difficult.

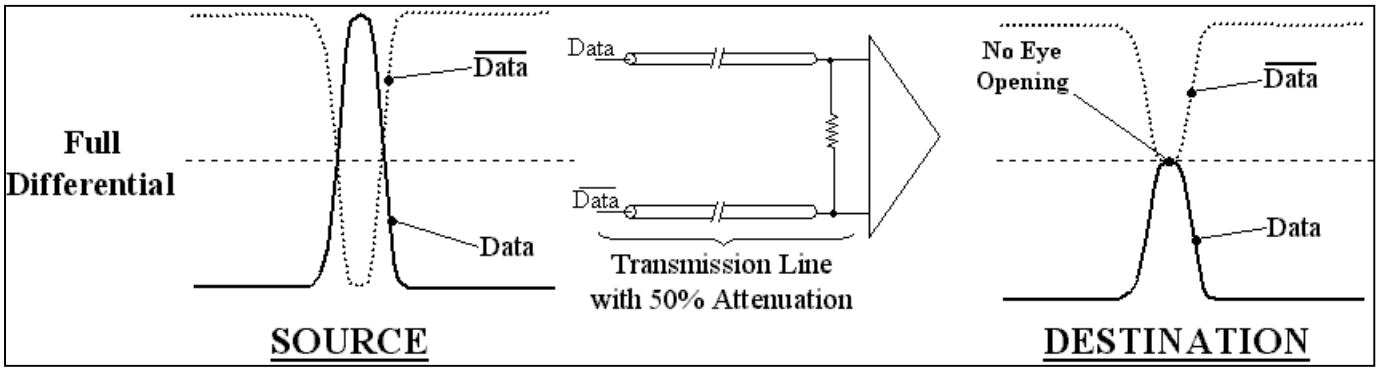

Figure 1: The effects of 50% attenuation on full differential signaling

B. Inter Symbol Interference (ISI)

The lone pulse problem also causes data dependent skew due to Inter Symbol Interference (ISI). Assuming that attenuation is less than 50% a lone pulse data eye opening exists even though it is reduced from the no-attenuation case. However, the time at which the differential signals cross each other on the lone pulse leading edge is delayed due to the slow edges which start their transitions from the maximum high and low levels. Likewise the crossing time for the trailing edge is advanced since the differential signals begin their transitions from intermediate levels. Thus the width of the lone pulse data eye is narrowed from both sides.

Now consider the case of a data pattern in which the data alternates every bit time. Since both D and /D are in steady state, the crossing points of leading edge and trailing edge are consistent. When the crossing points of the lone pulse case and the alternating data case are compared it can be seen that the crossing of the leading edge of the lone pulse is later and the trailing edge is earlier than the crossings for the alternating data case. The movement of crossing points for the various cases result in what is known as ISI jitter. The lone pulse has the distinction of achieving the maximum attenuation and the maximum ISI jitter. ISI jitter for a 20" trace of 50-Ohm impedance has been measured at approximately 130psec at 1.25 gigabits/sec increasing to approximately 180psec at 2.5 gigabits/sec. Assuming a measurement offset of about 20psec this jitter can be seen to increase approximately as the square root of the frequency as is predicted by skin effect theory.

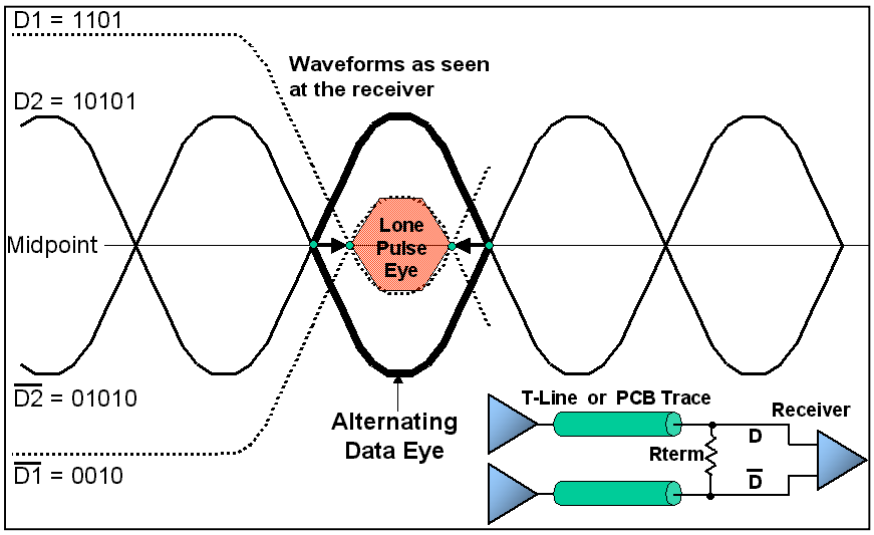

Figure 2: High frequency ISI jitter & attenuation for differential signaling

Figure 2 illustrates the reduced data eye opening due to attenuation and the crossing point movement due to ISI jitter. Two data patterns are shown, a lone pulse on data bit number one (D1 and /D1) and alternating data on data bit number 2 (D2 and /D2). The data eye opening of the lone pulse is shown as the small shaded area and the alternating data eye is shown by the thick outline. Notice that the lone pulse eye has shorter duration due to jitter and reduced amplitude due to attenuation. Figure 3 illustrates an even smaller data eye for pseudo-differential technology that uses a static voltage reference for the receiver to compare with a single date signal per bit. It is believed that this pseudo-differential technology will not operate reliably beyond a 10 to 15% attenuation level and, therefore, is unlikely to find use at data rates much beyond 1 gigabit/sec/pin at reasonable distances.

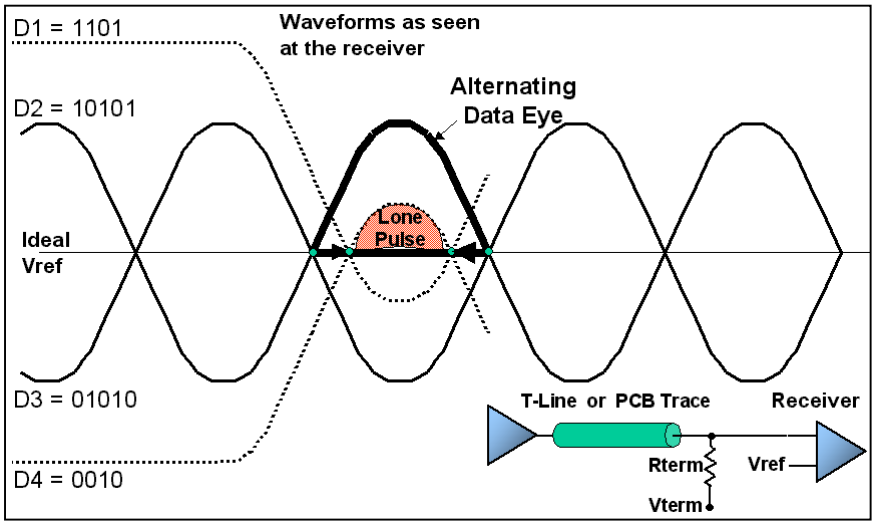

Figure 3: High frequency ISI jitter & attenuation for pseudo-differential signaling

C. The Effects of Attenuation and ISI using JAZiO

The effect of high frequency ISI jitter and attenuation for JAZiO technology is quite different. JAZiO uses a pair of alternating and complementary voltage and timing references (VTRs) that are compared with each single-pin per bit signal at the receiver. The VTRs transition every bit time and operate in the steady state. In fact these VTR signals appear identical to the D2 and /D2 signals of figure 2 and have no lone pulse problem. For JAZiO technology the lone pulse

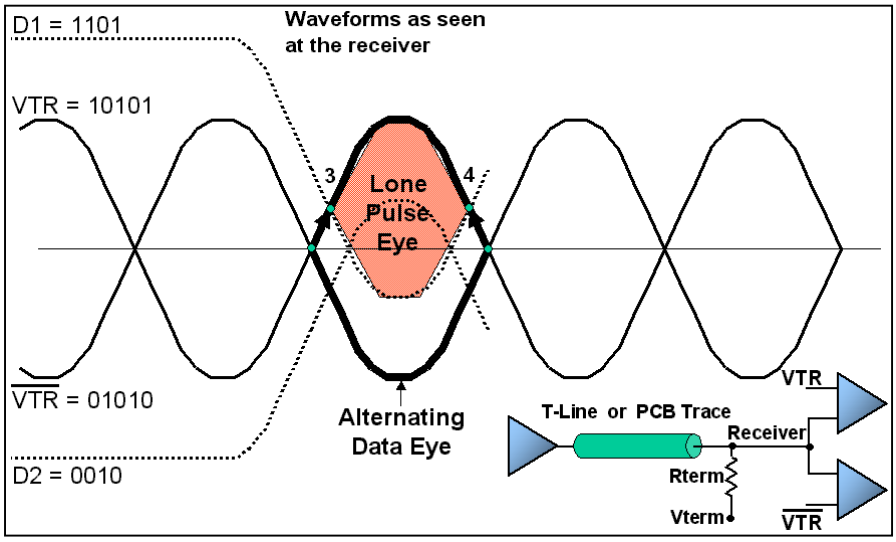

Figure 4: High frequency ISI jitter & attenuation for JAZiO signaling

problem is restricted to a data signal that has remained at a constant value for a few bit times. The eye diagram for JAZiO signaling using the same voltage swing and timing conditions as differential (figure 2) and pseudo-differential (figure 3) is shown in figure 4. The data eye formed by signal D1 crossing signal VTR is seen to be much larger than the eyes for full differential and pseudo-differential technologies which are superimposed in the same figure.

The reason for the larger eye opening in the JAZiO case is that the JAZiO data signal is compared against a steady-state VTR signal. One advantage is that the VTR is centered around the midpoint between the maximum high and low value and always makes a larger signal swing beyond the midpoint than a lone pulse makes. So the JAZiO data eye is taller. A second advantage is that the VTR signal does not experience ISI jitter and, therefore, the crossing point of D1 and VTR is not delayed as much on the leading edge as is the crossing point of D1 and /D1 in the full differential case and, likewise, is not advanced as much on the trailing edge. So the JAZiO data eye is wider. Another way to understand the difference is to recognize that full differential technology compares two signals both of which experience maximum lone pulse attenuation and ISI jitter while JAZiO technology compares one signal which experiences maximum attenuation and jitter to another signal which is immune to jitter and largely immune to attenuation. Although proponents of full differential technology believe strongly in the symmetry of the true and complement signals, we see here that this symmetry provides a "double dose" of attenuation and jitter problems while JAZiO only experiences a single dose. The data eye available at the receiver for the 50% attenuation case is shown in figure 5 and table 1 compares the differential (data eye amplitude) available at the receiver between JAZiO and full differential for various attenuation levels. It is important to notice that JAZiO has a larger data eye at all attenuation levels with a single data pin per signal and, so, provides a larger data eye while delivering **twice the data bandwidth per pin**.

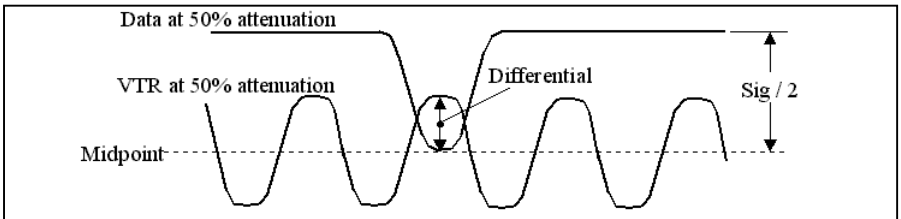

	Attenuation						
	0%	10%	20%	30%	40%	50%	
~ % Diff	100	80	60	40	20	0	Full Differential
~ Diff (mv)	300	240	180	120	60	0	
~ % Diff	100	80	60	40	20	0	Full Differential
~ Diff (mv)	600	480	360	240	120	0	
~ % Diff	100	81	63	47	32	17	JAZiO (change case)
~ Diff (mv)	600	486	378	282	192	102	

NOTE: Diff is magnitude of differential at receiver
% Diff is differential at receiver as a % of signal at source

Figure 5 & Table 1: Attenuation comparison between full differential and JAZiO

By studying figure 4 it can be seen that the largest contributor to the JAZiO expanded data eye is the larger amplitude due to the VTR signal transitions being centered around the midpoint. Therefore small variations in signal amplitude do not have a significant effect on the available data eye at the receiver. Differences in signal amplitude could be caused by cross talk, driver size, termination, or localized variations in power supply. By controlling the amplitude of one signal pair (VTR and /VTR) it is easy to obtain large data eyes at all receivers. Careful shielding on just the VTR pair reduces the effect of cross talk on the data eye at the receiver. This is one of the major advantages of JAZiO signaling.

D. Output Equalization or Pre-Emphasis

Since the lone pulse problems due to attenuation and ISI jitter are so pronounced for full differential and pseudo-differential technology, various techniques have been developed such as data encoding, output equalization, or pre-emphasis. Serial links use encoded data transmission so that there are never long periods of time when the signal is not changing. Thus, the lone pulse problem is reduced. However, this encoding costs latency, complexity, and bandwidth

reduction. In parallel data transmission this encoding of data can be very problematic and is almost never used. Thus, for interface standards using parallel data transmission (such as LVDS, SCSI), the lone pulse problem exists.

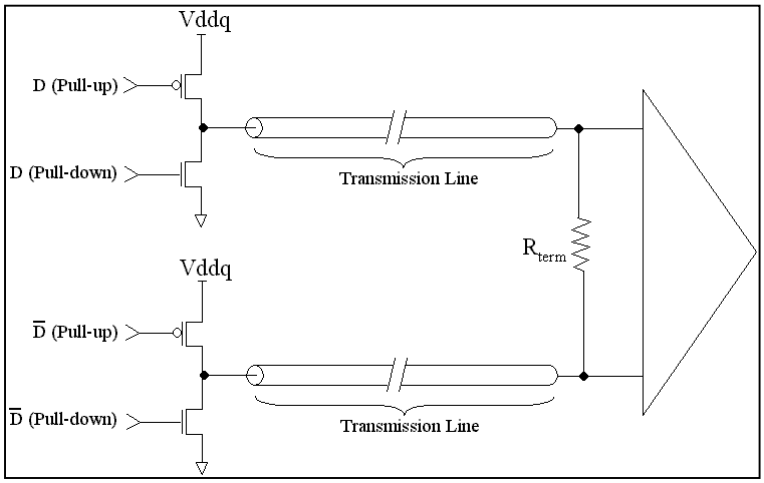

Figure 6: Full differential transceiver

The solution normally used in parallel data transmission is called output equalization or pre-emphasis. Here the output level driven from the transmitter is varied depending on the data pattern. If the data does not change for a few bit times the signal transmitter reduces the signal level so that a lone pulse can begin at a voltage level closer to the midpoint and, thus, provide a larger data eye. Pre-emphasis in differential signaling thus requires both signal and its complement to reduce the output drive strength allowing the differential swing to be reduced at the driver and the receiver in preparation for a lone pulse. The transceiver used in differential signaling uses a differential terminator at the receiver end as shown in figure 6. Since the terminator is between the signal and its complement it cannot reduce the swing at the receiver end because there is no current path to another source (like midpoint voltage termination). Therefore the output swing is generally reduced at the transmitter itself. This becomes complex and usually requires a driver that can generate precise voltages depending on the data pattern. When the output switches state it switches at normal strength and the net effect is increased drive strength from the steady state to the new state. Step down control circuitry is used to determine after a specified number of clock pulses how much drive strength should be stepped down and in how many increments. This increases the power compared to a push-pull driver in which only the driver size and termination determine the voltage swing.

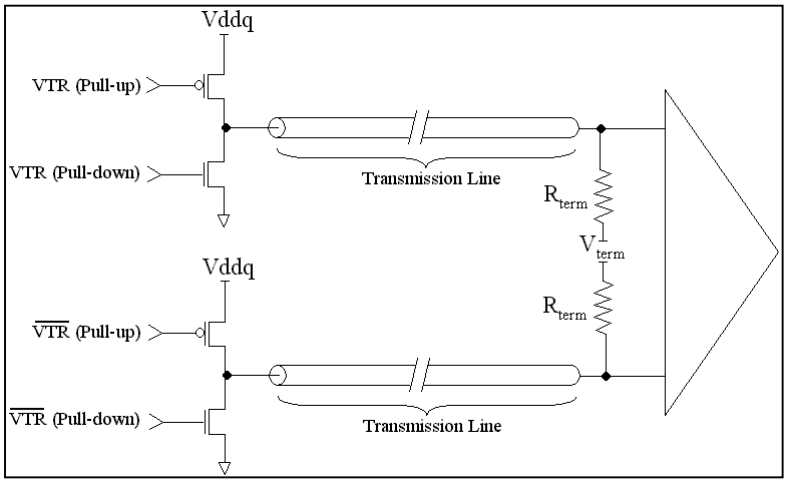

Figure 7: VTR transceiver

E. JAZiO Output Equalization

Although it has been shown in section II.C that JAZiO technology is inherently better able to deal with attenuation and ISI jitter than full differential or pseudo-differential technology, in this section we provide a very simple and effective method for enhancing JAZiO technology to deal even better with these problems.

JAZiO technology is based on receivers that compare each single ended signal with a pair of alternating voltage and timing references that are driven concurrently from the same chip. The simplified output driver has a pull-up to provide the desired high voltage level and a pull-down to provide the desired low voltage level. Methods for gate drive control that can be used to equalize signal slew rate by compensating for variation in process, temperature, and power supply are well known and understood for high-speed transmitters. The VTR pair is terminated at the receiver to a midpoint voltage reference, Vterm, which is approximately 0.8V as shown in figure 7. Since the termination voltage is centered between the maximum high and low levels, this method is called center termination. (Of course, a non-center, open drain with pull up termination scheme can also be used but the center termination scheme is described here for simplicity.) The termination resistance and driver impedance is matched to the line impedance and typically is approximately 50 ohms. Of course, the termination voltage and resistance can be different depending on the desired power dissipation, material of the PCB, etc. The single pin per data bit transceiver shown in figure 8 is similar. The signal swing is symmetrical around the Vterm and is approximately 300mV above and below Vterm (0.8V). As pointed out earlier, the VTRs change every bit time and have no lone pulse problem although the peak-to-peak amplitude is attenuated at the receiver. The VTRs are identical to any data signal which transitions every bit time and all of them have a symmetrical swing above and below Vterm with amplitude dependent on the attenuation. But a signal that has not changed for few bit

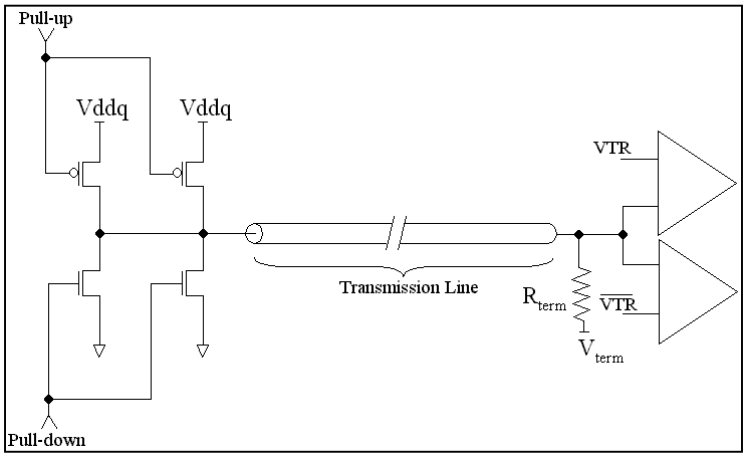

Figure 8: Single-ended data transceiver

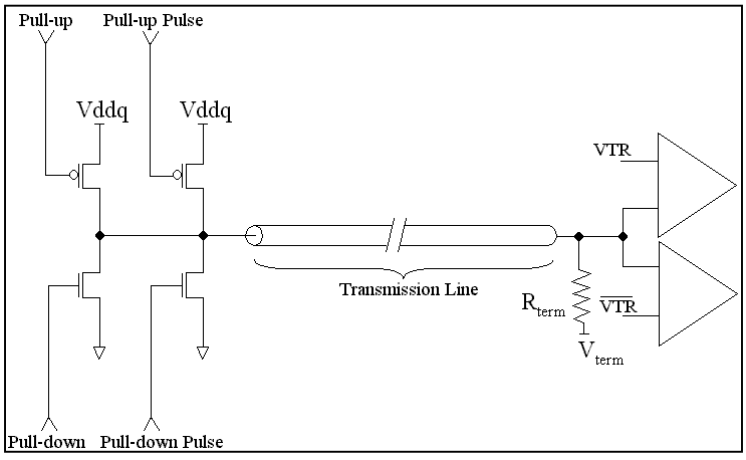

Figure 9: Single-ended data transceiver with output equalization

times will reach the maximum high or low level (lone pulse effect) and its next transition will only reach the midpoint voltage in the 50% attenuation case.

However, if a simple "turn off before turn on" scheme is used in the driver then center termination allows the lone pulse to make a transition beyond the midpoint because the termination resistor starts pulling the signal toward the midpoint as soon as the "turn off" occurs. So the lone pulse begins at this new level. The final amplitude of the lone pulse in the 50% attenuation case depends on how far the termination resistor has pulled the signal before the "turn on" occurs and will be somewhere between the midpoint if there is no time between "turn off" and "turn on" and the 75% level if there is maximum time between "turn off" and "turn on" to allow the terminator to pull the signal all the way to the midpoint before the lone pulse begins.

Figure 9 shows a transceiver with a very simple output equalization technique to reduce the effect of lone pulse for the single ended signals in the JAZiO signaling scheme. The output driver pull-up and pull-down are split into two pull-ups and two pull-downs both of which are turned on to deliver the normal drive strength for approximately one bit time for every signal transition. If there is no transition the appropriate pull-up or pull-down pulse is turned off allowing the output signal drive to be reduced. An implementation of output driver control is shown in figure 10. The discharge path of the signal is through the Rterm at the receiver. The widths of the pull-up and pull-down pulses that control one of each pull up/pull down pair can be programmable based on the amount of line attenuation present in the system. The programmable delay in figure 10 is approximately equal to a bit time and is programmable to various bit time lengths desired for the operation of the transceiver.

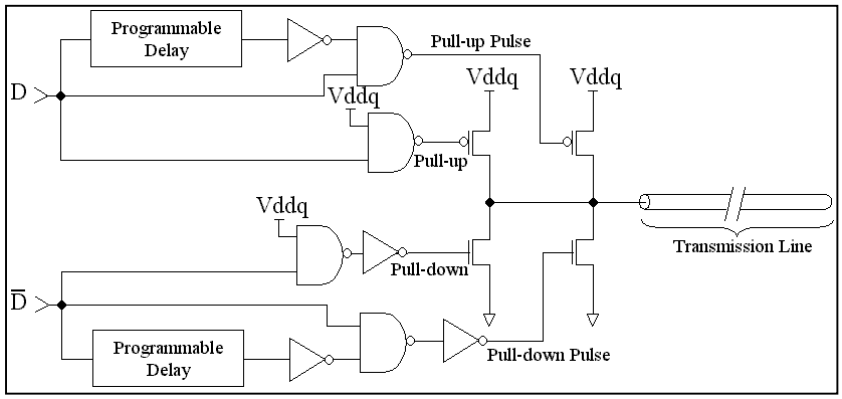

Figure 10: Output equalization control

Since the transceiver is intended for operation in high signal attenuation ranges and the output is being equalized to an intermediate voltage, it is preferred to have larger swing on the data signals compared to the VTRs. An appropriate termination voltage is Vterm=0.8V with VTR swing from 0.5V to 1.1V (300mV above and below Vterm) and data signal swing of 0.4V to 1.2V (400mV above and below Vterm) at the transmitter. The VTRs quickly settle to a steady state of slightly higher than 200mV swing (0.7V to 0.9V) at the receiver after their initial, start-up transitions and then data signaling can begin.

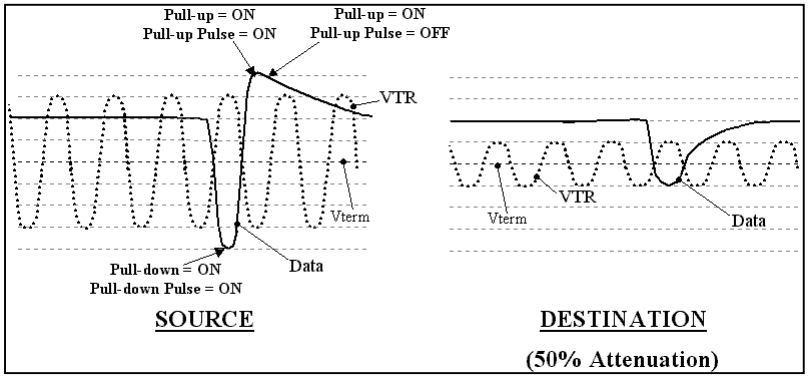

Figure 11: Effects of attenuation (lone pulse) using output equalization

Coming out of a power-up initialization sequence the data signals have been static and the pull up and pull down pulses are turned off. Therefore the data signals are stable at a voltage between the maximum high or low level and Vterm in preparation for the first (lone) pulse. These data levels are chosen to be about 1.0V and 0.6V which are slightly outside the maximum peak-to-peak levels of the VTRs and provide JAZiO technology with increased robustness in the no-change case.

At the transmitter the lone pulse voltage swing of a data signal is 600mV (1.0V to 0.4V for low going pulse or 0.6V to 1.2V for high going pulse). Assuming a 50% attenuation case, the voltage swing at the receiver is 300mV (1.0V to 0.7V or 0.6V to 0.9V). These levels are compared to the VTR signals with data eye amplitudes of 200mV (0.9V – 0.7V or 0.7V – 0.9V). This 200mV eye amplitude is sufficient for reliable high-speed operation especially due to the increased data eye width described in section II.C. See figure 11.

If the data signal transitions every bit time the output driver switches with full drive strength every bit time and the pull-up and pull down pairs are controlled identically. The waveforms for this case are shown in figure 12. The swing of the signal is higher at both the transmitter and the receiver. In cases where the data signal changes every other bit time and so on the data eye amplitude at the receiver is somewhat higher than 200mV, but less than the change-every-bit-time case.

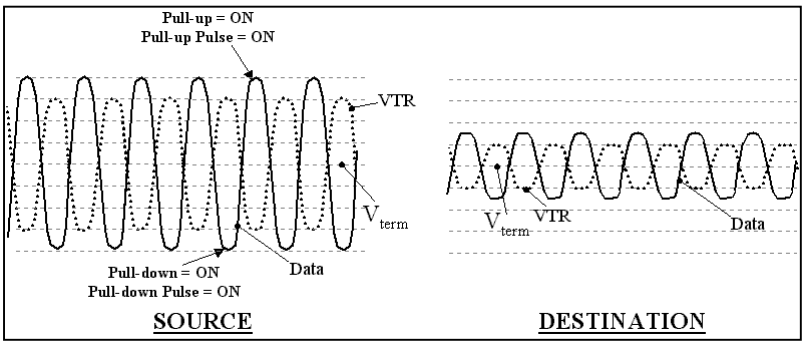

Figure 12: Effects of attenuation (continuous transitions) using output equalization

III. HIGH FREQUENCY ISSUES IN PARALLEL INTERFACES

A. Bit to Bit Skew

The transmission of parallel data over long distances is affected by differing propagation delays of the channels. These differences cause the arrival times of the individual data bits at the destination receiver to vary. The difference between the bit arrival times is referred to as bit-to-bit skew. When the channel lengths are kept relatively short or the transfer speeds are slow, the skew effect is not negligible. As channel lengths increase, the propagation delay differences may accumulate and the skew increases. Eventually, the increasing channel lengths result in bits from one word "drifting" into the bits of the next word. Once this occurs, the receiver has difficulty aligning the bits into their respective words.

One solution is to create individual serial data bit streams out of each channel. This scheme encodes the clock signal directly into the bit stream, recovers the clock signal at the receiver, and reconstructs the data word through signal processing techniques. This system requires complex signal processing at the receiver end, utilizes a valuable portion of the bandwidth of each channel encoding the clock, and increases latency.

Other schemes call for the receiver to digitize the small data eye at its input to an internal, full swing signal for each channel regardless of skew and then attempt to realign the parallel channels by somehow latching them into the same word. We call these schemes post digitization deskewing since they are performed after digitization by the receiving system. The problem with these schemes is determining the amount of skew since the data eye opening at the receiver is data pattern dependent and the digitization level conversion also creates some additional skew before the latch. Existing deskewing schemes compensate for skew by providing a programmable delay between the receiver and the latch that is programmed by comparing the received data to the expected data during a tuning pattern.

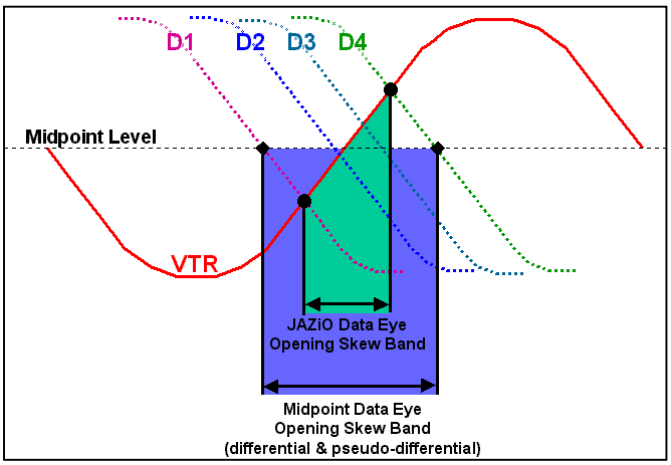

Figure 13: Skew effects on JAZiO versus other technologies

B. JAZiO Bit-to-Bit Skew

The effect of bit-to-bit skew with JAZiO technology is different than that for differential or pseudo-differential technology. The eye opening for both pseudo-differential and full differential technology always starts at the midpoint voltage level where the data signal crosses either the fixed Vref or the complement of the data signal respectively. Therefore, the width of the skew band of the eye opening exactly matches the skew of the various input data signals at the midpoint as shown in figure 13. For JAZiO technology, the width of the skew band of the eye opening is reduced because the data signals that come earlier than VTR cross VTR later than they cross the mid-point and data signals that come later than the VTR cross VTR earlier than they cross the midpoint. Therefore the data eye opening skew band for JAZiO is inside the skew band for the other technologies and up to 50% smaller allowing JAZiO to operate reliably in systems with larger data skew.

C. JAZiO Skew Monitor

Although in the previous section it has been shown that the data eye skew band for JAZiO technology is inherently smaller than for other technologies, in this section a method is presented for further enhancing JAZiO technology to deal with skew.

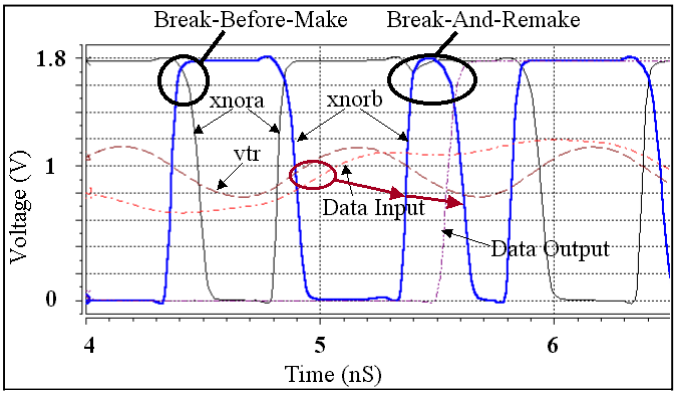

Figure 14: Receiver simulation at 2 gigabits/second

A good deskewing technique should offer an easy and accurate method of determining the amount of skew in each signal. It has been shown elsewhere [1] that a JAZiO receiver includes two exclusive-nor gates, XNORA and XNORB, that control multiplexers used to select between the two differential comparators that compare the data signal to each of the VTR signals. It is also shown in [1] that these XNORs are designed so that when it is time to switch the selection from one comparator to the other, the first comparator is disabled before the second comparator is enabled. We call this method "break before make" selection and spice waveforms are shown in figure 14. One result of this method is that a receiver

which experiences sufficient skew between the data signal and VTR may undergo a "break and remake" during which the first comparator is falsely disabled before being correctly re-enabled with the second comparator never being enabled. It is important to realize that this "break and remake" does not cause a failure in the receiver since correct data is always present at the receiver output. Rather, it represents a marginal situation in which skew is not bad enough to cause failure. If skew were to grow sufficiently worse then a failure would occur as the second comparator is falsely enabled. So this "break and remake" case represents marginal but still correct operation as shown in the middle timing diagram of figure 15.

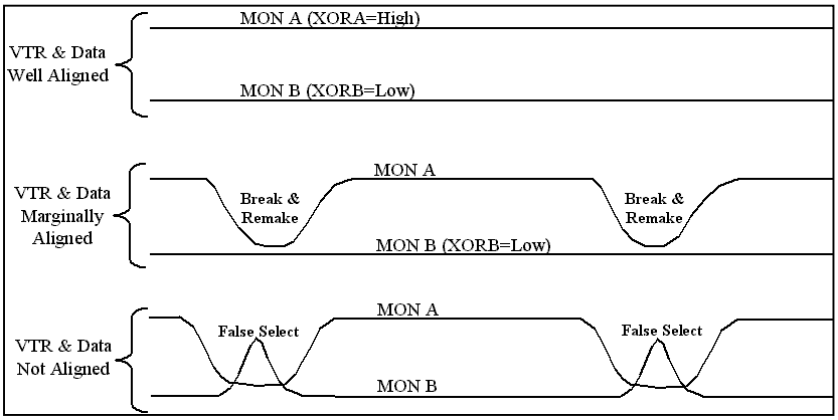

Figure 15: Alternating data pattern measuring & monitoring VTR to data alignment

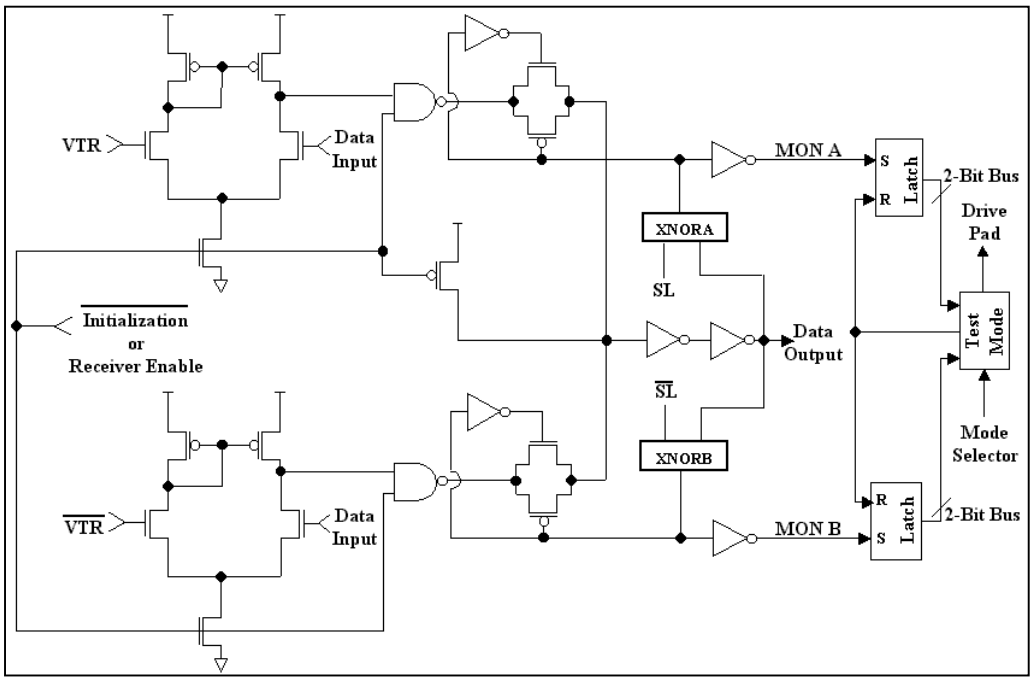

Figure 16: Receiver schematic with monitors

Figure 16 illustrates how the XNOR gates can be used to provide receiver monitors. The monitor logic uses latches to store the "break and remake" or "false select" pulses. By detecting a "break and remake" pulse on a XNOR during a bit time in which the XNOR should not switch (actually any bit time in which the new data is different than the previous data) it is easily determined when the receiver is marginal with respect to skew. This information can then be fed back to the transmitter by one of a variety of means. If transmitters are provided with programmable delays as shown in figure 17 then a self-tuning system minimizing skew is possible. Several methods of self-tuning are available:

1. Provide "tuning patterns" at initialization during which the programmable delays are optimized.
2. Interrupt normal operation periodically to rerun the "tuning patterns."
3. Provide periodic "tuning packets" during actual operation that continuously retune the delays.
4. Provide a portion of each packet during which the skew can be monitored and retune delays as needed.

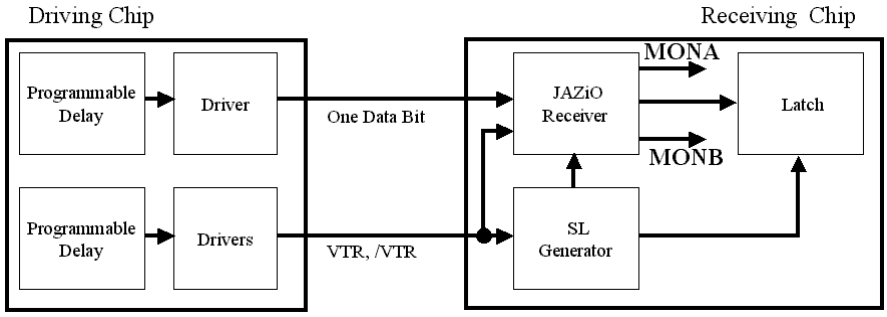

Figure 17: Block diagram of JAZiO tuning/monitoring technique

Of course, combinations of these methods can also be used. Depending on the system requirements a variety of solutions are available ranging from a simple power-up tuning to continuous, in-system retuning for high reliability systems which are not allowed to be taken down for maintenance. These monitors can also be used for testing each transceiver at various frequencies, with the other transceivers acting like noise sources for cross talk, coupling and simultaneous switching.

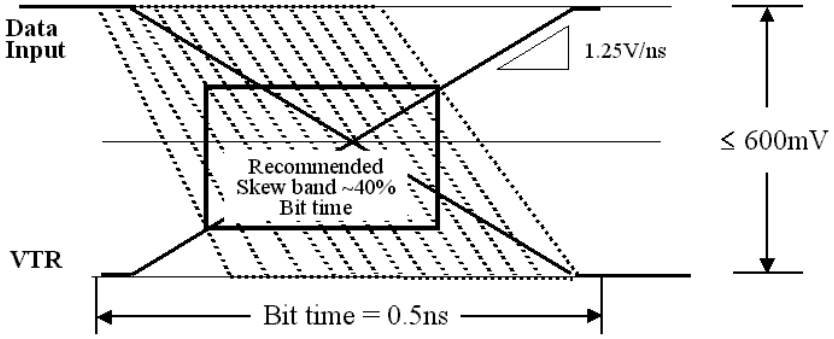

Figure 18: Recommended basic JAZiO skew band (~40% bit time)

Figure 18 shows a basic JAZiO skew band. Simulations determine that reliable operation can be easily achieved with a skew band of 40% of the bit time. If system effects such as mismatches between signal traces, package pins, drivers, pre-drivers, etc. combine to exceed this 40% then methods described in this section may be employed to bring the skew back within the 40% band.

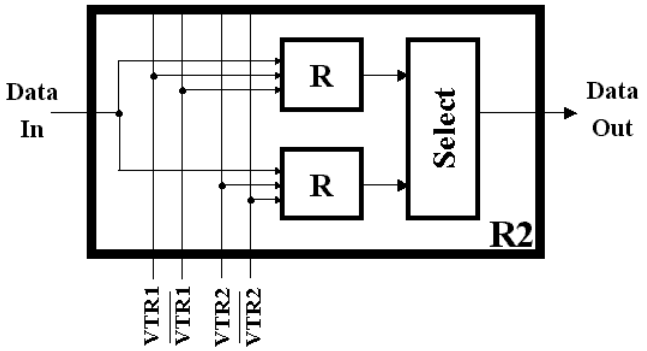

Figure 19: Dual VTR pair receiver block diagram

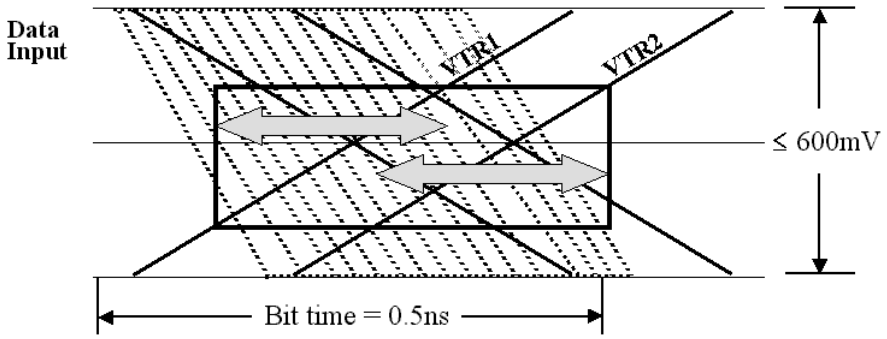

Figure 20: Improved skew with dual VTR pairs and overlapping skew margin

D. Dual VTR Pair Method

This section provides another method for dealing with bit-to-bit skew in systems with large skew. Two pairs of VTRs with a time delay offset can be used as shown in figure 19. In this case a five-input receiver is used which consists of two typical JAZiO three-input receivers and a selector that determines which of them is being used in any particular bit time. The monitors described in the previous section can be used to control the selector in order to allow the selection of the currently best performing receiver. With this method a widened skew band can be achieved as shown in figure 20. The band is increased to about 70% of the bit time allowing some overlap in the operating region compared to the 40% shown in figure 18. With a fixed delay between the two VTR pairs this method is excellent for dealing with fixed and constant bit-to-bit skew. It can also be used to deal with ISI jitter since it has been shown in section II.B that the leading edge of a lone pulse suffers delayed data eye opening while the trailing edge suffers advanced eye opening. Therefore it is possible to select the receiver corresponding to the later VTR pair when a lone pulse leading edge is anticipated while selecting the earlier VTR pair in the other cases.

It is also possible to combine the two VTR pair method of this section with the programmable transmitter delay methods of the previous section. The time delay offset between the VTR pairs can be optimized by methods described there. Additionally, still better tolerance to skew can be achieved by tuning delays on each data transmitter as well as the VTR pairs. In this case it is possible to widen the delay between the VTR pairs and tune the data transmitters to match the nearest pair resulting in reliable operation in systems with more than a bit time of skew as shown in figure 21.

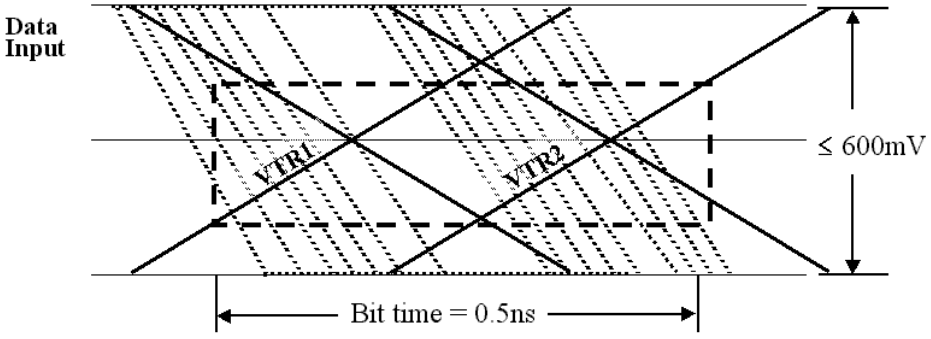

Figure 21: Improved skew with dual VTR pairs and non-overlapping skew margin

E. Advantages of JAZiO in long transmission lines for High Frequency Parallel Interfaces.

This section provides a summary of JAZiO advantages compared to full differential or pseudo-differential technology:

JAZiO offers three ways of reducing the effects of ISI.

1. The JAZiO data eye is wider in the lone pulse case because it begins earlier and ends later since the data signal is compared to VTR which has no jitter.

2. The JAZiO data eye is easily expanded even larger by a very simple output equalization scheme with center termination.
3. A two-VTR pair scheme can be employed which allows the later pair to be used on lone pulse leading edges.

JAZiO offers two ways of reducing the effect of attenuation.

1. VTRs operate in steady state and dominate the data eye amplitude reducing the effect of attenuation.
2. The simple, center termination, output equalization scheme further increases data eye amplitude.

JAZiO offers three ways of reducing the effect of bit-to-bit skew.

1. The JAZiO data eye skew band is narrowed since early arriving data signals cross VTR later than they cross the midpoint and late arriving data signals cross VTR earlier than they cross the midpoint.
2. The JAZiO receiver monitor with programmable transmitter delays can be used to reduce skew.
3. Two VTR pairs can be used with data signals aligned to the appropriate one.

IV. VARIOUS JAZIO SOLUTIONS

At this point it should be clear that JAZiO technology can be used on a variety of enhanced levels. Some of these levels are:

1. Basic JAZiO.
2. JAZiO with larger amplitude swing on the signals relative to VTRs.
3. JAZiO with skew monitors and programmable delays on data signals and VTRs.
4. JAZiO with dual VTRs and skew monitors.
5. JAZiO with dual VTRs, skew monitors and programmable delays on data signals and VTRs.
6. JAZiO with dual VTRs, skew monitors, output equalization and programmable delays on data signals and VTRs.

V. CONCLUSION

JAZiO technology is shown to be a very effective way to achieve high data bandwidth in "hostile" systems with large ISI jitter, large signal attenuation, and large bit-to-bit skew. In section II.C it is shown that the JAZiO data eye opening is much larger than the opening for full differential technology in systems with large ISI jitter and large attenuation even when JAZiO is transmitting at twice the data rate per pin. In section II.D a very simple and effective output equalization scheme with center termination is introduced which further expands the JAZiO data eye. In section III.B a bit-to-bit skew solution is provided with receiver monitors and programmable transmitter delays. In section III.D a dual VTR pair method is introduced which provides still more bit-to-bit skew solution. Taken together these methods provide a long-term roadmap of JAZiO extensions to higher and higher data rates.

It is important to understand that basic JAZiO technology is very simple and easy to implement. It is very low latency and inherently low power since it uses a single pin per bit with low switching levels. It is very low cost since the basic receiver is so simple and small.

It should also be noted that when comparing JAZiO technology to full differential technology one must not forget that in order to achieve equal data rates per pin that full differential signal pairs must operate at twice the frequency of JAZiO signals. All high frequency effects, such as ISI jitter and signal attenuation, introduced by the package, PCB, connectors, chip tester, etc., are exacerbated at double the frequency. Also, note that in section II.C JAZiO and full differential were compared at the same frequency with JAZiO winning the data eye battle by a large margin – **this is at twice the bandwidth per pin!**

For the reasons given here it is clear that basic JAZiO technology can far outperform its rivals without even requiring the enhancements outlined in this paper. When the enhancements are employed in a deliberate manner, as needed, JAZiO technology can still be scaled many years into the future, while other technologies exceed their limits of scalability (run out of gas).

REFERENCES

[1] E. Haq, et al, "JAZiO Signal-Switching Technology," IEEE Micro, January-February 2001.
[2] K. Krewell. "JAZiO: Slow Edges Can Run Fast-A Novel Approach to High-Performance Bus Interfaces," in Microprocessor Report, issue 2/21/00-02, February 2000.
[3] "JAZiO High Speed Digital I/O Signal Switching Technology," www.jazio.com (white paper), July 2000.