KANDOU BUS

KANDOU'S ENRZ INTERFACE

WHITE PAPER

Summary

In this white paper we will introduce the concept of Kandou's Ensemble NRZ (ENRZ) signaling, and how it relates to other high speed serial signaling methods such as NRZ or PAM-4.

Introduction

Wired communication between semiconductor devices and other IC components is facing serious challenges as a result of ever increasing demand for bandwidth and lower power consumption. The standard response of scaling transistor geometry is running out of steam because of physical limitations of analog circuitry and interconnect technology (packages, connectors and wires). Higher throughput has traditionally meant a higher frequency of communication: doubling the frequency means doubling the throughput – at least in theory. However, increasing the frequency comes with a large price tag: the loss of signal integrity (SI). Simple physics shows that high frequency signals suffer more attenuation than lower frequency ones. Signals arriving at the receiver may have only between a tenth and a hundredth of their original strength. Increasing the wire strength is not an option since this would lead to higher system power consumption. Moreover, such an approach would alleviate the problem only marginally, since more power for stronger signals to overcome noise may create more noise on the other wires. Another option is to try to limit the signal attenuation and this is typically hard and expensive to do.

A third option is a change of the signaling method. Traditionally, high speed links have used differential signaling (or NRZ, as this method is often called). While this method of signaling leads to robust communication, it suffers from a major disadvantage: two wires are needed to transmit one bit at every clock interval. In a pin-limited application (which covers many high speed applications), reaching a given aggregate throughput means running the clock at a higher frequency, and a higher clock frequency means larger attenuation of the signals, larger crosstalk effects, larger electromagnetic interference noise, and in general, a much larger headache to maintain signal integrity.

Changing the signaling method has been recognized by the industry, and has led to the introduction, and partial adoption of PAM-4 signaling within the IEEE 802.3 and OIF-CEI standards bodies. PAM-4 signaling is a generalization of NRZ in that correlated signals are sent across two wires. In contrast to NRZ, PAM-4 signals can have one of four possible voltage levels, leading to double the throughput per clock interval. PAM-4 has a number of advantages, and at the same time, a number of disadvantages. Because the throughput is doubled, PAM-4 signaling can halve the frequency for a given throughput as compared to NRZ. Therefore, the signals are subject to less attenuation. At the same time, however, the margins are lower -- from the start, PAM-4 has a 9.5 dB disadvantage compared to NRZ since the vertical opening of the eye at the transmitter is smaller by a factor of 3. The larger number of voltage levels at the output of the receiver comparators leads to higher susceptibility to intersymbol interference (ISI) noise, which reduces the horizontal eye width. Moreover, this larger number also leads to issues with decision feedback equalization (DFE). For the latter reason most PAM-4 receivers sample the data once, and perform all post-processing using digital rather than analog circuits. The reduced horizontal eye width makes correct sampling a challenging task due to the inherent jitter of the associated circuitry.

PAM-4 signaling is, however, not the only way forward. In this white paper we introduce 4-wire Ensemble NRZ (or ENRZ for short), a signaling method utilizing a low-skew group of 4 wires (rather than a low-skew group of two wires used by NRZ or by PAM-4). In contrast to NRZ which sends 2 bits over 4 wires in one clock interval, and in contrast to PAM-4 which sends 4 bits over one clock interval, ENRZ transmits 3 bits during one clock interval. At equal aggregate throughput, ENRZ experiences therefore less attenuation than NRZ, but more attenuation than PAM-4. Compared to PAM-4 signaling, ENRZ has one very big advantage however: when correctly implemented, it is far less susceptible to ISI noise than PAM-4 is. Therefore, even though the clock frequency used for ENRZ is higher than that for PAM-4, the horizontal opening of the corresponding eye is typically much larger, leading to a much higher margin, as will be described later in this white paper.

Kandou's ENRZ Solution

Imagine a device connected to another through two differential pairs. In every clock interval, 2 bits are transmitted over these four wires. To achieve an aggregate throughput of 112 Gbps, say, a clock is required that runs at a frequency of 28 GHz. This high frequency puts severe limitations on the link, even if the signals have to traverse a very short distance: typically, the amplitude of the signal is attenuated by a factor of 100 or more (corresponding to 40 dB or more of attenuation), and the inherent time-constants of the wires spreads the energy of the signal in time, causing subsequent bit signals to be corrupted. Though quite sought after by several industries, reliable and low-power communication at these speeds is far from a reality with today's methods.

Looking closely, the main culprit is the fact that we are spending too much for redundancy. That is, we are not efficient enough with the way we are using the wires: NRZ effectively transmits real data on half of the wires only. This means that the clock rate has to be very high to hit a given aggregate throughput. We can say that 50% of the wires are redundant when using NRZ. Do we really need that much redundancy? Looks like we do if we only have two wires. In this case the only other option is to use single-ended signaling, an option that is absolutely not viable for high data rates and differential circuit techniques are well known to solve common mode noise, power supply noise and crosstalk

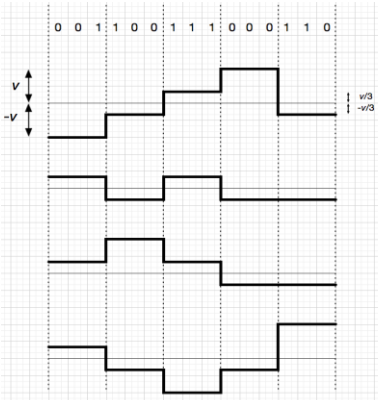

Figure 1: Example ENRZ waveforms on four wires

But can we consolidate if we have more than two wires? Fortunately the answer is yes, as Kandou's ENRZ 4-wire interface shows. ENRZ uses a transformation of the incoming signals. At every clock interval one of the four wires is driven high, and the other three wires are driven low to 1/3rd of the value of the high wire – or alternatively, one wire is driven low and the other three wires are driven high. With two bits of information we can choose which of the four wires is different from the other three, and with a third bit of information we can choose whether to drive that wire high or low. This way we can encode three incoming bits in an ensemble wave form on the 4 wires. Figure 1 provides an example. In the first clock interval (the right most interval), we are transmitting the 3 bits 1,1,0. The first two bits provide the binary representation of the unusual wire (wire 3, i.e., last wire) and the last bit says that the wire is driven high, i.e., to v. All the other wires are therefore driven to $-v/3$. In the fourth time interval the bit sequence 0,0,1 is transmitted. Again, the first two bits determine the position of the unusual wire, i.e., wire 0, and the last bit shows that this wire is driven to $-v$, while the other wires are driven to $v/3$. Another way of describing this is the orthogonal superposition of 3 signals onto 4 wires.

The simplicity of ENRZ is quite deceiving as it hides a number of crucial subtleties. For example, it looks like the extra bit per clock interval was gained by using a larger number of voltage levels (4 instead of 2), and hence, by losing noise margin. In reality, however, this scheme is actually binary once combined with the correct receiver structure, as we will see in the next section. The next section also shows that this receiver falls naturally off the theory that produced ENRZ in the first place. In fact ENRZ is a superposition of three different binary channels, and the detector "untangles" this superposition to get back the bit streams.

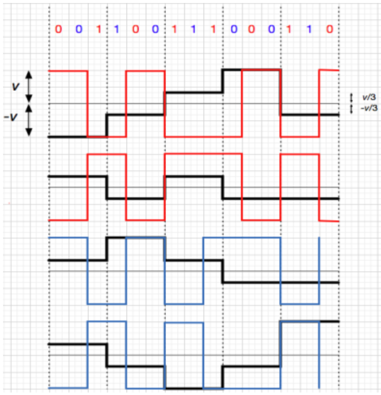

Figure 2: NRZ waveforms superimposed on ENRZ waveforms at equal throughput

To appreciate ENRZ and its advantages over NRZ, Figure 2 shows how the same bits could be transmitted with NRZ on the same number of wires. The two differential signals on the two pairs are denoted by the colors red and blue. Every second bit is transmitted on the first pair, and every other second bit on the second pair. In order to keep up with the speed, the differential signal needs to use a clock that has 1.5 times the frequency of the clock used for ENRZ. The signals used by NRZ will therefore attenuate much more, and will be subject to more noise, meaning that more power needs to be spent to recover them at the receiver. Also, the figure shows that NRZ will use more power to send the data: the amplitude of the signals transmitted is a good proxy for the power; ENRZ is much "quieter" in that respect and uses a less power for the transmitted signals. In fact, adding all the ways in which ENRZ saves power, an ENRZ implementation of the 112 Gbps link above can lead to significant improvements in terms of power even if someone managed to build a solution using NRZ.

ENRZ was designed to be minimally invasive to current high speed serial interconnect infrastructure. It serves as a bridge from today's chip-to-chip communication world to tomorrow's world where resources are consolidated across more than four wires (and hence power is decreased and speed is increased even further).

Because of the quaternary symbols on the wires in an ENRZ signaling scheme, the ENRZ driver is quite similar to a PAM-4 driver in the way the four signal levels are constructed. Since the signals are different than those of PAM-4 signaling, a different type of encoder is needed to extract driver information from the incoming three bits. The encoder is a very small piece of digital logic, less than 30 gates. The ENRZ transmitter also allows the inclusion of Tx FIR circuits as part of the overall channel in the equalization.

A link consisting of two skew matched differential lanes can be transformed into an ENRZ link using a minimally invasive procedure.

Many of the traditional SerDes circuits, such as those responsible for clock and data recovery (CDR) or decision feedback equalization (DFE) can be generalized for ENRZ signaling. Moreover, since they can be consolidated across four wires, they use less power than the corresponding circuits used in NRZ or PAM-4 based signaling. Another source of reduced power is the smaller Tx driver power used by ENRZ. In fact, ENRZ uses only 1/3rd of the driver power needed for three differential lanes.

A Closer Look at ENRZ

ENRZ signaling is derived from a mathematical transformation called the "Hadamard Transform." Its origin lies in a different view of NRZ signaling. In NRZ signaling a bit b, typically conceptually modulated as +1 or -1, is transmitted over two wires as b and $-b$. One way to view this signaling is by using a matrix transformation. Denoting the values transmitted on the two wires by x and y, the transformation in question is given by

$$(0, b)\begin{pmatrix}1 & 1\\ 1 & -1\end{pmatrix} = (x, y)$$

The matrix in the above equation has a number of interesting properties. To begin with, it is almost self-inverse, meaning that (twice) the original bit can be obtained by multiplying the right hand side with the matrix again. Moreover, any common mode (i.e., addition of both x and y by the same value) would appear in the first coordinate. Since this coordinate was 0 to begin with, and does not carry any information, it can be discarded. Therefore, in this case what is important is to take only the difference x-y in order to recover the bit b. This is what is typically done by a differential comparator.

A closer look at this matrix reveals the main properties that are needed to construct NRZ: the matrix should be close to orthogonal (i.e., the product of the matrix with its transpose should be diagonal, preferably with the same elements along the diagonal), and the sum of the columns of the matrix should be zero except at the first position. The second property is needed to ensure common mode resistance of the receiver. Moreover, it is needed to ensure that the sum of the values on the wires is 0, a property that is important for both the reduction of electromagnetic interference (EMI), and common mode resistance of the ensemble of wire values.

A class of matrices that satisfy this property is that of "Hadamard matrices." In fact, the matrix used to derive NRZ is the first member of this family. The next one is a 4 x 4 matrix, and the associated signaling, ENRZ, is given by the following formula:

$$\frac{1}{3}(0, a, b, c)\begin{pmatrix}1 & 1 & 1 & 1\\ 1 & -1 & 1 & -1\\ 1 & 1 & -1 & -1\\ 1 & -1 & -1 & 1\end{pmatrix} = (x, y, z, u)$$

Here, *a, b, c,* are the bits to be transmitted (modulated as +1 or -1) and *x, y, z, u,* are the values on the wires, forced to be between -1 and 1 because of the multiplication with the factor 1/3. The resulting vectors (*x,y,z,u*) are called the codewords of ENRZ signaling. Running through all possible values of *a, b, c,* the codewords are computed to be:

$$\pm\left(1,-\frac{1}{3},-\frac{1}{3},-\frac{1}{3}\right),\pm\left(-\frac{1}{3},1,-\frac{1}{3},-\frac{1}{3}\right),\pm\left(-\frac{1}{3},-\frac{1}{3},1,-\frac{1}{3}\right),\pm\left(-\frac{1}{3},-\frac{1}{3},-\frac{1}{3},1\right).$$

Another way of viewing the ENRZ codewords is as a superposition of the rows of the Hadamard matrix. Visualizing every row (starting with the second one) as a "spatial waveform," the signals on the wires will be superpositions of these waveforms multiplied with +1 or -1. These waveforms are shown on the Figure 3 on the right. The amplitude of each of the rectangular waveforms is $v/3$, where v is the largest voltage level allowed on the wires (corresponding to the nominal value of 1).

These spatial waveforms are "orthogonal" in the sense that the inner product of the corresponding vectors is 0. This is a consequence of the orthogonality of the Hadamard matrix. This condition allows for an easy detection and extraction of the bits: they can be detected by computing an inner product of the received signals with the waveforms on the right, or equivalently, with all rows (but the first) of the Hadamard matrix. In this respect, ENRZ is a simple generalization of NRZ: while for NRZ signaling detection is done by taking the inner product of the wire values with the second row of the 2 x 2-Hadamard matrix, for ENRZ we need to take the inner product with the last three rows of the 4 x 4-Hadamard matrix.

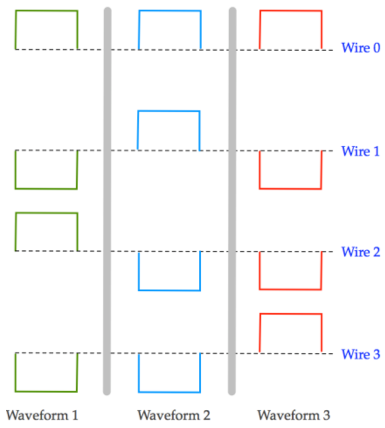

Figure 3: ENRZ waveforms superpositions.

This multiplication results in computing the values (*x*+*z*)-(*y*+*u*), (*x*+*y*)-(*z*+*u*), (*x*+*u*)-(*y*+*z*). Stated schematically, the receiver for ENRZ thus has the conceptual form given in Figure 4.

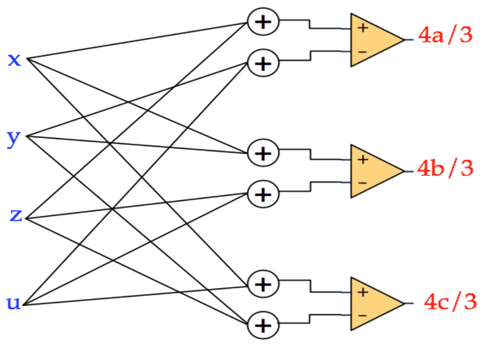

Figure 4: Conceptual comparator structure for ENRZ. Each comparator compares sums of two of the wires against sums of the other two.

The output of this receiver increases the swing of the incoming bits by 4/3 the same way that the output of the differential comparator increases the swing of the incoming bit by a factor of 2. With this architecture, ENRZ's vertical opening is about 3.5 dB worse than the vertical opening of NRZ at the same clock frequency. In almost all practical situations, the channel attenuation experienced by NRZ is much more than 3.5 dB worse than the attenuation experienced by ENRZ. The vertical eye of ENRZ is therefore in most practical situations much better than the vertical eye of NRZ at the same throughput per wire.

Some Properties of ENRZ

ENRZ has a number of interesting properties as discussed below.

Intersymbol Interference (ISI). ISI is a frequency dependent effect caused by the inertia of the wires to facilitate the propagation of electromagnetic waves. A rectangular wave form created by the transmitter for one unit interval (UI) will disperse into subsequent (and even previous) UI's as a result of this phenomenon. The received signal at a given UI is therefore a combination of the signal transmitted at the corresponding UI, and the residual signals of the previous UI's. Typically, the effect of ISI is worse for signaling methods that employ more than 2 levels on the wires. However, this is not true for ENRZ. Because the output of the comparators is binary, the signaling scheme has the same vulnerability to ISI as NRZ would have for the same clock frequency. The exact derivation of this important fact is beyond the scope of this note, though. Kandou Bus has developed a FOM (Figure of Merit) for the effect of ISI. The higher the FOM, the higher is the susceptibility of a coding scheme to ISI noise. NRZ and ENRZ both have an FOM equal to 1, while PAM-4, for example, has an FOM equal to 3.

Transmitter Driver power. The power consumed by the ensemble transmitter driver using ENRZ signaling is proportional to the sum of the currents on the wires from one end to the next. We can introduce a FOM for driver power consumption, normalized to be equal to 1 for one differential lane. For the same peak-to-peak signal on an ENRZ interface, the FOM for the driver power consumption of ENRZ is also equal to 1: At any point in time, there is a current of strength 1 in one direction. But since ENRZ transmits three bits per time interval, this means that the driver power consumption per bit for ENRZ is one third that of NRZ. This reduction of power comes at the price of lower swing at the receiver compared to NRZ. As was mentioned above, this reduction of the swing is typically compensated by the fact that the channel attenuation experienced by NRZ at 1.5 times the frequency is much higher than the same for ENRZ.

Wire routing and crosstalk analysis. Routing wires as loosely coupled pairs with some distance between the pairs to avoid crosstalk between the pairs (the way it is usually done for NRZ) is one of the optimal routing methods for ENRZ signaling. In this case, the three binary sub-channels at the output of the comparators don't experience any crosstalk among one another. The reason for this is hidden in the structure of the Hadamard matrix, and will not be further detailed here. Another routing constraint to consider is wire skew. In differential NRZ signaling schemes attentions is paid to close skew matching of the individual n- and p-wires of a pair, although pairs are often not matched. For ENRZ signaling the usual differential skew matching constraints are now applied to the ensemble of 4-wires. Thus existing wire routing rules can be applied simply to ENRZ signaling. Of course many systems include one or more connectors in a channel. For ENRZ a connector should be chosen that manages skew across its pins; such connectors are readily available in the market.

Electromagnetic interference (EMI) noise. One of the most interesting properties of ENRZ is its superior immunity to far field EMI noise. Analysis and simulation of representative channel structures shows that ENRZ EMI performance is better than the corresponding NRZ EMI performance for the same data throughput, and well below FCC limits.

Performance on a Measured Channel

The performance of ENRZ signaling versus NRZ and PAM-4 signaling is demonstrated in this section using Kandou's statistical eye analyzer tool "KEYE™". The channel used for these simulations is TE Connectivity's STRADA Whisper™ Backplane Channel (http://www.ieee802.org/3/100GCU/public/channel.html). This channel consists of 2 loosely coupled differential pairs with a trace length of 30 inches on a Megtron 6 HVLP board. The two pairs are identical, with no crosstalk between them. No alien NEXT/FEXT was considered. The trace geometry is stripline and the trace width is 6 mils. The differential insertion loss of this channel is approximately 32dB at a frequency of 25GHz.

For the statistical simulations we used Tx FIR with 1 pre- and 1 post-cursor. The CTLE and Rx noise filter was implemented according to 802.3bj D2.2 (with poles and zeros changed relative to the baud rate). In addition, a 2-tap DFE was used. A single-ended peak-to-peak amplitude of 600 mV was used in all our simulations. An aggregate throughput of around 50 Gbps was targeted for all three signaling schemes, leading to a transmission of 25 GBaud for NRZ (hence 50 Gbps over 4 wires), 12.5 GBaud for PAM-4, and 16.7 GBaud for ENRZ (leading to an aggregate throughput of 50.1 Gbps over the 4 wire interface).

Performance Comparison of NRZ, PAM-4 and ENRZ for a 50Gbps Throughput

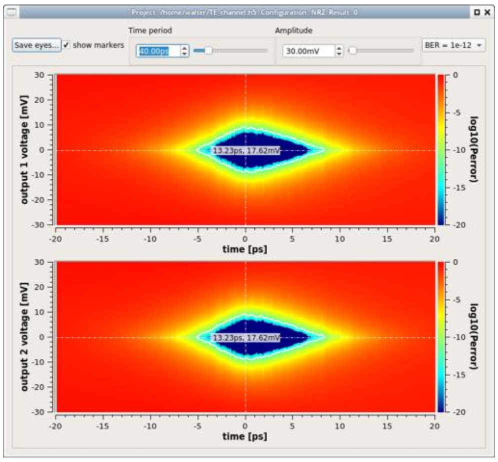

Figure 5:
NRZ Signaling Statistical Eyes

Channel Eye	Horizontal Opening	Vertical Opening
Eye 1	13.23 ps	17.62 mV
Eye 2	13.23 ps	17.62 mV

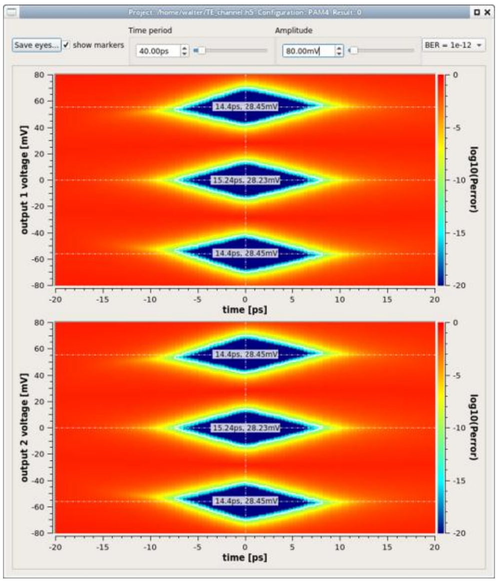

Figure 6:
PAM-4 Signaling Statistical Eyes

Channel Eye	Horizontal Opening	Vertical Opening
Eye 1	14.40 ps	28.45 mV
Eye 2	15.24 ps	28.23 mV
Eye 3	14.40 ps	28.45 mV
Eye 4	14.40 ps	28.45 mV
Eye 5	15.24 ps	28.23 mV
Eye 6	14.40 ps	28.45 mV

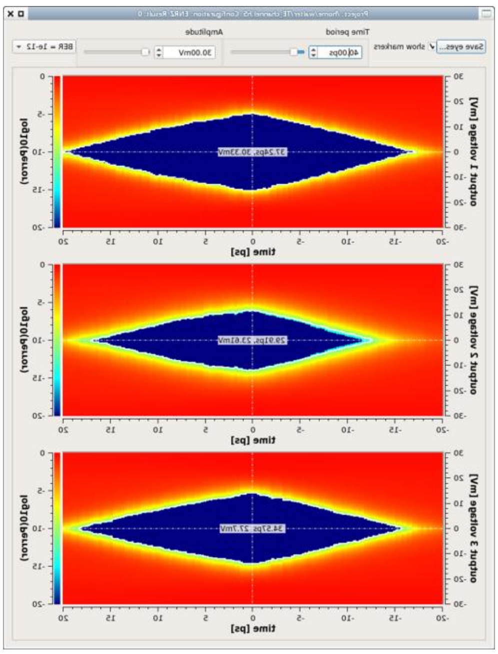

Figure 7:
ENRZ Signaling Statistical Eyes

Channel Eye	Horizontal Opening	Vertical Opening
Eye 1	37.24 ps	30.33 mV
Eye 2	29.91 ps	23.61 mV
Eye 3	34.57 ps	27.7 mV

Figure 5 shows the statistical eye diagrams for NRZ signaling on the channel described above. Each of the eye diagrams corresponds to one of the pairs.

Figure 6 shows the results of the simulations for PAM-4 signaling on the same channel as described above. In this case, to each pair of wires there are 3 associated eye diagrams, one for each level and the eye diagrams are identical for each pair.

Finally, Figure 7 shows the simulation results for ENRZ. In this case, the comparators used for detection don't compare sums of wire values against one another (as described in Figure 2), but instead they compare averages of two wire values against the other two values. This results in smaller vertical eye openings that could be achievable. Nevertheless, it can be seen that the vertical eye opening is reasonably large, and the horizontal eye opening is the largest among all the signaling methods compared, with values ranging between 29 and 37 psec. The horizontal width of the ENRZ eye is almost double that of PAM-4 despite the higher clock frequency used for ENRZ. This is primarily due to the fact that ENRZ is much less susceptible to ISI because it is inherently a binary coding scheme (despite the quaternary values used on the wires).

Another interesting point to note is the difference in horizontal eye opening between the second eye and the other two eyes. The main reason for this difference is that the second eye corresponds to the bit modulated via the row (1,1,-1,-1) of the Hadamard matrix. For optimal transmission of this bit the common mode of the signals should be passed along through all stages of the channel. This particular channel experiences a higher loss for the common mode than for the differential mode.

The table below summarizes the results above relative to one another by comparing the worst-case eye width/height of ENRZ with those of NRZ and PAM-4 as ratios.

	ENRZ vs NRZ	ENRZ vs PAM-4
Eye width	2.26	2.07
Eye height	1.34	0.83

Current and Emerging Standards

Industry Standards organizations such as OIF-CEI and IEEE 802.3 have created standards for high speed serial links of many years and generations of technology. As demand for higher data throughput increase new approaches are required by industry and the Kandou ENRZ signaling technology is a candidate for emerging Standards and Kandou is actively engaged with the Standards organizations.

Conclusions

This white paper has introduced a new signal coding technology, Ensemble NRZ (ENRZ), for high speed serial interfaces. The technology has been shown to have significant advantages in pin efficiency and signal performance over existing NRZ and PAM-4 signaling thus offering system power and data throughput gains whilst continuing to use known interconnect technologies.

About Kandou Bus

Kandou Bus (http://www.kandou.com) is a fabless semiconductor IP company specializing in the design of high-speed, pin-efficient, energy-efficient serial links, SerDes, and associated technologies based on correlated data coding and signaling techniques. Kandou Bus follows an IP business model with three typical elements of non-recurring engineering revenues for custom work, and licensing fees/royalties in case of productization.

Kandou Bus has development centers in Switzerland and the UK.

For more information please contact info@kandou.com.

V1.0, Jan-14.