Measuring and Evaluating Best Source Selection

Diarmuid Corry MSc. (Eng)
ACRA CONTROL INC

ABSTRACT

To properly evaluate and characterize the performance of a bit synchronizer we need to apply a known data stream and then adjust several interference parameters to measure the effect on synchronization performance: white noise, offset and gain variations and frequency and phase shifts.

The task becomes more complex when we consider the performance of a best source selector (BSS) which combines the performance of two or more bit synchronizers to achieve better bit error rates and more consistent synchronization than can be achieved with one alone. Each of the parameters (noise, offset, gain, phase) are often different for each bit synchronizer, and may vary over time. In addition the incoming bit streams can drift in time (possibly 100s of bits) with respect to each other.

This paper discusses how these parameters are measured, and looks in particular at the problem of evaluating a BSS. Results showing the performance that can be achieved when aligning and combining multiple streams are presented and discussed.

1. INTRODUCTION

The theoretical bit error rate (BER) for a bit synchronizer in the presence of additive white Gaussian noise has been understood since the 1940s. However, it is only in the last decade or so that devices operating close to that limit have been available. This was made possible due to the availability of high speed analog to digital converters and logic capable of implementing phase locked loops (PLLs), matched filters and Inter-symbol Interference (ISI) equalizers in real time.

One approach to improving RF data links is to use superior bit shapes (for example ARTM Tier II) to send more bits across existing bandwidth allocations – perhaps using the additional bits for error correction.

Another technique is to use diversity techniques to combine signals from multiple receivers and/or transmitters. This papers looks briefly at issues involved in testing traditional bit synchronizers and how these techniques can be applied to source selectors/combiners.

2. ONE dB FROM THEORY

In the 1940s Shannon et al showed that there is theoretical best case bit error rate (BER) for symbols transmitted in noisy channels – more noise means more bit errors for the same signal power. This is traditionally graphed as in Figure 1 below, where Eb is the energy per bit and No is the white noise power.

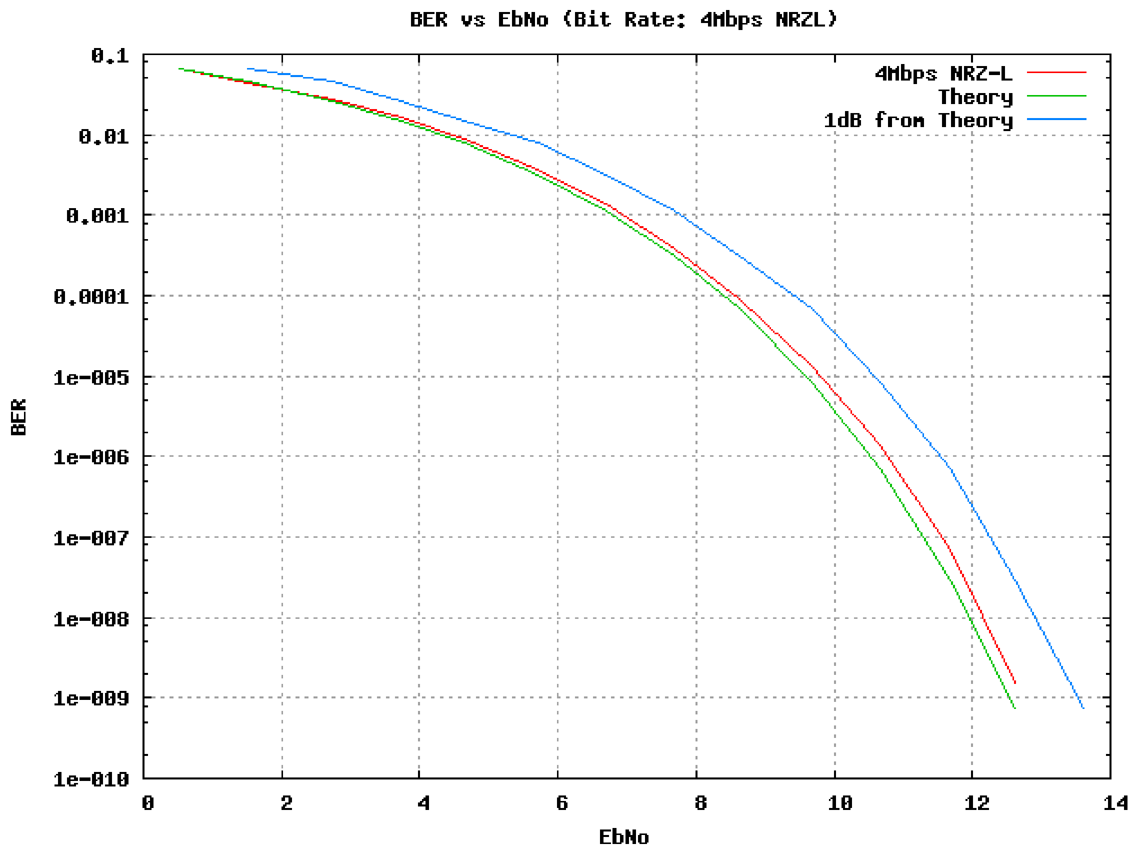

Figure 1 Bit Error Rate (BER) versus Eb/No – Theory, Theory + 1dB, practice

The graph also shows the BER for a device "1dB of theory" – a decade ago this was about the limit of what was technically achievable at acceptable cost. The graph also shows the typical BER for a modern bit-synchronizer using state-of-the-art techniques (shown in red). To put these noise levels and bit error rates in context the graphs below show oscilloscope displays of the same NRZ-L signal with noise levels (Eb/No) of 3.7 dB and 12.6 dB respectively.

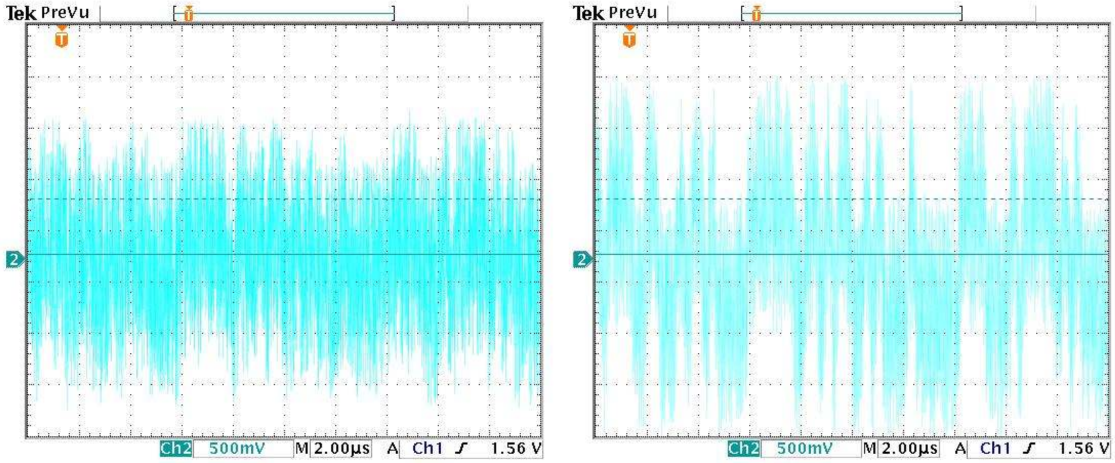

Figure 2 4Mbps RNZ-L PCM stream with Eb/No of 3.7dB and 12.6 dB

Modern bit-synchronizers rely on two techniques to achieve results close to theory: matched filters and decision directed equalization.

From the start it was known that to achieve close to theory a matched filter was required, for example the traditional "integrate and dump" technique matches a rectangle pulse shape well. When "integrate and dump" is applied to PCM shapes filtered using a 6^{th} order Bessel filter, or a root raised cosine shape, it will cause a BER curve typically 0.5dB from theory. It is also worth mentioning that today's fast A/Ds mean that matched filters can be implemented digitally.

Another source of error in traditional bit synchronizers is due to the inherent inter-symbol interference (ISI) of many PCM pulse shapes. Rectangle shapes and root raised cosine shapes do not have this problem but "real life" shapes such as those caused by Bessel filters do. If a bit is preceded by a "1" and followed by a "0" then the effects cancel – however, if the bit ahead and behind are the same then the threshold used to choose between "1" and "0" changes.

One method to combat this is knows as decision directed equalization. Here the decision threshold used for each bit depends on what the previous and following bit is assumed to be. Often these bit values are also not known with great confidence so the bits before those have to be used we work backwards and so on.

3. OTHER METRICS

When implementing these techniques, the task of testing and validating a bit-sync presents some problems. For example, when graphing BER curves we found that in certain circumstances our results were “better than” theory. Spectral analysis of the so-called white noise sources from three vendors showed that they were not true White Gaussian Noise. In the end, equipment from a fourth vendor provided good results.

Additive white Gaussian noise (AWGN) is not the only metric associated with bit synchronizers. What follows is a brief discussion of some other metrics.

Acquisition range and number of bits

This metric is concerned with the number of bits required to achieve an appropriate BER after a signal has been applied to the input of bit synchronizer and its dependency on signal bitrate variance. For example if a bit synchronizer is programmed for 4Mbps and you suddenly apply a signal at 4Mbps ±1, 2, 3 or 5% of that, this metric measures the number of bits it takes to achieve a target BER.

Tracking range

This metric is concerned with the bit sync’s ability to track a percentage deviation from the bit rate. For example, if a bit synchronizer has acquired a signal that subsequently varies to ±20% of the bit rate with a modulation frequency of 0.1%, this metric measures how well the bit-sync responds.

Baseline drift

This metric is used to determine the bit sync’s response to a baseline drift. For example, if a sinewave of 0.1% of the bit rate is added with the same amplitude as the signal this metric determines the effect on the BER.

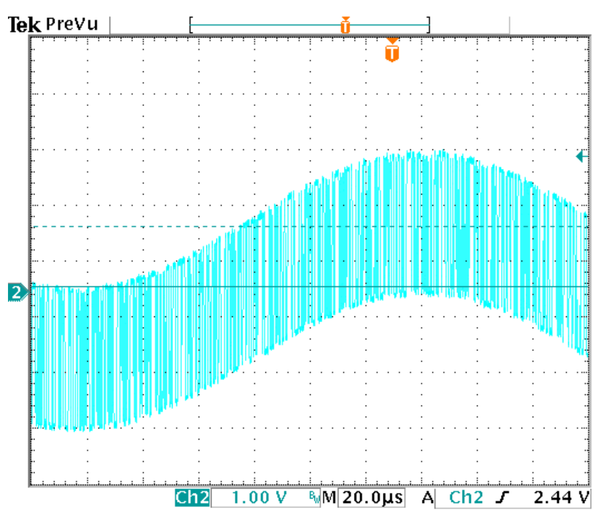

Figure 3 PCM with baseline drift

Gain drift

This metric is used to determine the bit syncs response to a gain drift. For example, if the amplitude is modulated so that it varies from 50% to 200% with a frequency of 0.1% of the bit rate, this metric measures the effect on BER.

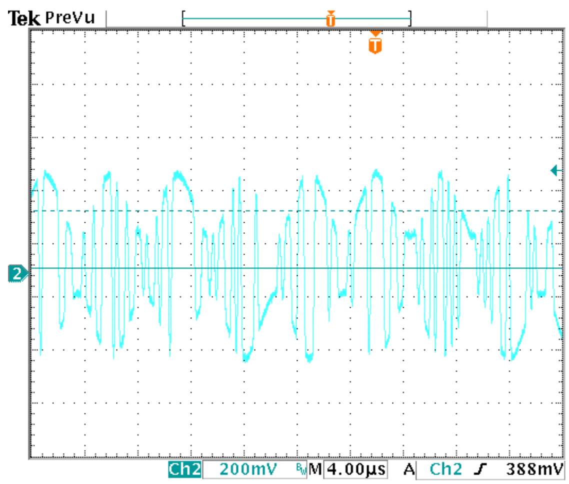

Figure 4 Amplitude modulation or fading

Consecutive "1"s

This metric is concerned with how a bit sync reacts to patterns of 1s and 0s. Take the three scenarios for different encoding schemes as examples:

- If, for NRZ-L, there is 128 consecutive "1"s or "0"s will the bit synchronizer see 127 or 129?

- If for NRZ-L there are consecutive "1100" bit patterns will the PPL lock to half the bit rate.

- If, for BI0-L, there are consecutive "1010" bit patterns will the PLL lock to half the bit rate.

4. TEST SETUP

The tests were conducted using a single channel modern bit synchronizer which was fed from a custom designed and built test jig. This jig is shown in Figure 5 below. It consists of a PCM encoder with programmable PCM code (RNRZL, BIO-L), programmable bit shape (rectangle, Bessel, RRC) and programmable bit rate. The PCM encoder outputs 128 consecutive 1s every 1023 bits.

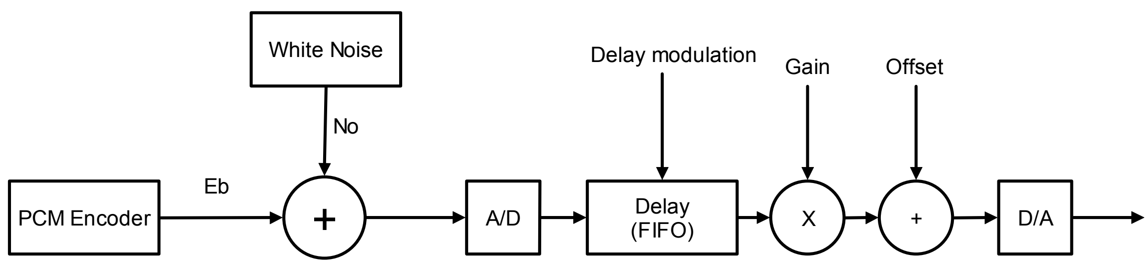

Figure 5 Simple test jig to add Noise, delays, gain and offset drifts

This is followed by a programmable noise source where Gaussian white noise is added.

The high speed A/D samples the signal plus noise and delays it by a number of bits, then modulates the amplitude and adds an offset. The number of bits delayed, the modulation amplitude and offset amplitude are each programmable as is the rate of change of each. Two tests were conducted to examine the effect of different metrics on the BER.

The BER rate graph shown if Figure 6 was obtained for the following test conditions

- a 4Mbps RNRZ-L stream 1088 bits long
- an offset of ±100% amplitude @ 4kHz (baseline drift)
- a delay of ±20 bits @ 10Hz

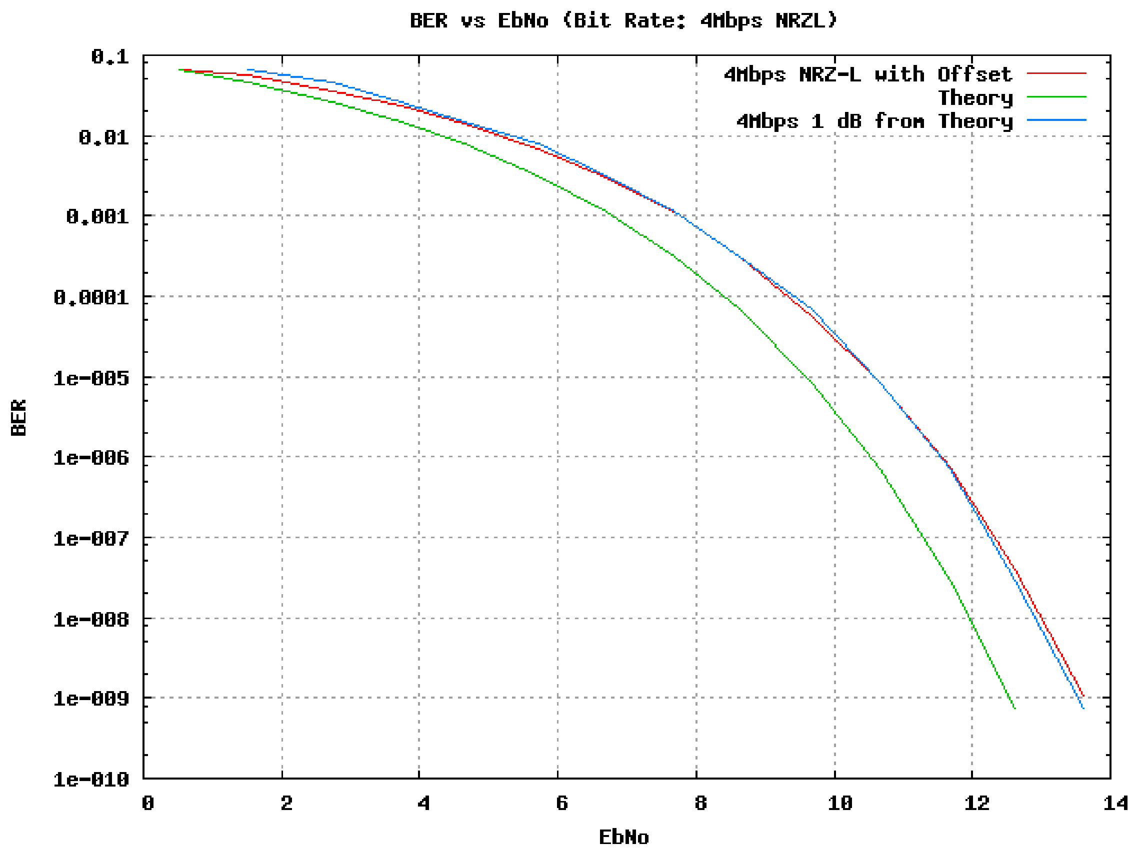

Figure 6 Bit Error Rate (BER) versus Eb/No – Theory, Theory + 1dB, signal with offset of ±100% amplitude @ 4kHz

The BER rate graph shown if Figure 7 was obtained for the following test conditions

- a 4Mbps RNZ-L stream 1088 bits long
- a 30% amplitude modulation @ 200kHz
- a delay of ±20 bits @ 10Hz

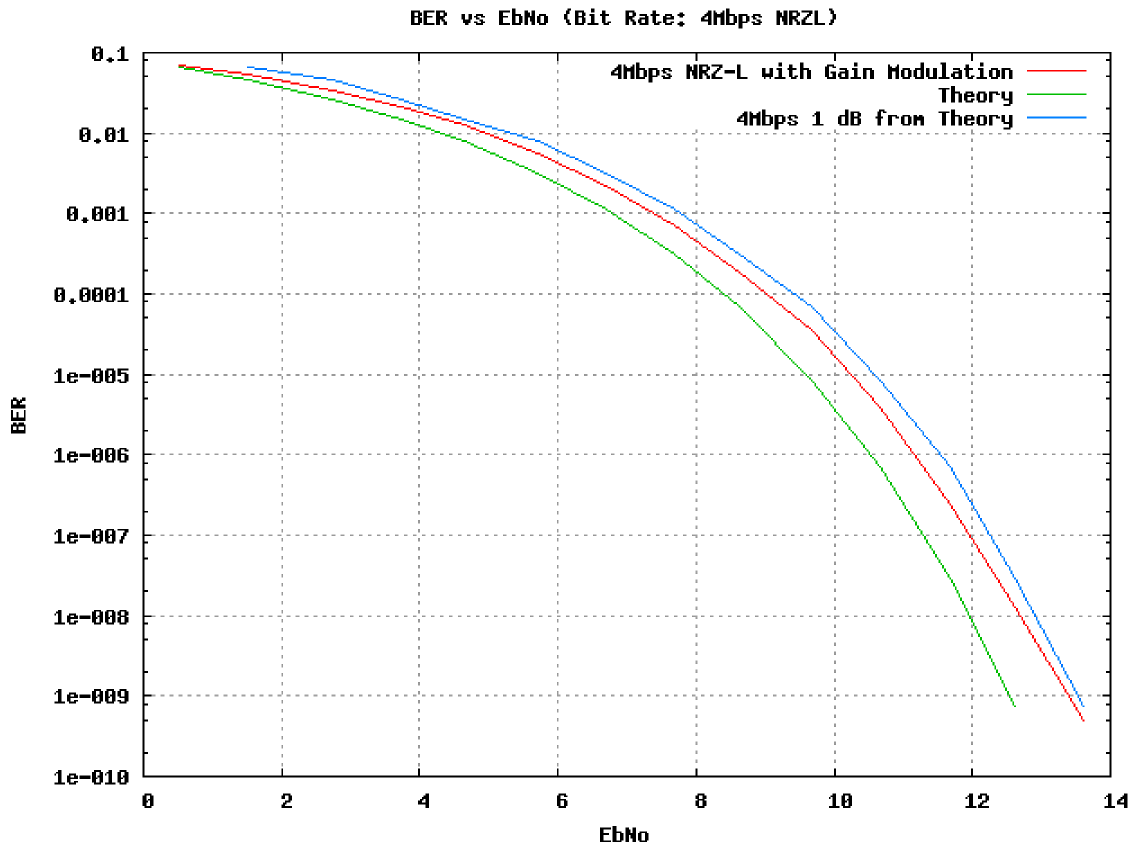

Figure 7 Bit Error Rate (BER) versus Eb/No – Theory, Theory + 1dB, signal with 30% gain modulation @ 200kHz

5. SOURCE SELECTION

As mentioned above there is a theoretical limit to how few bit errors can be obtained from single PCM stream in the presence of noise, but what if there were more than one PCM stream with the same information?

Some FTI programs transmit data on two carrier frequencies, one transmitter on the top front of the plane and one on the bottom, this is called frequency diversity. Some receivers on the ground are placed apart or at different altitudes, this is called spatial diversity. Some receivers support polar diversity.

On the ground there is hardware that uses these two sources to recover the transmitted signal. "Best source selectors" switch to the stream with the best signal to noise ratio while "best data selectors" switch to the stream with the fewer data (syncword) errors. In both of these cases care must be taken on switching such that downstream decoders remain in lock. Best data selectors require decryptors on each channel if encryption is used.

Correlated source selectors time align different streams then vote on a bit by bit bases – possibly using error metrics to weight the voting of each channel.

Smart source selectors is the term given to devices that do not simply switch or vote but rather use the best of both worlds – plus soft bit values.

A smart source selector first selects only those channels with a signal level above a certain threshold, then selects only those within 5dB of the best Eb/No ratio. The selected streams are then time aligned and each bit combined using the soft bits for each stream. Note: this scheme works with encrypted data and downstream decoms are not affected by stream switching.

6. SMART SOURCE SELECTION IN ACTION

TO test the effectiveness of smart source selection, a single PCM stream was split in two and, using two test jigs, independent white noise, offsets, gains and delays added to each instance of the stream. The two new streams were used as inputs to the smart source selector and the resultant BER is graphed. The result is shown in Figure 8 below.

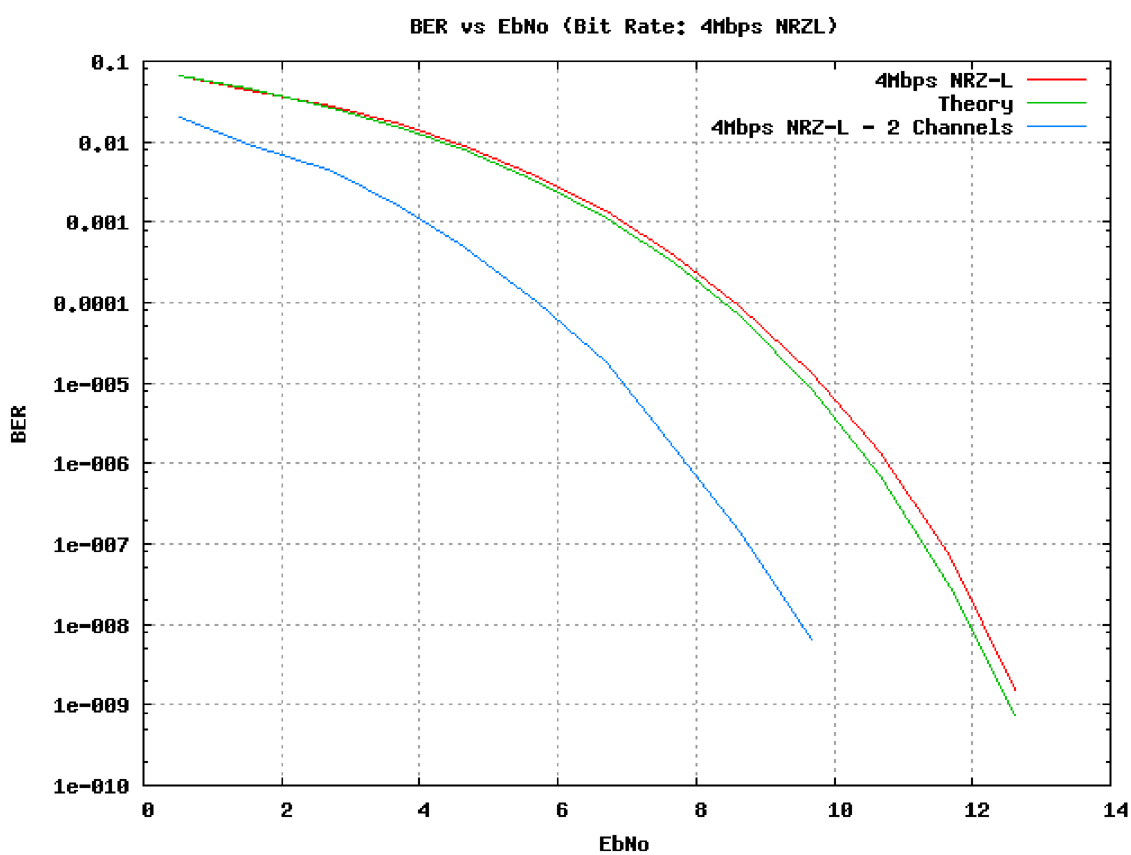

Figure 8: BER for a smart source selector compared to either input on its own

For two channels the resultant error rate is approximately 3dB better than for either channel on its own. For 4 channels it would be 6dB, for 8 it would 9dB. This is an improved strategy over best source or best data selectors as they output a resultant BER that is equal only to the better of the two, not better than either.

One explanation for this is that if both channels have one error per thousand and these are independent, then the combination would usually only be wrong when both are wrong (one per million). Put another way, with two channels you have twice the signal power but the noise power, being independent, remains the same.

7. SOME LESSONS LEARNT

Tests show that if the “poorer” signal has an Eb/No within 5dB of the “better” signal than the results of combining remain better than either signal alone. Because the aligner is designed to operate with encrypted data it is better to ensure that there no long strings of “1”s or “0”s if NRZ-L are used so that “false” alignment is avoided when the signal is first acquired. However, even if this does happen the correct alignment will happen on the first transition.

The additional improvement with two inputs combined takes approximately 1000 bits due to the time taken to align the two streams initially. Until aligned, BER performance is the same as a single channel bit-sync.

Initial results indicate that any benefit in incorporating best data techniques for stream selection may not justify the additional complexity.

8. CONCLUSION

With modern digital techniques, today’s bit synchronizers achieve errors rates very close to theory. In particular high speed A/Ds, 100% digital PLLs, matched filters and modern equalization techniques mean that to improve performance other techniques must be explored. One approach is to use digital transmitters and receivers to transmit more bits over the same bandwidth and perhaps use the extra bits for error correction. Another approach explored in this paper is to use frequency, spatial or polar diversity techniques and smart source selection to obtain fewer errors with existing RF infrastructure.