OBTAINING SUPERIOR PERFORMANCE FROM DUAL-CHANNEL RECEIVERS USING BEST-CHANNEL SELECTION

Jim Uetrecht
Quasonix, Inc.
West Chester, OH, 45069
jimu@quasonix.com

ABSTRACT

A Pre-Detection Maximal-Ratio Combiner (Pre-D MRC) provides optimal performance in the face of additive noise but sub-optimal and even potentially degraded performance in the face of other common channel impairments such as multipath. Using Data Quality Metric (DQM) information from the Pre-D MRC and its two input channels, a Best-Channel Selector (BCS) can correlate the three demodulated data streams and select the best available data on a bit-by-bit basis. Thus, the BCS can produce a single receiver output with optimal performance in additive noise and superior performance across all types of channel impairment – fulfilling the promise the Pre-D MRC cannot always keep. Further, since the BCS need only accommodate small and predictable latency differences between its inputs, its local performance may exceed that of an external source selector designed to handle seconds of delay between channels. This paper describes the BCS and presents performance results from several test scenarios.

INTRODUCTION

Many telemetry applications employ dual-channel receivers to provide signal diversity. Common forms of diversity include polarization diversity, using the right-hand and left-hand circularly polarized feeds from a single antenna; frequency diversity, using two independent channels to receive the same data; and spatial diversity, using two spatially separated antennas to receive the same data. In any case, it is generally possible to combine the received signals from Channel 1 and from Channel 2 to create a single combined signal that can be demodulated with fewer bit errors than either Channel 1 or Channel 2. This is the role of the Pre-D MRC.

Because the Pre-D MRC can produce an improved signal, it is common to assume the Combiner data output will always be better than the Channel 1 or Channel 2 data output. However, typical telemetry channels suffer from signal impairments that may break this assumption. Relying solely on the Combiner data output can therefore lead to avoidable communication errors.

Several strategies may be considered to alleviate this problem. These strategies fall into two main categories: either change the combiner from traditional Pre-D MRC to a structure better able to avoid some forms of degradation, or add a Post-D source selector to select Channel 1 or Channel 2 output if the Combiner output is degraded. The first strategy sacrifices optimal performance of the Pre-D MRC, and the second strategy requires connection to and processing of the Channel 1 and Channel 2 data outputs. This latter approach, while preferable, sacrifices the simplicity of connecting to a single data output – unless the source selector can be integrated with the receiver. This is the role of the BCS.

MAXIMAL-RATIO COMBINING

It has long been known [1] that a Pre-D MRC provides optimal combining performance in the face of additive white Gaussian noise (AWGN). If the only impairment suffered by a telemetry channel were noise, the Pre-D MRC would always give the best possible answer. Unfortunately, this is simply not the case. A typical telemetry channel may be subject to multipath, interference, and other forms of signal corruption that can challenge or fool even a sophisticated Pre-D MRC.

To understand how these impairments degrade Pre-D MRC performance, first consider the basic structure of the Pre-D MRC:

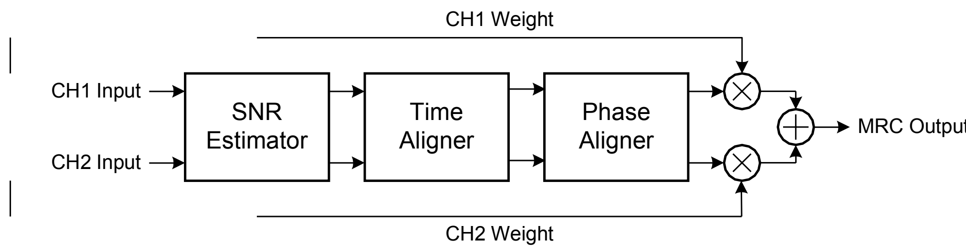

Figure 1 Combiner Operation.

To operate optimally, the Pre-D MRC must obtain an accurate estimate of the received signal-to-noise ratio (SNR) for both channels, align the two channels in time and in phase, and then sum the two channels with weights based on their relative SNRs. An error in any stage of this process will lead to suboptimal performance.

In the case of multipath, all three stages of the process are subject to failure. In particular, SNR estimation will probably not reflect an accurate measure of signal impairment, especially if the estimate is based solely on signal strength. Time alignment may also deteriorate as multiple copies of a signal arrive with different delays. Worse yet, when Channel 1 and Channel 2 experience different multipath, the Pre-D MRC output will be composed of all rays from both channels, which may be more difficult to equalize or demodulate than either channel by itself.

Likewise, interference can severely degrade SNR estimation. A strong interfering signal with a high SNR may fool the Pre-D MRC into weighting the interferer well above the desired signal on both channels, effectively throwing away all useful received information.

It is also possible for propagation effects to result in non-combinable received signals. For example, space-time coded (STC) waveforms are transmitted as two separate orthogonal signals, P0 and P1, that each contain the same data. With linearly polarized transmit antennas, P0 and P1 are normally received equally by right-hand and left-hand circularly polarized receive antenna feeds, so Channel 1 and Channel 2 both normally receive a mix of P0 and P1. However, certain multipath scenarios may cause the P0 and P1 signals to arrive at the receiver with circular polarization. In this case, the right-hand feed receives only P0 and the left-hand feed receives only P1, or vice versa. When this occurs, Pre-D MRC phase alignment of Channel 1 and Channel 2 fails catastrophically, as the two orthogonal signals cannot possibly be aligned.

BEST-SOURCE SELECTION

When the Pre-D MRC fails to operate optimally, the result is evident as reduced data quality and ultimately as bit errors. With the advent of data quality metric (DQM) and data quality encapsulation (DQE) standards, it is possible for a best-source selector (BSS) to observe this reduction in Combiner output quality and select Channel 1 or Channel 2 output data instead [2,3,4].

Assuming the cost and space for a BSS is justified, this seems like a nearly ideal solution. However, there are a few issues with this approach related to using all channels (Channel 1, Channel 2, and Combiner) from the same receiver.

Most simply, for a BSS to select Channel 1 or Channel 2 data rather than Combiner data, the BSS requires access to the Channel 1 and Channel 2 data streams in addition to the Combiner data stream. Connecting these streams consumes physical wiring resources and/or network bandwidth, either of which may be in short supply, or at best, inconvenient to provision and maintain.

More subtly, the optimal BSS [5] implements the log-likelihood ratio (LLR) decision criterion

$$\sum_{n\in\mathcal{N}_0} \log\left(\frac{1-p^n}{p_n}\right) \gtrless \sum_{n\in\mathcal{N}_1} \log\left(\frac{1-p^n}{p_n}\right), \qquad (1)$$

where $\mathcal{N}_0$ is the set of source indexes for which a given demodulated bit is a 0, $\mathcal{N}_1$is the set of source indexes for which a given demodulated bit is a 1, and $p_n$is the bit error probability (BEP) for source n. This criterion is predicated on the assumption that source errors are statistically independent, which is clearly invalid when the sources are Channel 1, Channel 2, and the Combiner from the same dual-channel receiver. Worse, the degree to which the Combiner BEP is dependent on Channel 1 or Channel 2 BEP is unknown. Therefore, under some circumstances the optimal BSS may yield suboptimal performance.

Finally, correlating source selectors are typically designed to handle potentially large delays between input sources, as much as several seconds. However, the differential arrival delays from a single dual-channel receiver are guaranteed to be small (at least, local to the receiver). Assuming finite processing power, a smaller correlation window should lead to faster and/or stronger correlation.

BEST-CHANNEL SELECTION

The potential issues with using a BSS to cover for the Pre-D MRC are alleviated by placing source selection within the receiver and leveraging intrinsic knowledge of channel attributes, in particular, relative channel delays and potential BEP correlation. A high-level block diagram shows how the BCS augments traditional combiner functionality:

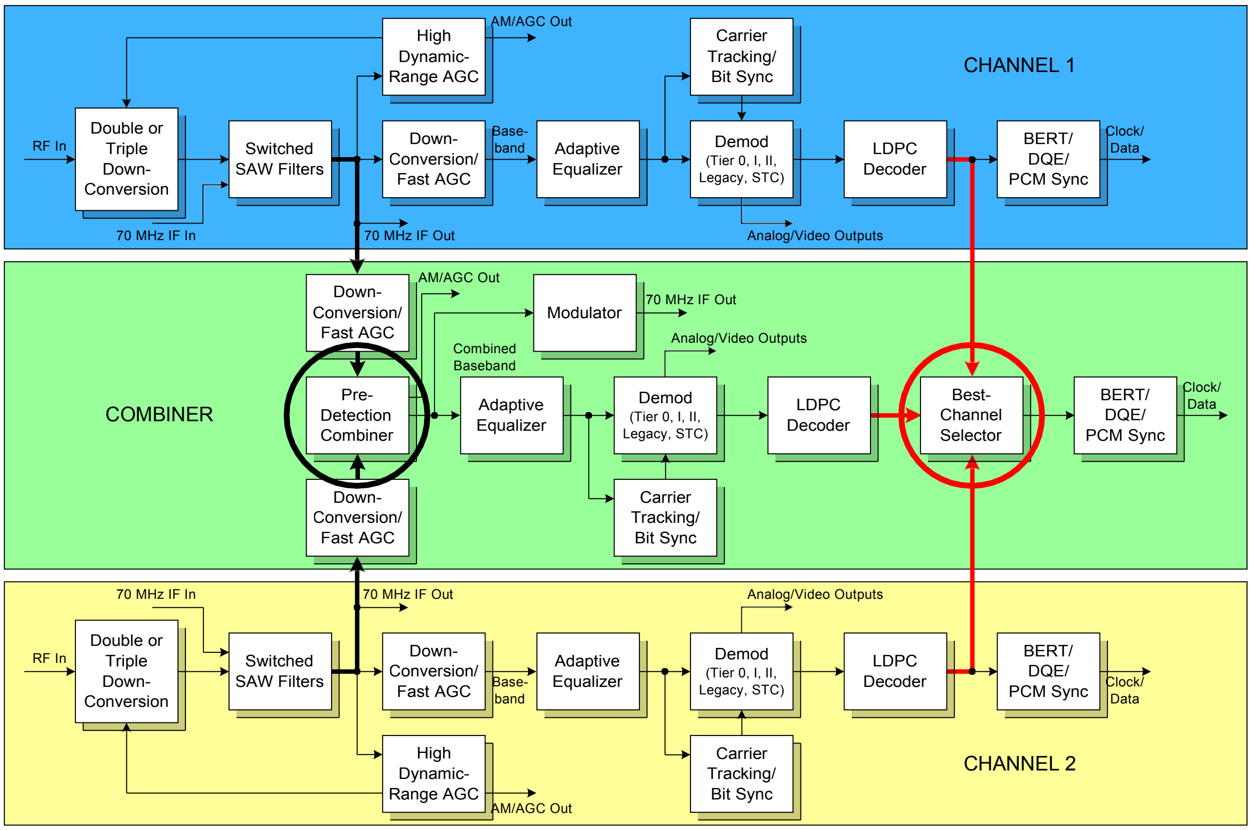

Figure 2 Dual-Channel Receiver Block Diagram.

The Pre-D Combiner maintains the form of an MRC, so it performs optimally when AWGN is the primary signal impairment.

The BCS looks like a mini correlating BSS:

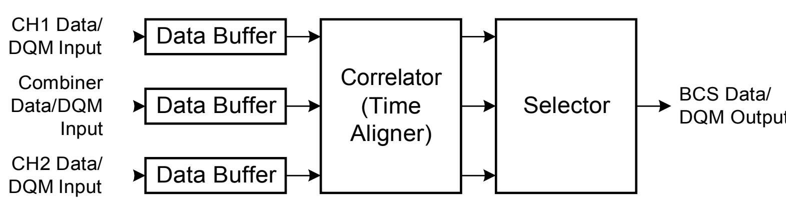

Figure 3 Best-Channel Selector Block Diagram.

The BCS correlates the Channel 1, Channel 2, and Combiner bit streams so that selection can switch at any time without introducing a slip in the BCS data stream. This process is enhanced by knowledge of the maximum possible delay difference between channels, which is very small and also predictable. Thus, the BCS exhibits exceptional dynamic performance, as dropped channels can be correlated quickly and accurately the moment they recover.

Once the channels are correlated, the BCS uses each source’s DQM to select the best data on a bit-by-bit basis. This is where the BCS has the biggest advantage over the Pre-D MRC. DQM is computed based on full knowledge of the waveform’s phase structure and the resulting likelihood that the received signal matches some valid (but unknown) bit sequence, so the BCS decision criterion is extremely robust.

In the present BCS implementation, the selection process can take one of two forms. The default selection criterion simply selects the channel with highest DQM. The alternate selection criterion is the LLR criterion from equation (1). Based on extensive testing, the max-DQM criterion generally performs comparably to the LLR criterion within a dual-channel receiver. More importantly, while the LLR criterion may unpredictably select a result worse than the best channel, the max-DQM criterion by definition cannot. Therefore, the LLR criterion is provided for test purposes only.

As a final step, the BCS generates an accurate DQM for the composite data stream, which can then be encapsulated for use by an external BSS. So, only one receiver output is needed to provide the best possible receiver data. In this way, the BCS and BSS can be best utilized for their respective strengths: highly dynamic selection to obtain superior performance from each dual-channel receiver (BCS), and highly elastic selection to allow distributed receivers to fully cover a range (BSS).

TEST RESULTS

Several test scenarios have been used to verify the efficacy of the BCS. A few of these are described here.

AWGN Test This test is straightforward: it applies signals with a known level of added noise and measures accumulated bit errors. Channel 1 and Channel 2 operate normally, but the Combiner is manually configured to ignore its Channel 2 input, so Combiner errors are highly correlated with Channel 1 errors.

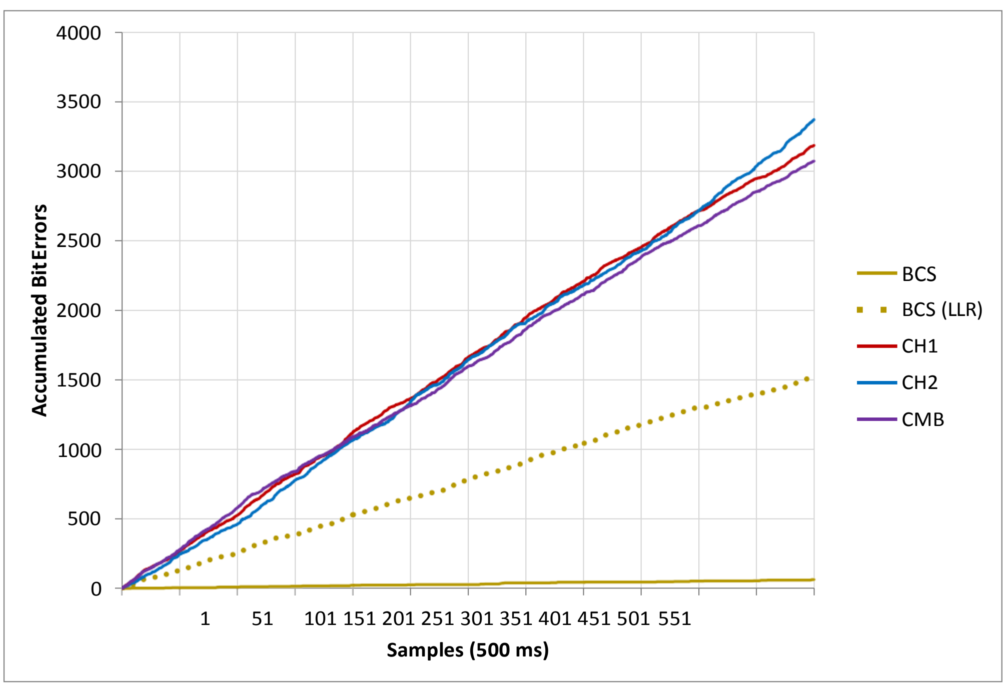

Figure 4 AWGN test accumulated bit errors – Combiner degraded due to unused CH2 input.

Figure 4 shows the results of this test with SOQPSK at 5 Mb/s and E_b/N_0 = 12 dB. Clearly, the BCS output is at least as good as the best channel. In fact, it is so much *better* than the best channel, the result is not immediately intuitive. However, this feat is readily explained.

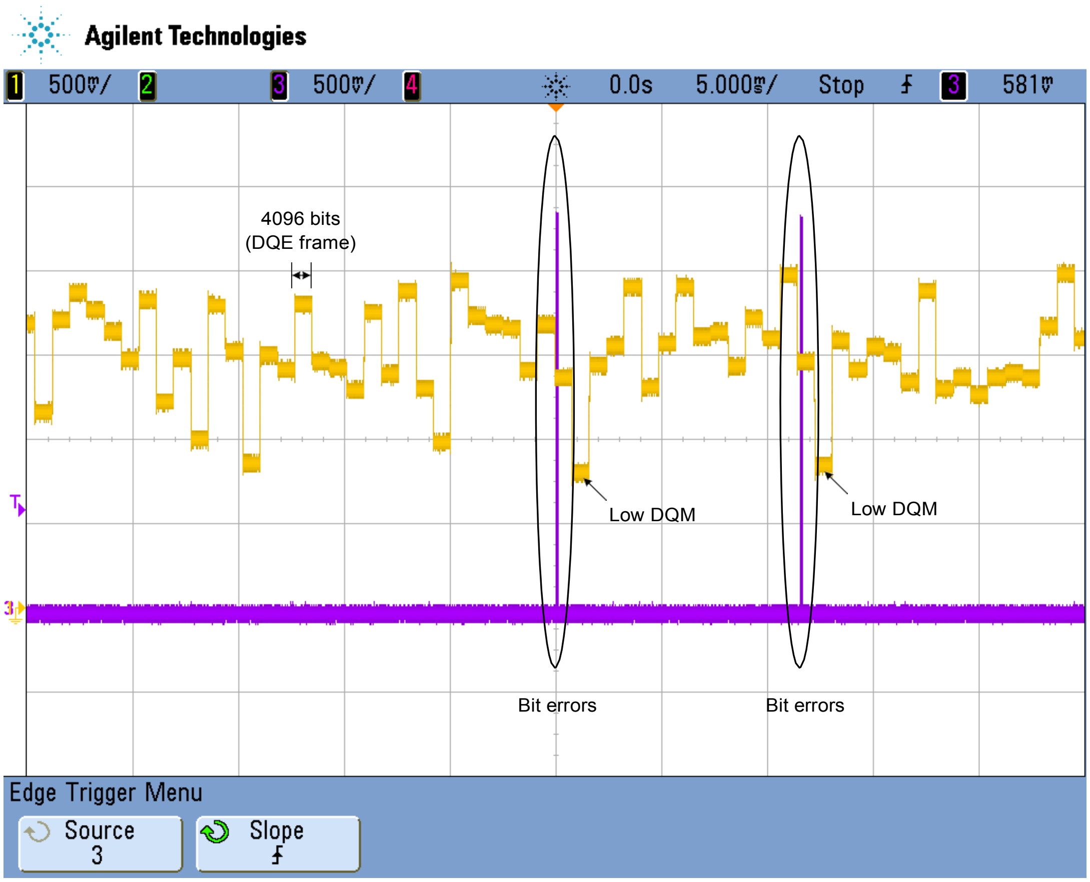

Figure 5 AWGN test real-Time DQM and bit errors.

Figure 5 shows a typical real-time capture of DQM and actual errors detected by a bit error rate tester (BERT) under the same test conditions as Figure 4. The DQM value is represented by the yellow trace, ranging from 0.0 V (BEP = 50%) to 2.5 V (BEP < 10^{-12}). Bit errors are represented by the purple trace, ranging from 0.0 V (no error) to 2.5 V (error). In this particular instance, two pairs of bit errors occurred, one pair at time t = 0.0 s, and another pair at time t = 11.5 ms.

Note that the DQM was lowest immediately following the bit errors; that is, the DQM was lowest corresponding to the DQE frames in which these bit errors would appear. Since DQM values are logarithmically related to BEP, these lower values dramatically delineate frames in which bit errors are likely. And because this delineation occurs in real time, the BCS can confidently select alternate sources on a frame-by-frame basis. Thus, "best" is virtually instantaneous best, not just average best. When errors are sparse, as in this example, the benefit is substantial.

Returning to Figure 4, note that both BCS selection criteria, max-DQM and LLR, were tested. As expected, the max-DQM criterion significantly outperforms the LLR criterion due to the high degree of correlation between Channel 1 and Combiner bit errors.

STC Test This test synthesizes STC waveforms to mimic the real-world propagation phenomenon mentioned previously. In reality, the P0 and P1 signals may acquire circular polarization, which leads to the circularly polarized feeds at the receive end delivering mostly P0 to one channel and mostly P1 to the other channel. In this test, P0 and P1 power is varied to achieve a similar result:

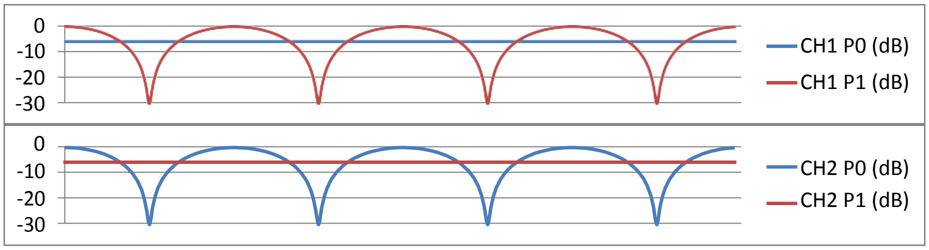

Figure 6 STC test time-varying input power.

During the oscillation, Channel 1 sees the ratio of P1/P0 power vary from roughly 4:1 to 1:250; Channel 2 sees the same ratios of P0/P1 power. The test is run near sensitivity, so Channel 1 and Channel 2 incur bit errors at the bottom of each oscillation. Due to combiner gain, the Combiner incurs few bit errors due to the power drop; however, at the same time, the Combiner cannot phase align its two inputs (P0 only from Channel 1 and P1 only from Channel 2) since they are orthogonal, so it has its own burst of bit errors. Thus, this scenario demonstrates the effect of uncorrelated Combiner errors.

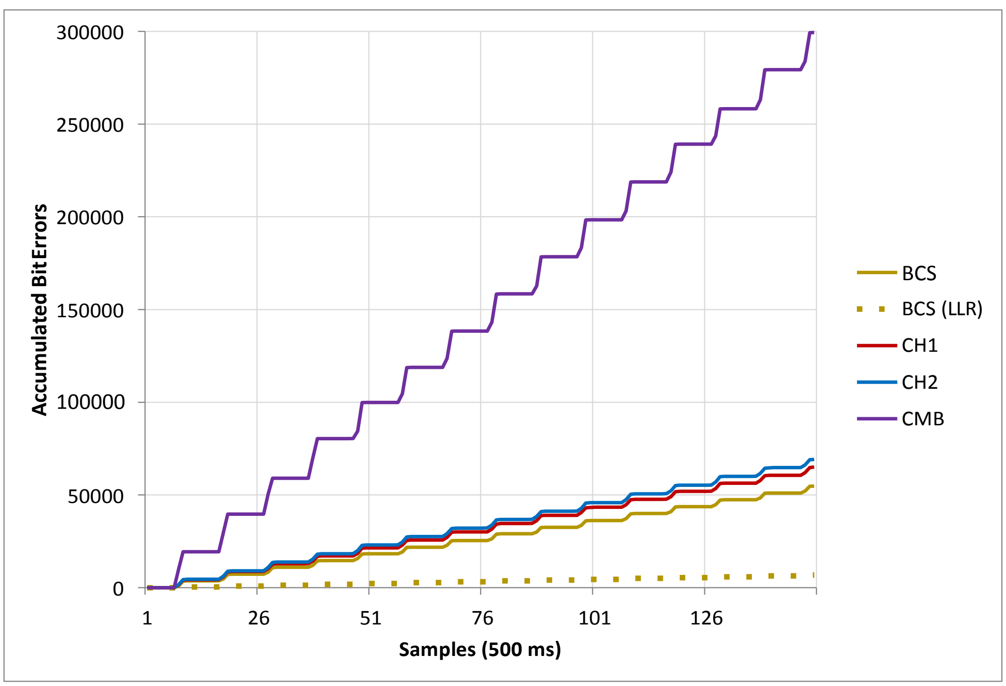

Figure 7 STC test accumulated bit errors – Combiner degraded due to periodic P0-only + P1-only orthogonal inputs.

Figure 7 shows the results of this test with STC at 5 Mb/s and net input power at -91 dBm. Again, the BCS output is better than Channel 1 and Channel 2, and far better than the Combiner, which is failing to do the impossible: combine two orthogonal signals. Once more, both BCS selection criteria were tested. In this case, the LLR criterion outperforms the max-DQM criterion due to the negligible correlation between Combiner bit errors and Channel 1 and Channel 2 bit errors.

Multipath Test This test generates independent time-varying multipath on the Channel 1 and Channel 2 inputs. Due to severity of the multipath, the test is run with adaptive equalization enabled. This scenario clearly shows potential for Pre-D MRC performance degradation due to a common channel impairment.

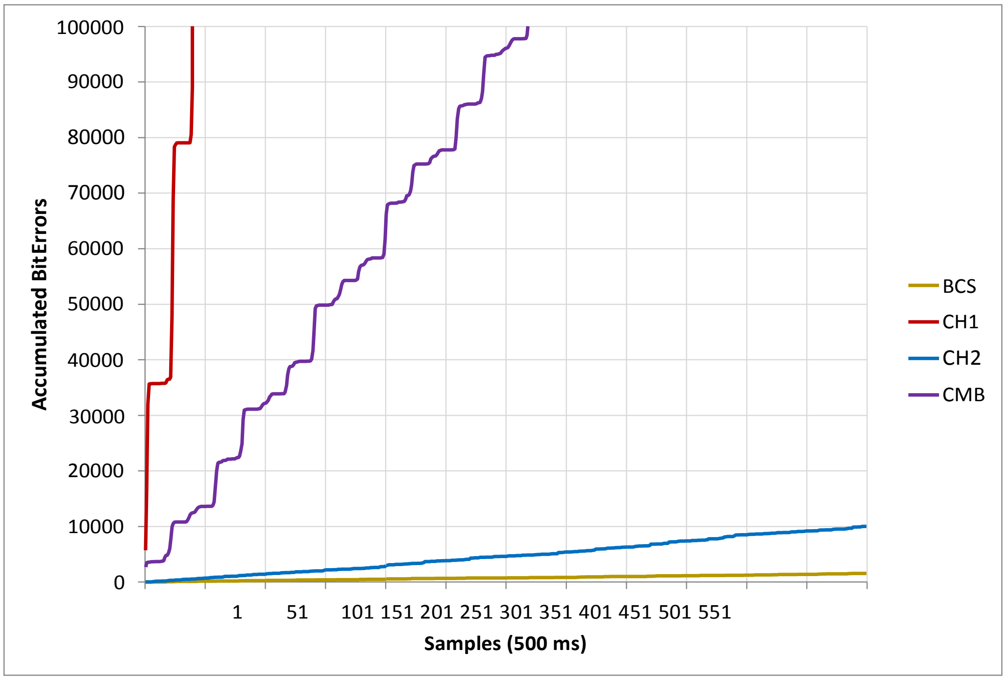

Figure 8 Multipath test accumulated bit errors.

Flight Test This test applies recorded signals from a 43-minute flight test at the Yuma Proving Ground. The recording includes two complete RF outages. So, this scenario demonstrates synchronization recovery as well as performance during challenging real-world link conditions.

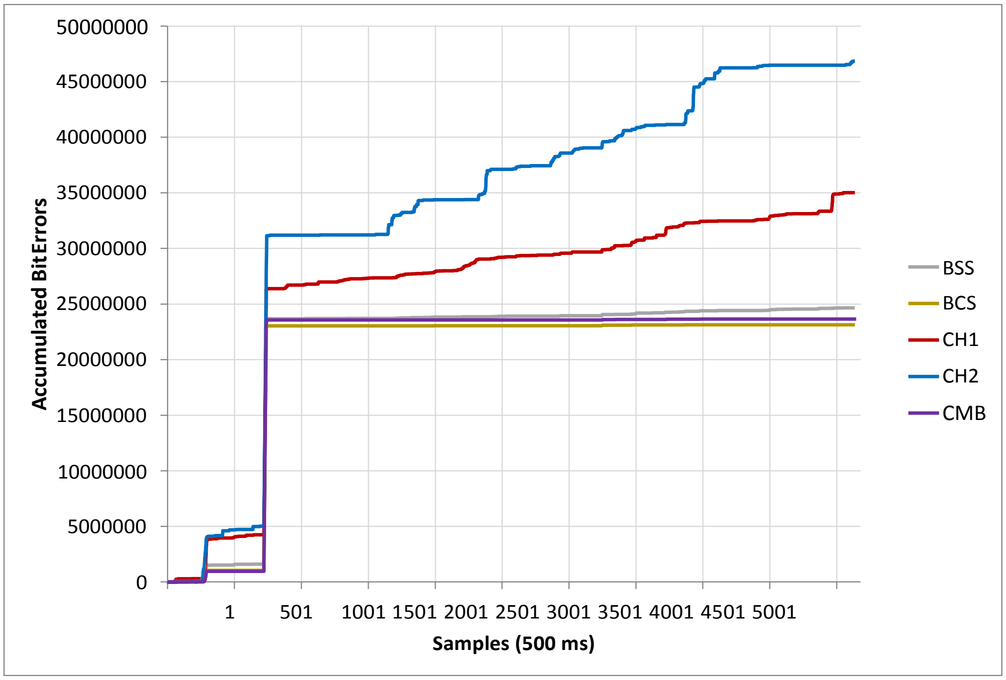

Figure 9 Flight test accumulated bit errors.

Figure 9 shows that the BCS again outperforms all channels, despite excellent performance from the Combiner.

Figure 9 also shows performance results for an independent BSS configured to use the same inputs as the BCS, that is, data and DQM from Channel 1, Channel 2, and the Combiner. It is worth noting that the BCS fared well in comparison with the BSS, while also noting that a single-receiver application is not typical for, and does not showcase the strengths of, a BSS.

Break Frequency Test This test implements the equivalent of the IRIG 118 break frequency test [6]. In this test, the Channel 1 and Channel 2 inputs experience periodic fades 180 degrees out of phase. As one or the other channel fades, the Combiner weights its output toward the opposing channel, avoiding the fade and maintaining a low bit error rate. The fade rate is increased to find the limit at which the Combiner can no longer keep up with estimating and implementing optimal weighting, at which point its bit error rate exceeds a predefined threshold. This fade rate is called the combiner break frequency.

The primary channel impairment in this test is AWGN. Therefore, the Pre-D MRC is optimal, and the BCS would ideally select only the Combiner output data. However, at the bottom of a fade, the opposing channel is essentially the only signal present in the Pre-D MRC output. Therefore, the DQM for the opposing channel and for the Combiner are essentially equal, and small variations in the DQM estimate may lead the BCS to occasionally select Channel 1 or Channel 2 rather than the Combiner. This should have no more impact than possibly degrading bit error rate by a tiny amount.

Consider, however, the case in which the fade is sufficiently deep to cause sync loss in the faded channel. If the BCS fails to properly correlate channels, and rapidly enough, then when the faded channel recovers, any selection other than the Combiner will lead to significant degradation. So, this test demonstrates the dynamic correlation performance of the BCS.

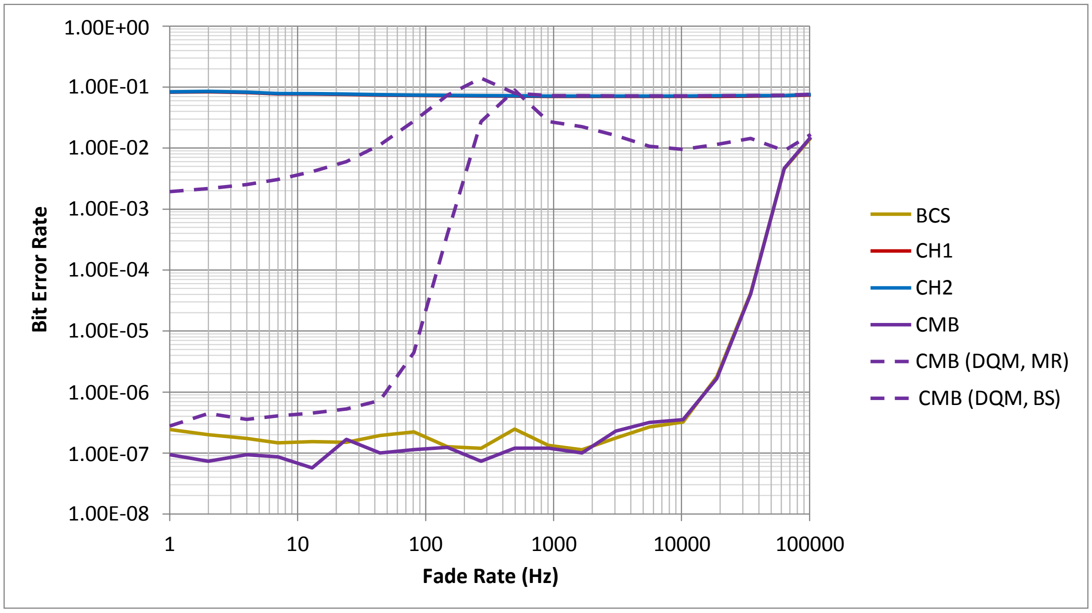

Figure 10 Break frequency test results.

Figure 10 shows results of the break frequency test for PCMFM at 5 Mb/s with a 20 dB fade depth – more than deep enough to induce sync loss. The BCS performs nearly ideally, showing no degradation due to correlation failure.

As an aside, note that Figure 10 also shows break frequency test results for two alternate combiner implementations. Given availability of an accurate DQM, it may be tempting to try replacing the Pre-D MRC with a combiner using DQM-based weighting. Clearly, the weighting could be tuned to emulate maximal ratio operation (DQM, MR in the figure), or it could be best source, based simply on the higher DQM (DQM, BS in the figure). This approach seems like a relatively easy way to achieve the objective of the BCS.

Unfortunately, ease of implementation comes at a price. As shown above, dynamic performance suffers dramatically with either of these implementations. This is because the DQM is calculated over a frame of received bits. By the time the Combiner receives the fed-back DQM, it is processing bits in a subsequent frame. Under static channel conditions, no harm is done. But under dynamic channel conditions, the guidance provided by the delayed DQM may be slightly misleading or even categorically incorrect.

A corollary to the preceding statement is that the BCS operates as well as it does in part because its selection criterion is applied directly to the bits the criterion is derived from. This may sound like an obvious statement of desired operation, but its importance cannot be overstated.

CONCLUSIONS

The well-known Pre-D MRC provides optimal performance under some channel conditions but not others. This paper has presented a means to mitigate those cases where Combiner performance falls short: a BCS that can dynamically select the best data from Channel 1, Channel 2, or the Combiner in a dual-channel receiver. The BCS uses DQM information to form its decision criterion, and it generates an output that contains accurate DQM information for the composite selected data stream. Thus, the BCS provides a single output from a dual-channel receiver that reliably supplies data superior to Combiner-only output.

The BCS does not replace a BSS for range-wide source selection from multiple distributed receivers. It has been demonstrated, however, to provide excellent performance, particularly under dynamic channel conditions. Because it is not constrained to correlate channels with very large delay skew, and it is aware that its input channels are likely correlated, the BCS is ideally suited for source selection at the receiver level.

ACKNOWLEDGEMENTS

The author would like to thank his talented and supportive teammates at Quasonix, many of whom made the BCS possible. In particular, Mark Geoghegan and Bob Schumacher planted the seed for the concept and steadfastly nurtured it until it could bear fruit. Gregg Wood – as usual – helped solve the thorniest implementation issues. And Greg Wells finished bringing the implementation to life. Most of all, Terry Hill contributed equal amounts of brilliance and patience, always expanding and never doubting what could be accomplished.

REFERENCES

[1] L. Kahn, "Ratio Squarer," in *Proc. IRE (Correspondence)*, Vol. 42, Nov. 1954.
[2] T. Hill, "Metrics and Test Procedures for Data Quality Estimation in the Aeronautical Telemetry Channel," in *Proc. Int. Telemetering Conf.*, Vol. 51, Oct. 2015.
[3] Range Commanders Council Telemetry Group, Range Commanders Council, White Sands Missile Range, New Mexico, *IRIG Standard 106-17: Telemetry Standards*, 2017.
[4] M. Geoghegan and R. Schumacher, "Performance Results Using Data Quality Encapsulation (DQE) and Best Source Selection (BSS) in Aeronautical Telemetry Environments," in *Proc. Int. Telemetering Conf.*, Vol. 53, Oct. 2017.
[5] M. Rice and E. Perrins, "Maximum Likelihood Detection from Multiple Bit Sources," in *Proc. Int. Telemetering Conf.*, Vol. 51, Oct. 2015.
[6] Range Commanders Council Telemetry Group, Range Commanders Council, White Sands Missile Range, New Mexico, *IRIG Standard 118-12: Test Methods for Telemetry Systems and Subsystems*, Vol. 2, 2012.