

Packet error rate and bit error rate non-deterministic relationship in optical network applications

L. B. James, A. W. Moore, A. Wonfor, R. Plumb, I. H. White, R. V. Penty
University of Cambridge, 15 JJ Thompson Avenue, Cambridge, CB3 0FD, United Kingdom
lbj20@cam.ac.uk

M. Glick, D. McAuley
Intel Research Cambridge, 15 JJ Thompson Avenue, Cambridge, CB3 0FD, United Kingdom

Abstract: The non-deterministic relationship between Bit Error Rate and Packet Error Rate is demonstrated for an optical media access layer in common use. We show that frequency components of coded, non-random data can cause this relationship.

©2005 Optical Society of America

OCIS codes: (060.2330) Fiber optics communications; (060.4250) Networks

This paper illustrates that, when operating at low receiver power, a commonly used (M,N) block-coding system, (8B/10B), causes a non-deterministic relationship between packet error rate and bit error rate (BER). Further, we show that at lower power, as is expected for systems operating in more complex and-or higher speed environments, a DFB laser has significant dependence related to the frequency of coded data. While a pseudo-random BER test may successfully achieve a desired error rate, repeated testing using real data and a common (M,N) block code results in frequency components that can cause a poorer error rate.

Optical Networking Context

We assert that the condition of low receiver power is increasingly likely as networks become more complex, with longer fibre lengths, optical switching systems and higher data rates. Ethernet in the first mile [1], along with a new generation of switched optical networks, are examples of this trend. Motivating our study is an investigation of Optical Packet Switching (OPS) constructed using a switched optical data path based upon semiconductor optical amplifiers (SOAs) [2]. In this work we observe that the data path between the sending and receiving end-systems consists of a significant numbers of devices such as SOAs, wavelength multiplex and de-multiplex units. The result is that the smaller power budget needed for higher data rates and designs with increasing numbers of optical components is forcing us towards what traditionally have been technical limits. In addition to restrictions on the power budget due to network complexity, we focus upon low power results because we assert that lower receiver power is a natural consequence of systems using higher data rates. While an increase in bit-rate requires a proportional increase in transmitter power, fibre nonlinearities impose limitations on the maximum optical power able to be used in an optical network.

We selected 8B/10B (M,N) block coding as the basis for our work [3]. This codec is widely used in many varied systems; it converts 8 bits of data for transmission (ideal for any octet-orientated system) into a 10 bit line code. We investigate Gigabit Ethernet on optical fibre (1000BASE-X [4]) under conditions where the received power is sufficiently low as to induce errors in the Ethernet frames. Following Jain [5], we limit frame size to less than 1512 octets where the Function Redundancy Check (FRC) within Ethernet is sufficiently strong to catch all errors.

Packet error rate versus BER

In past work we illustrated how bit errors are position independent but have a dependence upon the encoded data [6]. We found that the errors occur uniformly across any data packet, independent of packet size, and that there are no correlations evident between the positions of errors within the frame. We interpret this result as confirming that errors are highly localised within a frame and from this we are able to assume that the error-inducing events occur over small (bit-time) time scales. Further, we compared BER and packet error rate results, noting that frames containing different data contents lead to substantially different BER performance. Importantly, the relationship between the test data and BER results has little connection with the packet error rates for the same test data. This past work illustrated that the BER is not a good indicator of packet error, nor was packet error a useful indicator of BER. Our work presented here investigates why, for (M,N) block-codes such as 8B/10B, and the DFB laser, line-level measurements such as BER may not relate to packet error rate.

With commonly available test environments unable to evaluate BER, coding errors and Packet Error Rate, we used a combination of traditional BERT equipment (Agilent parts 70841B and 70842B) measuring a directly modulated 1548nm laser subjected to variable attenuation. The BERT was programmed with a series of bit sequences, each corresponding to a frame of Gigabit Ethernet data encoded as it would be for the line in 8B/10B. Purpose-built code is used to convert a frame of known data into the bit sequence suitable for the BERT. Against the BERT results, tests were conducted using a custom built environment (described in [6]) to evaluate both packet error and errors arising in the (M,N) coding layer.

Cause and Effect

Figure 1 illustrates the relationship between errors and the value of the each octet and with the preceding data octet, for packets carrying pseudo-random data. Figure 1(a) shows the error frequencies for the *current octet* X_i (the correct transmitted value of octets received in error) on the x-axis, versus the octet which was transmitted before each specific errored octet, X_{i-1} , on the y-axis. Figure 1(b) shows the preceding octet and the octet before that: X_{i-1} , X_{i-2} . Vertical lines in Figure 1(a) are indicative of an octet that is error-prone independently of the value of the previous octet. In contrast, horizontal bands indicate a correlation of errors with the value of the previous octet. It can be seen from Figure 1(b) that while there is a correlation between errors and the value in error or the immediately previous value, there is no apparent correlation with octets before this.

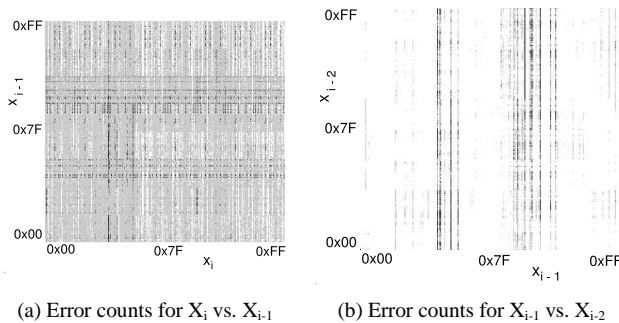


Fig. 1. Error counts for pseudo-random data octets




Fig. 2. Contrasting FFTs for low-error and high-error data octets

Consider the octets which are most subject to error, along with the 8B/10B codes used to represent them. In the pseudo-random packet data, the following ten octets give the highest error probabilities (independent of the preceding octet value): 0x43, 0x8A, 0x4A, 0xCA, 0x6A, 0x0A, 0x6F, 0xEA, 0x59, 0x2A. It can be seen that these commonly end in A, and this causes the first 5 bits of the code-group to be 01010. The octets not beginning with this sequence in general contain at least 4 alternating bits. Of the ten octets giving the lowest error probabilities (independent of previous octet), which are 0xAD, 0xED, 0x9D, 0xDD, 0x7D, 0x6D, 0xFD, 0x2D, 0x3D and 0x8D, the concluding D causes the code-groups to start with 0011. Fourier Transforms (FTs) were generated for data sequences consisting of repeated instances of the code-groups of 8B/10B, examples of which are shown in Figure 2. Examining the FTs of the code-groups for the high error octets, the peak corresponding to the base frequency (625MHz, half the line rate) is pronounced in most cases, although there is no such feature in the FTs of the code-groups of the low error octets. Illustrating this property, Figure 2 contrasts the Fourier Transforms for examples of low-error and high-error data octets.

The 8B/10B codec defines both data and control encodings, and these are represented on a 1024x1024 space in Figure 3(a), which shows valid combinations of the current code-group C_i and the preceding one C_{i-1} . The regions of valid and invalid code-groups are defined by the codec's use of 3B/4B and 5B/6B blocks [3]. In Figure 3(b) the octet errors found in pseudo-random data have been displayed on this code-space. It can be seen that errors tend to be clustered and that the clusters correspond to certain features of the code-groups. Two groups of clusters have been ringed; those that are indicated as $C_i=0011...$ represent those codes with a low-error suffix. In contrast the ringed values indicated as $C_i=010101...$ are the error-prone symbols with a suffix of 0xA.

For (M,N) block codes we observe that the likelihood of an error occurring in an octet depends not only on the value of that octet but the value of the preceding octet as well. Certain code-groups are more subject to error than others, and these code-groups are clustered together due to the nature of the coding system (Figure 3(b)). Such clustering leads to certain groups of octet values being more vulnerable to error once encoded. In addition, the nature of the coding scheme means that a single bit physical layer error can give rise to up to 4 bits of error at the decoded, octet layer [3].

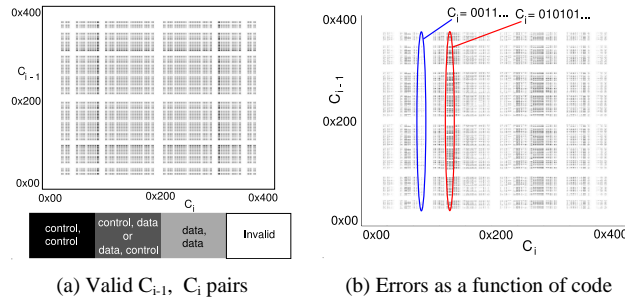


Fig. 3. Codebook for 8B/10B represented in 1024x1024 space




Fig. 4. Eye diagram: DFB laser at 1.25Gbps NRZ

We now interpret that result from the perspective of the physical device: electrically, semiconductor lasers are simple diodes, but the interaction between electron and photon populations within the device makes the modulation response complex. A first-order representation of the laser and driver may be obtained via a pair of rate equations, one each for electrons and photons, but DFB lasers at frequencies above 1 Gbit/s (e.g., 1.25 GHz for Gigabit Ethernet) need multiple coupled equations in order to account for spatial variations within the laser [7]. A significant range of behaviour is possible as bias, drive conditions, and physical structures vary. With ideal bias, just at threshold, some lasers have sufficient "memory" to react to the high frequency energy in 10101010 strings; resulting in a significant eye closure. Modelling, illustrated in Figure 4, confirms this result. The effect is small, but enough to increase the probability of error for such a data block. In addition, laser drive control loops, receiver timing loops, and the more sophisticated bandwidth limiting filters in receivers will, in principle, be disturbed slightly by particular bit sequences, and hence give increased error rates for those sequences.

Conclusions

The design of optical networks must consider the physical layer, its physical coding sub-layer and the combined impact upon higher level network protocols. We observe that the errors for an (M,N) block code in a low power regime are not uniform. We show that while a pseudo-random BER test may show low error rate, using real data and a (M,N) block code results in frequency components that cause non-deterministic error and a poorer overall result. This conclusion is contradictory to the assumptions of a significant body of coding and protocol work. We identify failures that through a combination of non-uniform data and error non-uniformity, lead to poor performance and potential undetected errors. This content-specific effect is particularly insidious because it occurs without a total failure of the network.

References

- [1] IEEE, "IEEE 802.3ah -- Ethernet in the First Mile," 2004, standard.
- [2] L. James, G. Roberts, M. Glick, D. McAuley, K. Williams, et al., "Wavelength Striped Semi-synchronous Optical Local Area Networks," in London Communications Symposium (LCS 2003), Sept. 2003.
- [3] A. X. Widmer and P. A. Franaszek, "A DC-Balanced, Partitioned-Block, 8B/10B Transmission Code," IBM Journal of Research and Development, vol. 27, no. 5, pp. 440-451, Sept. 1983.
- [4] IEEE, "IEEE 802.3z -- Gigabit Ethernet," 1998, standard.
- [5] R. Jain, "Error Characteristics of Fiber Distributed Data Interface (FDDI)," IEEE Transactions on Communications, vol. 38, no. 8, pp. 1244-1252, 1990.
- [6] L. B. James, A. W. Moore, and M. Glick, "Structured Errors in Optical Gigabit Ethernet," in Passive and Active Measurement Workshop (PAM 2004), Apr. 2004.
- [7] J. E. Carroll, J. Whiteaway, and R. Plumb, Distributed Feedback Semiconductor Lasers, ser. IEE Circuits, Devices & Systems Series. Co-published by the IEE and SPIE Press, 1998, no. 10.