Improvement in Bit-Error Rate for Optoelectronic Multicomputer Interconnection Networks Using Cyclic Redundancy Code Error Detection

Phil May, Jeffrey Cross, Abelardo Lopez-Lagunas, *Member, IEEE,* Brent Buchanan, D. Scott Wills, *Member, IEEE,* Nan Marie Jokerst, *Member, IEEE,* and Martin Brooke, *Member, IEEE*

***Abstract*—This letter presents testing results of an integrated optoelectronic (OE) channel employing *hop-by-hop* error control circuitry based on cyclic redundancy codes (CRC) to improve the effective bit-error rate (BER). The use of OE interconnect in place of wires in multicomputer networks becomes more attractive as channel bandwidth and power efficiency are increased. But these improvements must be accomplished while maintaining an acceptable channel BER. Test results of an integrated OE channel incorporating CRC-based error control circuitry demonstrate a BER reduction of two orders of magnitude while incurring a 20% bandwidth overhead. This may lead to higher bandwidth and higher efficiency OE interconnects.**

***Index Terms*—Error detection codes, multicomputer interconnection networks, optoelectronic channels, wormhole routing protocols.**

I. Introduction

OPTOELECTRONIC (OE) technology offers unique architectural alternatives to wire-based interconnect for applications such as multicomputer networks. However, OE-based interconnect must be competitive with respect to density, channel bandwidth, bit-error rate (BER), and power efficiency. Interrelationships between these metrics, especially bandwidth, BER, and power, complicate efforts to improve overall channel performance. This letter presents experimental results of an integrated error control approach to help improve the effectiveness of OE-based interconnect for multicomputer network applications.

Research being conducted at Georgia Tech [5] addresses the development of a scalable interconnection network incorporating through-wafer optoelectronic channels. This network employs integrated error control circuitry based on cyclic redundancy codes (CRC) [7]. CRC's have been extensively used in long haul telecommunication applications [8], [9] to validate data integrity when the possibility of a transmission error exists. We propose a new technique for applying CRC's to wormhole-routed multicomputer interconnection networks [3], [6]. Rather than providing *end-to-end* data validation on a low-error channel, we are employing CRC's in *hop-by-hop* data validation to provide an acceptable effective BER using a lower power, higher bandwidth OE channel. This is accomplished by extending the existing ACK/NACK protocol commonly used in wormhole routing to retransmit data found to be in error by the CRC circuitry.

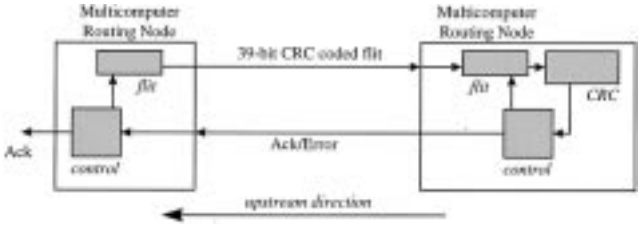

Fig. 1. Wormhole routing flit transfer protocol can be used to initiate flit retransmission in the presence of errors.

This letter provides test results for an optoelectronic data link using CRC error control circuitry integrated with the optoelectronics. We experimentally demonstrate a significant BER improvement by integrating CRC error control in the channel. This improvement allows the delivered BER of the channel to be much lower than the physical BER (raw error rate of the channel). Using integrated CRC error control, the optical channel can be operated at a lower power or higher link speeds, which can potentially offset the overhead of the CRC bits transmitted with the data.

II. Flit-Based Error Control and Wormhole Routing

A data transfer protocol used in wormhole routing can support flit-based error detection and retransmission, as shown in Fig. 1. In wormhole routing, a *flit* is the smallest data unit over which flow control is performed. In a multicomputer interconnection network, a flit moves from one routing node to the next on its way to the destination, signaling to the flits behind it that the way is clear using an acknowledgment protocol. This requirement for an acknowledgment signal can be expanded to support error control over the link by incorporating an error signal that requests retransmission of a flit. If, after a flit has been transmitted across the link, the Ack/Error signal indicates success, the transmitting node (leftmost node in Fig. 1) signals to the node upstream from itself to transmit the next flit. If, on the other hand, an error is detected, the Ack/Error signal indicates an error and the transmitting node retransmits the same flit and forgoes

Manuscript received December 9, 1996; revised February 13, 1997. This work was supported by the National Science Foundation under Contract ECS-9422452, Contract ECS-9402723, and Contract ECS-9058144, by the members of the Georgia Tech Manufacturing Research Consortium, by the Ford Motor Company, by Army MICOM under Contract DAAH01-92-D-R005, and by DuPont through a donation of polyimide PI2611.

The authors are with the School of Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0250 USA..

Publisher Item Identifier S 1041-1135(97)04052-4.

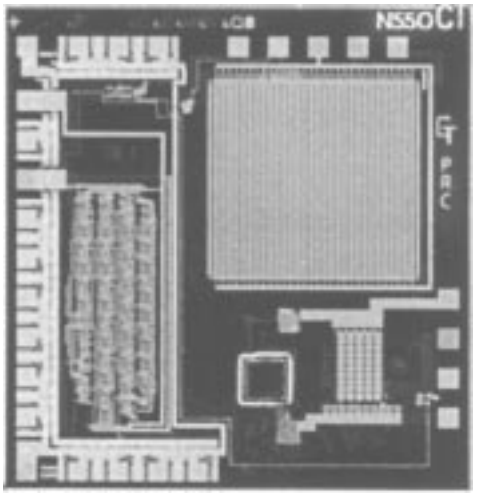

(a)

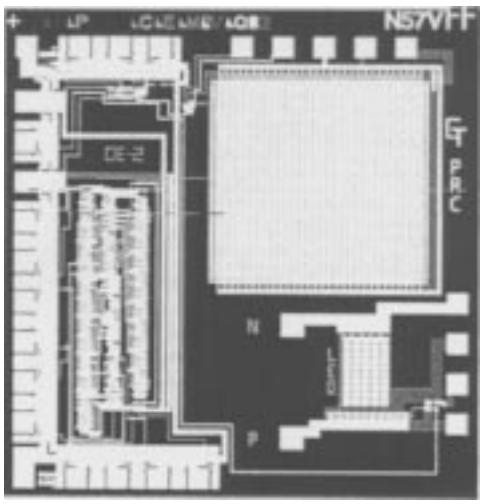

(b)

Fig. 2. Photomicrographs of the integrated transmit chip and the receive chip used in CRC BER testing (CRC circuitry in lower left).

signaling the upstream node. This process is repeated until the flit is successfully transmitted. In this approach, only errors that are so severe that the CRC circuitry can not detect them are propagated through the network. The particular CRC that is used in this design is guaranteed to detect all one, two, and three bit errors in the flit, and any burst (sequential) errors up to seven bits in length [6].

III. Test Description

To evaluate the integrated OE channel and error control circuitry, a transmitter and receiver chip were fabricated. These chips differ in the function of their CRC error control circuitry (encoding versus decoding). Their circuitry includes the CRC encoder/decoder along with an silicon emitter driver and detector amplifier. A bipolar junction transistor (BJT) Si detector and thin film emitter integration site are also included. The chips were fabricated in 2.0-μm SCMOS through the MOSIS foundry service. After chip fabrication, a thin-film emitter was integrated onto the unpackaged transmitter die. Then each chip was bonded into a 40-p-i-n ceramic DIP package and tested. Fig. 2 shows photomicrographs of the

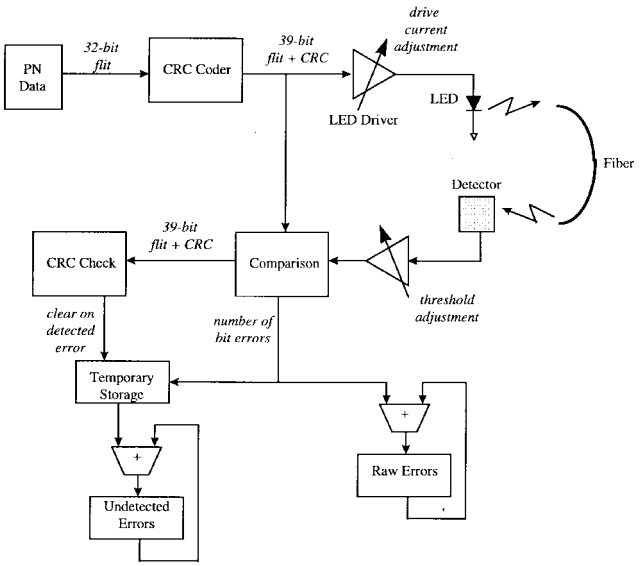

Fig. 3. The test setup measures both raw BER and effective BER simultaneously by using CRC results to convert the raw number of errors to an effective number of errors.

unbonded chips, including the integrated thin film emitter on the transmitter.

For this system, a p-i-n $Al_{0.3}Ga_{0.7}As$–GaAs–$Al_{0.3}Ga_{0.7}As$ double heterostructure light-emitting diode (LED) was fabricated for integration with the CRC circuit. An AlAs layer was used for the etch stop layer or sacrificial layer enabling the separation of the epitaxial device layers from the growth substrate [1], [10]. The bottom contact for the LED on the CRC circuit was metallized with Ti–Au (500–1500 Å), and the thin-film LED was aligned and bonded to this pad. DuPont 2611 polyimide was used to electrically isolate the top and bottom contacts, and windows in this polyimide for the top contact were subsequently opened using reactive ion etching. Finally, the top contact of AuGe–Ag–Au (800–800–1800 Å) was deposited onto the LED/circuit [2].

The CRC receiver circuit uses an array of silicon phototransistors as detectors with the base current controlled by the input optical signal [2]. The individual phototransistors are composed of a small n-type central emitter diffusion, a large area p-type base diffusion, and a ring electrode n-type collector. Due to the CMOS circuit process, only small junction depths are available that limits the absorption depth. In addition, the absorption coefficient for silicon is low compared to GaAs, and as a result, the responsivity of the Si detectors is small. Thus, the use of a BJT with the photocurrent injected into the base produces a gain in the photodetector that compensates for the low responsivity.

A functional block diagram of the test setup is shown in Fig. 3. The testing involved the transmission of CRC coded, 32-b data flits from the transmitter chip through a fiber link to the receiver chip.[1] The flits were also sent to the receiver via a wire link, and the two flits were compared to identify the number of bit errors introduced in the optoelectronic link. This number was added to a running total, and then converted into

[1] Although the target application for the CRC error control circuitry is a through-wafer link, a fiber link was used to characterize the error control capabilities, simplifying the test setup.

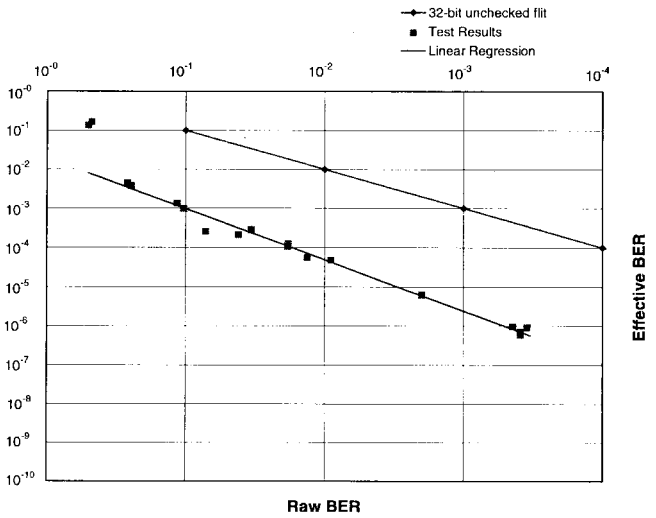

Fig. 4. Test results indicate a greater than two order of magnitude BER improvement using CRC codes with flit retransmission. Error bars are negligible on this plot.

a raw BER number for the channel (measured errors divided by elapsed time). The receive CRC circuitry also checked the received flit for errors. Since the error detection capability of the CRC is limited, not all possible error combinations are detectable. If the CRC circuitry was unable to detect the errors in the data, the number of errors identified in the wire-fiber comparison was added to a second running total which was then converted into an *effective* BER number. If, however, the CRC detected the errors, the number of errors for that flit were ignored by the second running total (effective BER number). This models the effect of retransmitting flits found to be in error by the CRC circuitry. Errors detected by the CRC circuitry can be eliminated (by retransmission) from the channel. The degradation of channel bandwidth due to retransmission is negligible for raw link BER values less than 10^{-4} [6]. The overhead of transmitting the CRC data accounts for a 20% loss of available channel bandwidth. Errors not detected by the CRC circuitry are propagated to the destination, contributing to the effective channel BER. The data used to generate the flits was generated using a p-n sequence generator of the type described in [4].

To introduce a controlled number of errors on the optoelectronic link, the LED bias current was adjusted to vary the intensity of the light produced. As part of this procedure, the threshold of the detector was also adjusted to produce the best combination of emitter and detector biasing. This setting resulted in the lowest raw BER number given the intensity of the transmitted light. The adjustment points were chosen to produce a varied distribution of data points, and multiple data points were taken at each adjustment point over several days to verify the repeatability of the experiment.

IV. Test Results

The results of this testing are presented graphically in Fig. 4. This graph shows, on a log–log scale, the measured raw and error controlled BER's. There are two important features of the data in Fig. 4 and the lines fitted to these points. First, the corrected BER data points are a *minimum* of two orders of magnitude lower than the line indicating the raw BER of the link. Second, the line fit to the data points has a greater (magnitude) slope than the raw BER line which may extrapolate to further improvements in BER for lower raw BER links. The CRC link described herein should achieve the accepted industrial standard BER of less than 10^{-12} running at an uncorrected BER of 10^{-8}, (an extension of the current results), however, the time required to perform such experiments at our current link operating rate is prohibitively long. Future work will explore faster links with tests being performed at lower BER's. The two points above the line fit to the data indicate that the curve is swinging toward the point (1, 1) where the raw and effective BER are equal.

V. Conclusion

The results of the testing described in this letter indicate that a significant improvement in BER (two orders of magnitude) is achievable at modest cost (a 20% reduction in link bandwidth) by the use of integrated CRC error control circuitry and flit level retransmission in multicomputer interconnection networks. This improvement in effective BER allows the use of higher raw BER OE channels while still satisfying data integrity requirements. This allows optoelectronic links to be operated with greater channel bandwidth and power efficiency, making optoelectronic interconnect more attractive to system designers considering OE technology alternatives.

Acknowledgment

Microfabrication was performed in the Microelectronics Research Center (MiRC), Georgia Institute of Technology, Atlanta.

References

[1] C. Camperi-Ginestet, M. Hargis, N. M. Jokerst, and M. Allen, "Alignable epitaxial liftoff of GaAs materials with selective deposition using polyimide diaphragms," *IEEE Photon. Technol. Lett.*, vol. 3, pp. 1123–1126, 1991.

[2] J. Cross, A. Lopez-Lagunas, B. Buchanan, L. Carastro, S.-C. Wang, N. M. Jokerst, S. Wills, M. Brooke, and M. A. Ingram, "A single fiber bidirectional optical link using colocated emitters and detectors," *IEEE Photon. Technol. Lett.*, vol. 8, no. 10, pp. 1385–1387, Oct. 1996.

[3] W. J. Dally, "Performance analysis of k-ary n-cube interconnection networks," *IEEE Trans. Comput.*, vol. 39, pp. 775–785, June 1990.

[4] P. D. Hortensius, R. D. McLeod, and H. C. Card, "Parallel random number generation for VLSI systems using cellular automata," *IEEE Trans. Comput.*, vol. 38, pp. 1466–1472, Oct. 1989.

[5] W. S. Lacy, J. L. Cruz-Rivera, and D. S. Wills, "The offset cube: A three-dimensional multicomputer network topology using through-wafer optics," *IEEE Trans. Parallel Distrib. Syst.*, submitted for publication.

[6] P. May, M. Lee, S. T. Wilkinson, O. Vendier, Z. Ho, S. W. Bond, D. S. Wills, M. Brooke, N. M. Jokerst, and A. Brown, "A 100 Mbps, LED through-wafer optoelectronic link for multicomputer interconnection networks," *J. Parallel Distrib. Comput.*, special issue on parallel computing with optical interconnects, to be published.

[7] E. Prange, "Cyclic error-correcting code in two symbols," Air Force Cambridge Res. Center-TN-57-103 Rep., Cambridge, MA, Sept. 1957.

[8] A. S. Tanenbaum, *Computer Networks.* Englewood Cliffs, NJ: Prentice-Hall, 1981.

[9] S. B. Wicker, *Error Control Systems for Digital Communication and Storage.* Englewood Cliffs, NJ: Prentice-Hall, 1995.

[10] E. Yablonovitch, T. Gmitter, J. P. Harbison, and R. Bhat, "Extreme selectivity in the lift-off of epitaxial GaAs films," *Appl. Phys. Lett.*, vol. 51, pp. 2222–2224, 1987.