A Chipset for Gigabit Rate Data Communication

Richard C. Walker, Tom Hornak, Chu-Sun Yen
Hewlett-Packard Labs, 1501 Page Mill Road, Palo Alto, CA 94303

Kent Springer
Digital Equipment Corporation

Abstract

A gigabit rate data link consisting of four custom silicon bipolar chips for transmitting parallel data between elements of a distributed computer system has been developed. A transmission rate of 16 bits in parallel at 50MHz, or with encoding overhead, a serial rate of 1Gbit/s is demonstrated. The link utilizes a new encoding scheme that is bandwidth efficient. Unlike other links, the phase/frequency locked loop also provides frame synchronization and requires no trimming for data retiming, either in production or later.

Introduction

Telecom applications have pioneered the use of fiber equipment in voice transmission, but use of fiber media for high speed computer communication has been limited. Computer data transmission requires special hardware to provide serialization and deserialization of the data to be compatible with the serial nature of the fiber media. Additional constraints on the data stream for a fiber optic system include DC balance (to facilitate AC-coupled laser biasing and receiver design) and some scheme for recovering a clock from the data stream. To date, there has been a lack of high speed interface chips to provide these functions.

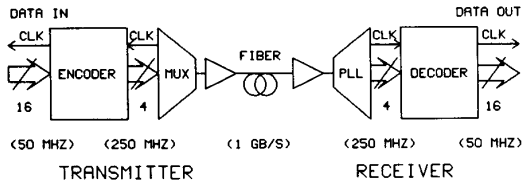

Fig. 1. Block Diagram of one half of a duplex link.

We have designed a set of four ICs using silicon bipolar technology for a gigabit rate data link: an Encoder chip, a Multiplexer chip, a PLL/Demultiplexer chip and a Decoder chip. These chips provide the necessary functionality to form a high bandwidth fiber optic data link for point to point communication (Fig. 1).

Line Code

Coding schemes that satisfy the needs of clock recovery and DC balance are a tradeoff between coder complexity and bandwidth utilization. Simple Manchester coders suffer a bandwidth penalty of 1/2: two symbols are sent for each received bit. Other codes such as 5b/6b [1] and 4b/5b used in FDDI [2] are more bandwidth efficient, but are complex to implement, usually requiring a ROM. By using an efficient code, more of the raw link bandwidth is available. Our code is only moderately complex to encode, simple to decode, and has the same efficiency as the 4b/5b code.

The new code used in this link achieves DC balance by transmitting the 16 bit data words either inverted or non-inverted (Fig. 2) in such a way that the accumulated offset of the transmitted words is bounded. Only one additional bit per data word is required to indicate the polarity of transmission to the receiver.

To allow the transmission and recognition of non-data words, another indicator bit is needed. Fill words are sent on the link when no data is available, and are discarded by the decoder. Control words, the other type of non-data code, are available for maintenance of the link.

Finally, two more bits are added which are always "01". This "0" to "1" edge, called the *Master Transition*, is used by the phase locked loop as the phase reference for both bit and frame clock recovery. The frame clock is used by the demultiplexer for frame alignment. Thus a total of four extra bits per word are required in this coding scheme for any length of data word. Fig. 2 shows the coding format for the 16 bit data word.

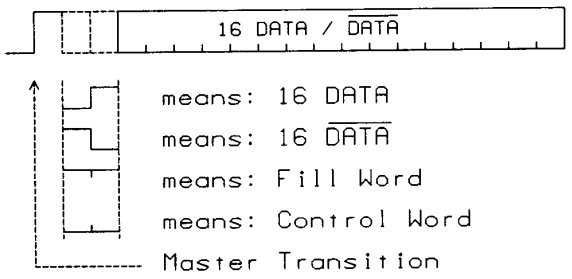

Fig. 2 Format of 16b/20b Line Code

The DC balance performance of this code can be judged by the baseline wander with two 100KHz poles due to AC-coupling in the link, for example: one at the transmitter and one at the receiver. Fig. 3 shows a simulated histogram for 100K bits of random data, coded, at 1Gbit/s. For these conditions the baseline wander at the receiver slicing circuit is less than ±1% of the eye opening. The maximum accumulated offset in the serial data is guaranteed to be less than 24 bits.

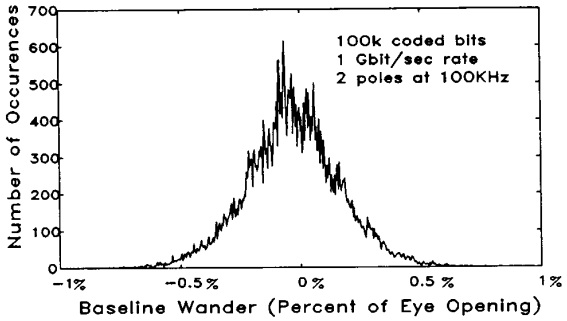

Fig. 3 Histogram of baseline wander at receiver slicing circuit.

CH2771-4/89/0000(0288)$1.00 © 1989 IEEE

Encoder Chip

A simplified block diagram of the encoder is shown in Fig. 4. As each word enters the encoder, the sign of its offset is compared with the sign of accumulated offset, and is conditionally inverted if that will tend to restore the balance. Inversion takes place if the two signs are equal. The offset of the encoded word is then accumulated before being transmitted.

In addition to the sign circuit, there are sections for clock division, fill/control word generation, and serialization. The 4 bit data path is pipelined to provide minimum delay between stages, operating with a 2.5ns minimum cycle time. Input and output latches are used to provide maximum timing margin for external connection of the chips.

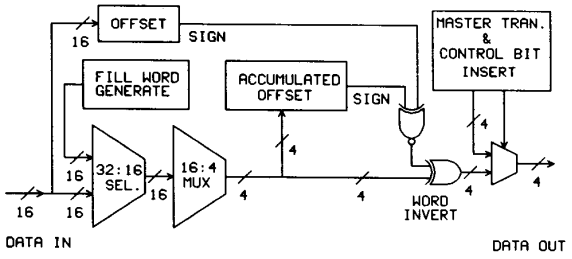

Fig. 4 Encoder Block Diagram.

Multiplexer Chip

A 4:1 time division multiplexer serializes the output of the Encoder chip at up to a 2Gbit/s rate. To maximize the serial data rate, the multiplexer clocks the the output on both edges of the clock. A 1 Gbit/s output data rate therefore is achieved with only a 500MHz clock. The MUX chip also provides a divided down clock for the Encoder chip. Fig. 5 shows the transmitted frame structure at 1Gbit/s.

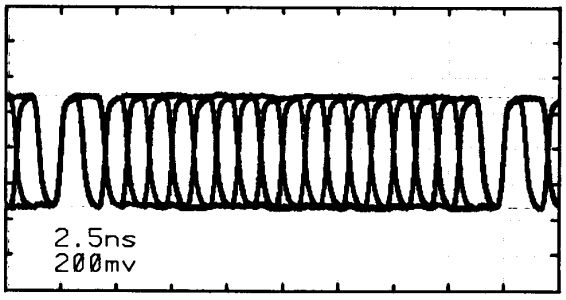

Fig. 5 MUX output eye diagram, triggered at the frame rate, showing transmitted frame and *Master Transition* at 1Gbit/s.

Phase Locked Loop / Demultiplexer Chip

Initial synchronization between the transmitter and the receiver is achieved by sending a special training fill word containing only a single rising edge per frame, which is the *Master Transition*. The phase locked loop in the PLL chip uses these edges to lock the receiver voltage controlled oscillator (VCO) to the transmitter clock. A combined frequency detector and phase detector is used, as shown in Fig. 6.

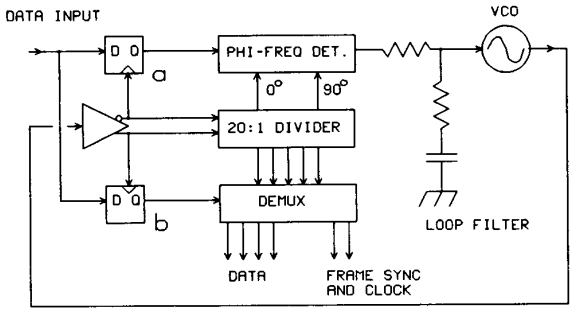

Fig. 6 PLL Block Diagram.

A common element of the frequency and phase detectors is a D-latch (labeled "a") that is clocked by the VCO clock. The frequency detector operates by sampling the output of the D-latch at two points 90° apart (at the frame rate). The sampling is controlled by a frame clock, obtained by dividing the VCO clock by 20. When the input is the special training fill word, the signals at these two points determine the direction that the VCO frequency should be adjusted. The addition of a frequency detector allows the use of a low-cost, non-precision, on-chip VCO with potentially hundreds of MHz center frequency uncertainty.

The frequency detector is first used to bring the VCO frequency to within the pull-in range of the loop. After that, the frequency detector is disabled, and the phase detector output is used to control the VCO. Samples of the D-latch output taken at the frame rate provide a binary quantized representation of the phase error.

Once phase-locked, the D-Latch samples the data stream in the vicinity of the *Master Transition*. If the VCO drifts slightly higher in frequency, the latch will start sampling prior to the *Master Transition* and will begin seeing all 0's. A slightly lower VCO frequency, on the other hand, will cause the latch to sample the high portion of the *Master Transition* thereby outputting all 1's.

The VCO alternates between two frequencies, bracketing the input frequency with a duty cycle such that the average frequencies of the VCO and data are equal. Without a loop filter, this leads to a high jitter in the recovered clock. By adding a appropriate lag-lead filter to the loop, an acceptable amount of tracking jitter can be achieved without limiting the frequency detector capture range. Fig. 7 shows the measured jitter of the recovered clock, showing a hunting jitter of less than 3.5 ps RMS.

Once the loop is phase locked, data transmission may be be safely enabled because the phase detector samples the input only in the vicinity of the *Master Transition*. Since transmission must be continuous to maintain phase lock at the receiver, non-data codes are sent on the link when no data is available, or when desired for the maintenance of the link.

Data retiming is achieved on the same chip with a second D-latch (labeled "b") clocked with the complement of the VCO clock. Because the clock-to-Q delay of the retiming D-latch is matched to that of the phase detector D-latch, the circuit is guaranteed to sample the data in the middle of the bit cell once the loop is locked, with no adjustments required.

Data from the retiming flip-flop is demultiplexed 1:4 to be sent to the decoder chip. Along with the data, a phase adjustable nibble clock is output to allow for chip to chip delays. To provide for frame synchronization, the PLL chip outputs a sync pulse every 5 data nibbles. This pulse occurs coincident with the control nibble.

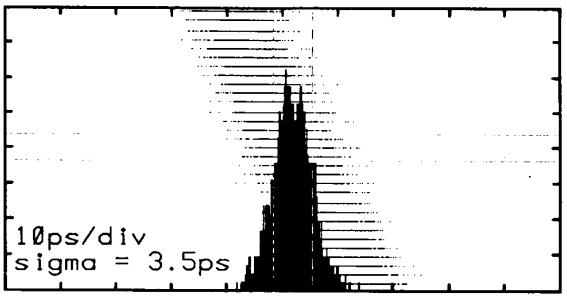

Fig. 7 Phase jitter histogram of the recovered clock in loop-back mode at 1Gbit/s, showing that the loop hunting jitter is less than 3.5ps RMS.

The PLL chip is configured to utilize either an on-chip or an off-chip VCO. An on-chip voltage controlled ring-oscillator similar to that published [3], with a tuning range of 0.9 to 1.5GHz is available for use with the PLL. This circuit will replace the off-chip VCO in future experiments.

Decoder Chip

The decoder provides the companion functions of deserialization and decoding. A block diagram is shown in Fig. 8. No processing is done on the frame until it is in full parallel form. A framing pulse generated in the phase locked loop synchronizes the clock divider phase to the incoming frame and correctly aligns the demultiplexer. If the frame contains data, the added coding bits are used to enable or disable the word inverter. If the word is a fill or control code, it is identified and decoded. Like the encoder, all data inputs and outputs are latched.

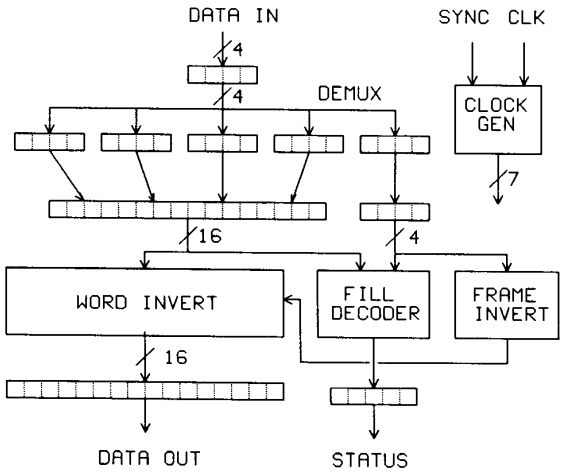

Fig. 8 Decoder Block Diagram.

Link Startup and Maintenance

Some of the non-data codes are used during the lock up interval. An end to end handshake ensures that both ends of a full duplex link are frequency and phase locked before data is transmitted.

The consistent presence of the two *Master Transition* bits is monitored by the decoder chip to detect a locked condition. If loss of lock is detected, then the loop startup procedure is automatically re-initiated.

Implementation

An overview of the four chips described is given in Table 1. The Encoder and Decoder chips are fabricated in a 5GHz f_t bipolar process, using Emitter Function Logic [4] standard cells. They contain 2500 and 2000 active devices, respectively. Packaged parts run up to 400MHz, capable of supporting a serialized line rate of 1.6Gbit/s. The MUX and PLL chips are implemented in a 10GHz f_t bipolar process using full custom ECL design. They contain 350 and 950 active devices, respectively. A PC Board containing all four chips has been developed and link speeds of 1Gbit/s have been achieved.

chip	function	size (mm)	power
Encoder	Encoding 16:4 multiplex	4x5	3W
MUX	4:1 multiplex	1.3x2	1W
PLL	Clock Recovery 1:4 demultiplex	2.3x2.6	1.5W
Decoder	Decoding 4:16 demultiplex	4x4	2W

Table 1. Chipset overview

Acknowledgments

We thank the Hewlett-Packard Santa Clara Tech Center for fabrication of the circuits described. We also thank Rasmus Nordby, Doug Crandall, Craig Corsetto, Steve Hessel and Dave DiPietro for their contributions in the early phases of the project.

References

[1] J. Tani, D. Crandall, J. Corcoran, T. Hornak,"Parallel Interface IC's for 120Mb/s Fiber Optic Links", *ISSCC Digest of Technical Papers*, Vol. 30, p. 190,191, February 1987.

[2] "FDDI Physical Layer Protocol (PHY)", Draft Proposed American National Standard, X3T9/85-39, Fig. 5-2., May 10, 1985.

[3] K.E. Syed and A.A. Abidi,"Gigahertz Voltage-Controlled Ring Oscillator", *Electronics Letters*, Vol. 22, No. 12, p. 677-679, 5th June 1986.

[4] Z.E. Skokan,"EFL Logic Family for LSI", *ISSCC Digest of Technical Papers*, Vol. 16, p. 162-163, February 1973.