HEWLETT PACKARD

A 1.5 Gb/s Link Interface Chipset For Computer Data Transmission

Richard C. Walker, Tom Hornak, Chu-Sun Yen
Joey Doernberg, Kent Springer
Instruments and Photonics Laboratory
HPL-90-105
July, 1990

data communication; computer links; fiber optic link; interface; line coding; clock extraction; Gbps LAN

A set of four ICs provide encoding, multiplexing, clock extraction/demultiplexing and decoding for gigabit-rate serial data transmission. These chips form a high bandwidth data link for point-to-point communication. A new line code is implemented that provides DC balance, efficient encoding, framing, and simple clock extraction. Embedded in the code is a fixed transition used by the phase/frequency locked loop (PLL) for simple clock extraction and frame synchronization. Unlike other links, our PLL requires no trimming for data retiming, either in production or later. An on-chip voltage-controlled oscillator (VCO) with a tuning range of 1.1 to 1.6 GHz is available for use with the PLL. With this chip set, we have demonstrated a transmission rate of 16 bits in parallel at 75 MHz, or with encoding overhead, a serial rate of 1.5 Gbit/s.

Kent Springer is presently with Digital Equipment Corporation.

A 1.5 Gb/s Link Interface Chipset For Computer Data Transmission

Richard C. Walker, Tom Hornak, Chu-Sun Yen, Joey Doernberg, Kent Springer[*]

Hewlett-Packard Labs, 1501 Page Mill Road, Palo Alto, CA 94303

1. Introduction

Telecom applications have pioneered the use of fiber equipment in voice transmission. Because of the high cost of fiber optic transmitters and receivers and problems with skew between multiple channels over long distances, serial transmission is generally used for telecom data transmission [1]. However, parallel data interfaces are preferred for convenient connection to computer equipment. The use of fiber media for high speed computer communication has been limited by the lack of Link Interface chips required to adapt the parallel format of the computer data into the serial bit streams required for fiber transmission. The Link Interface chipset must guarantee DC balance of the serialized bit stream (to simplify laser biasing and facilitate receiver design), and provide some scheme for recovering a clock from the data stream. Although there have been several GaAs chipsets published [2, 3], to date, there has been a lack of high speed silicon bipolar interface chips to provide these functions. Silicon bipolar process is a preferred solution because of its long established reliability records.

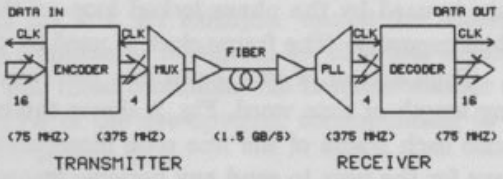

Fig. 1. Block Diagram of one half of a duplex link.

[*]Presently with Digital Equipment Corporation.

We have designed a set of four ICs using silicon bipolar technology for a gigabit rate data link: an Encoder chip, a Multiplexer (MUX) chip, a PLL chip and a Decoder chip [4]. These chips provide the necessary interface functionality to form a high bandwidth fiber optic data link for point-to-point communication (Fig. 1). The architecture of the link is largely determined by the line code design which is discussed next.

2. Line Code, Clock and Frame Synchronization

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 [5] and 4b/5b used in FDDI [6] are more bandwidth efficient, but are complex to implement, usually requiring a ROM. By using an efficient code, less raw link bandwidth is required. Our code is only moderately complex to encode, simple to decode, and can be highly efficient [7].

The new code used in this link achieves DC balance by transmitting each M-bit data word either inverted or non-inverted in such a way that the accumulated offset of the transmitted words is bounded. An extra coding bit is prepended to each M-bit word to indicate the polarity of the word to the receiver.

To allow the transmission and recognition of non-data words, such as Fill and Control 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 are available for out-of-band signalling.

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 PLL 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 a 16 bit data word. Because each frame of the line code incorporates a reference transition, it is not necessary for the user to send any periodic frame-sync words, as is the case with 4b/5b and 8b/10b codes. This allows the link to be conveniently used in a synchronous environment where the insertion of extra frame-sync words is impossible.

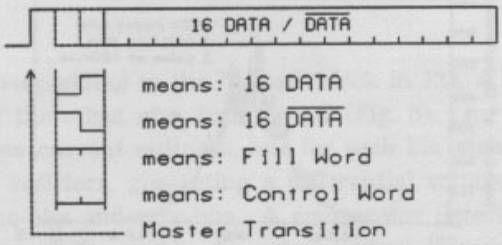

Fig. 2 Format of Line Code for 16 Bit Data Word.

Long term DC balance of the transmitted data permits the regulation of laser bias current by simply maintaining a fixed average optical power. In the receiver a balanced code permits AC-coupling without eye-degrading baseline wander and jitter.

The DC balance performance of a given code can be characterized by the worst case longterm disparity of ones and zeroes sent on the link over time. For a code, with guaranteed maximum disparity of n bits, transmitted at a data rate R through a channel with k low-pass poles, each with a time constant τ, the percentage peak to peak baseline wander at the receiver slicing circuit will be approximately

$$\approx \frac{kn}{2R\tau} \times 100\%.$$

At low data rates R, complex codes with very low disparities n are required to control the baseline wander. At higher data rates, for a given low-pass time constant, longer disparities n are tolerable. This allows a trade-off of coding complexity and disparity for Gbit/sec data links.

The DC balance performance of our code can be judged by the baseline wander with two 100 kHz poles 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 with M=16, at 1 Gbit/s. For these conditions the baseline wander at the receiver slicing circuit is less than $\pm 1\%$ of the eye opening, corresponding to a maximum accumulated disparity of 24 serial data bits.

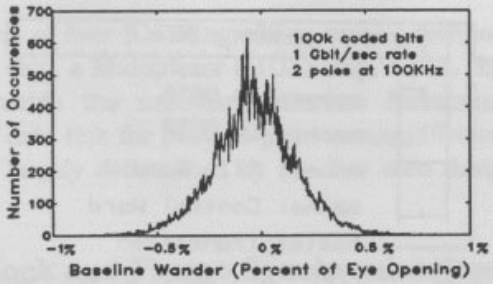

Fig. 3 Histogram of baseline wander at receiver slicing circuit, M=16.

3. Encoder Algorithm

A simplified block diagram of the Encoder is shown in Fig. 4. As each word enters the Encoder, the sign of its offset is detected and compared with the sign of accumulated offset in an XNOR gate. Each nibble of the word is then inverted if the signs are the same, and are not inverted if the signs are different. 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 2.5 ns minimum cycle time. Input and output latches are used to provide maximum timing margin for external connection of the chips.

The hardware for encoding of the data is simple, consisting of exclusive-or gates.

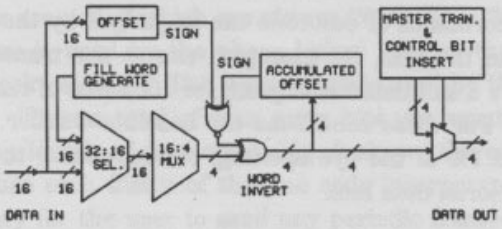

Fig. 4 Encoder Block Diagram.

3.1. Sign Circuit

An analog circuit, corresponding to the "offset" block in Fig. 4, was implemented to determine the sign of the offset of a 16 bit word (Fig. 5), i.e., whether it has more ones or zeroes. Sixteen current switches, one for each bit, steer a unit current into one of two summing resistors, generating a differential voltage proportional to the difference between one-bits and zero-bits. A comparator detects the polarity of the voltage, giving the sign of the word. An extra current source shifts the threshold from the balanced condition so that a word with 8 ones and 8 zeroes will produce a known sign.

This embedding of an analog circuit in an all-digital chip gives a faster, lower power implementation. The requirement for matching of current sources is tighter than for the logic gates but is still within the capabilities of a conventional digital bipolar process.

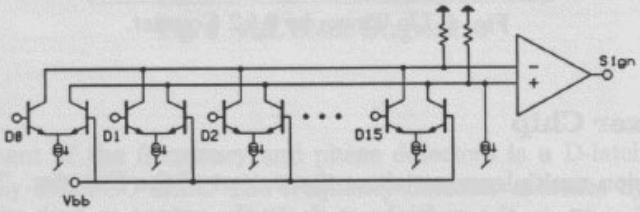

Fig. 5 Sign Circuit.

3.2. Accumulator

In a simple minded approach to the accumulator, the serialized bit stream could be applied to an up/down counter: counting up for ones and down for zeroes. The MSB of the counter would give the sign of the accumulated offset. Since the feedback forces the offset toward zero, the required size of the counter is bounded: it is guaranteed not to overflow. Unfortunately, running a counter at the full serial speed would be difficult and wasteful of power. Also, the serialized data is not available on Encoder chip, since final 4:1 multiplexing is done on the MUX chip built with a higher speed process.

Another approach would be to accumulate the offsets at the 16 bit parallel level, processing a whole frame at a time. The hardware required is formidable, including a 6 bit full adder and a number of smaller adders. The complexity and power of this approach was also prohibitive.

The accumulator was instead implemented with a lookup table and an up/down counter at the 4 bit level (Fig. 6). When the data is examined in 4 bit nibbles, the contribution to the offset can only be ± 4, ± 2 or 0. A special counter that can be incremented or decremented by one or two satisfies this function. It must operate at 1/4 of the serial rate, or five times the parallel rate, nominally 375 MHz. For each 4 bit nibble, the data addresses a 16 by 5 lookup table that generates the control signals for the counter. A data nibble of 0001, for example, would activate the -1 control line and decrement the counter by one.

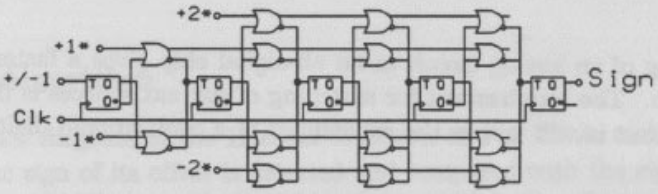

Fig. 6 Up/Down by 0,1,2 Counter.

4. Multiplexer Chip

A 4:1 time division multiplexer serializes the output of the Encoder. To minimize the required clock frequency, the multiplexer clocks the output on both edges of the clock. A 1.5 Gbit/s output data rate therefore is achieved with only a 750 MHz clock. The MUX chip also provides a divided down clock for the Encoder chip. Fig. 7 shows the measured transmitted frame structure at 1.5 Gbit/s.

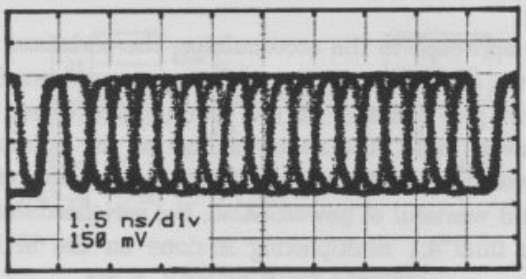

Fig. 7 MUX output eye diagram, triggered at the frame rate, showing transmitted frame and *Master Transition* at 1.5 Gbit/s.

5. 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 VCO to the transmitter clock. A combined frequency detector and phase detector is used [8], as shown in Fig. 8.

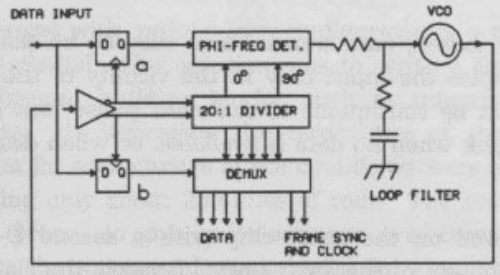

Fig. 8 PLL Block Diagram.

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 phase detector output is used to control the VCO and the the frequency detector is disabled. Samples of the D-latch output taken at the frame rate provide a binary quantized representation of the phase error.

Once phase-locked, D-Latch "a" 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 would lead to a high jitter in the recovered clock. By adding an appropriate lag-lead filter to the loop, an acceptable amount of tracking jitter can be achieved without limiting the frequency detector capture range. Fig. 9 shows the jitter of the recovered clock, using the on-chip ring oscillator VCO, measured with an HP 54120A Digital Oscilloscope. The hunting jitter is 10.8 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, Fill words 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, provided that the bit clock has a 50% duty cycle.

Data from the retiming latch "b" 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 *Master Transition* nibble.

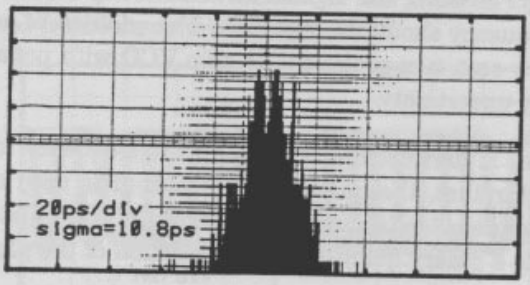

Fig. 9 Phase jitter histogram of the recovered clock in loop-back mode at 1 Gbit/s.

An on-chip voltage controlled ring-oscillator similar to that published [9], with a tuning range of 1.1 to 1.6 GHz is used with the PLL [10].

5.1. PLL Simulation

Because we chose to use a non-linear phase detector, closed form analysis of the PLL dynamic characteristic was difficult. Straightforward simulation of PLL acquisition was also troublesome due to the wide range of phenomena to be simulated, ranging from 1ns clock cycles to 500 microsecond long lock up transients. The 6-7 order of magnitude range in time scale made SPICE simulations prohibitively long to run.

To be able to experiment with various loop configurations, a relatively accurate and fast simulator was essential. Our solution was to write a time stepping simulator that operated on idealized building blocks, such as integrators, low pass filters, amplifiers, slicers, etc. By leveraging the capabilities of the C compiler and re-compiling the program for each change of the circuit, we were able to make the simulator quite small, using only about 250 lines of code. The resulting C program was simple to write and proved to give quite accurate predictions of measured performance. Fig. 10 shows the correspondence between our simulated and measured PLL frequency acquisition.

The program operates by keeping track of the voltage at each node of the circuit block diagram as an entry in an array of floating point numbers. We keep track of node values of both the current and previous time step. At each simulation time-step, the future value of each node is calculated based on the present and previous value. The node values are then recorded, the array of present values is copied to the previous array, and the future values are copied into the present array.

Euler integration techniques are used to handle building blocks such as integrators and filters [11]. As an example, a single pole RC low-pass filter with a time constant equal to τ, can be implemented as

$$V_{out}(t+\Delta t) = \frac{V_{out} + \frac{\Delta t}{\tau} V_{in}(t)}{1 + \frac{\Delta t}{\tau}}$$

Non-linear blocks such as limiting amplifiers, hysteresis elements, integrators modulo 2π... etc, which are difficult to model in SPICE are easy to handle.

Because it would be so inefficient to simulate every cycle of the VCO, we adopted a strategy where the high frequency nodes of the circuit are simulated by voltages which are proportional to the frequency of the node waveform. For example, a 1 GHz input clock would be simulated by a 1×10^9 volt signal rather than a sinusoidal waveform with a 1 ns period.

The VCO is simulated as a limiting amplifier which converts an input tuning voltage into an output voltage which is proportional to the actual output frequency. The loop frequency error is modeled by the output voltage of an amplifier which takes the difference between the incoming frequency (voltage) and the VCO frequency (voltage). The frequency error voltage is then integrated modulo 2π by an integrator to give a voltage proportional to the loop phase error. Finally, the binary quantized phase error signal representative of the D-latch phase detector output is produced by slicing the loop phase error signal.

With these simulation techniques, the overall loop lock-in transient can be simulated on an HP 9000/350 UNIX Workstation in about 3 minutes of real time.

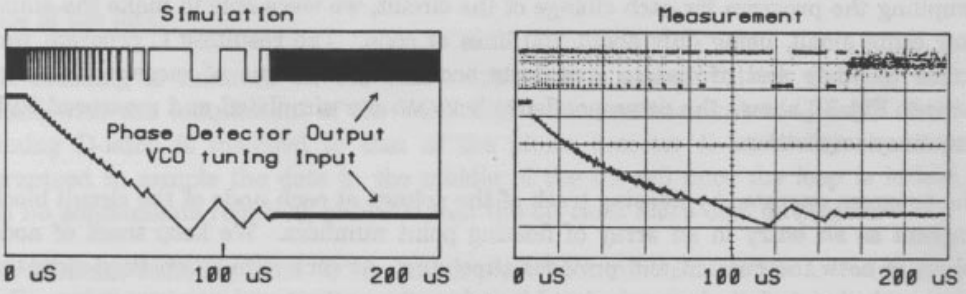

Fig. 10 Comparison of Simulated and Measured PLL Frequency Acquisition.

6. Decoder Chip

The Decoder provides the companion functions of deserialization and decoding. A block diagram is shown in Fig. 11. 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.

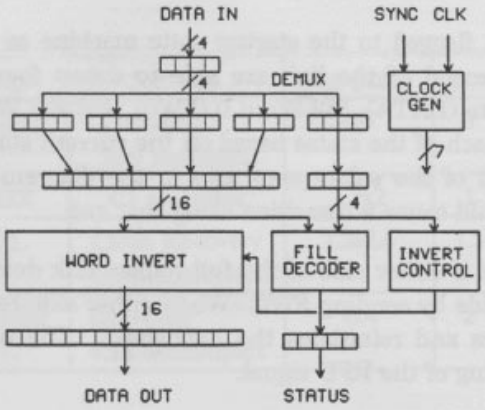

Fig. 11 Decoder Block Diagram.

7. Link Startup and Maintenance

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

Figure 12 is a state diagram describing the startup handshake procedure for a full duplex link. Both the near and far ends of the link independently follow the state diagram of Fig. 12. At power up, each end of the link enters the sequence at the arc marked "START".

Each node in the state machine has three notations. The top notation is either "FDET" or "PHASE". FDET stands for Frequency detect mode, and implies that the Frequency detector has been enabled on the PLL Chip. When the chip is in this mode, it is important that no data is being sent, as the frequency detector is only able to lock onto one of the special training Fill words: FW0, or FW1. The PHASE notation means that the PLL has been switched to phase detect mode and is ready to allow data transmission. The middle notation in each state bubble is the Fill word which is currently being sent by the node's Encoder and MUX. The last notation is the Ready For Data (RFD) status on the Encoder chip. When RFD is low, the Encoder signals the user to hold off any incoming data while it sends Fill words. When RFD is high, data is sent if available, and otherwise Fill words are sent.

The consistent presence of the two *Master Transition* bits is monitored by the Decoder chip to detect a locked condition. If the Decoder detects an unlocked

condition, then this is flagged to the startup state machine as a Frame Error. The Decoder chips at both end of the link are able to detect four different conditions: Frame Error (FE), Data (DATA), Fill Word 0 (FW0), and Fill Word 1 (FW1). Transitions are made from each of the states based on the current status condition received by the Decoder. Each of the subsequent arcs in the diagram are labeled with the relevent state that would cause a transition along that arc.

It is worth noting that if either side of the full duplex link detects a Frame Error, it will notify the other side by sending FW0. When either side receives FW0, it follows the state machine arcs and reinitiates the handshake. The user is notified of this action by the deasserting of the RFD signal.

This startup protocol provides the user with a guarantee that no data will be delivered before the link will accept any data. The use of a training sequence handshake avoids any false lock problems inherent in PLL systems which attempt to lock onto random data with wide-range VCOs.

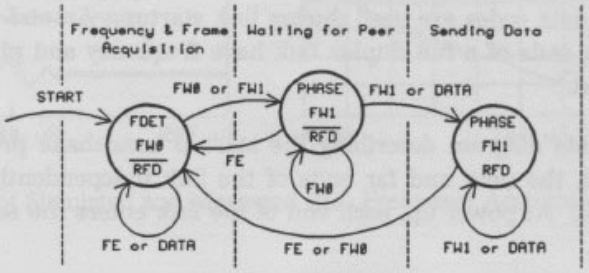

Fig. 12 Link Start Up Procedure.

8. Implementation

An overview of the four chips described is given in Table 1. The Encoder and Decoder chips are fabricated in a 5 GHz f_t bipolar process, using Emitter Function Logic [12] standard cells. They contain 2500 and 2000 active devices, respectively. Packaged parts run up to 400 MHz, capable of supporting a serialized line rate of 1.6 Gbit/s. The MUX and PLL chips are implemented in a 10 GHz 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 1.5 Gbit/s have been demonstrated.

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

Table 1. Chipset overview

9. Acknowledgments

We thank the Hewlett-Packard California Design 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.

10. References

1. Wood, I. C., M. S. J. Mudd, D. G. Taylor and P. H. Saul, A 565/680Mb/s Multiple Coder/Decoder for Optical Fiber Trunk Line Systems, *ISSCC Digest of Technical Papers 28*(1985), 156-157.

2. *Preliminary Hot Rod High-Speed Serial Link (data sheet)*, Gazelle Microcircuits, September 1989.

3. Ewen, J. F., D. L. Rogers, A. X. Widmer, F. Gfeller and C. J. Anderson, Gb/s Fiber Optic Link Adapter Chip Set, *GaAs IC Symposium Technical Digest*, November 6-9, 1988, 11-14.

4. Walker, R. C., T. Hornak, C. Yen and K. Springer, A Chipset for Gigabit Rate Data Communication, *Proceedings of the 1989 Bipolar Circuits and Technology Meeting*, September 18-19, 1989, 288-290.

5. Tani, J., D. Crandall, J. Corcoran and T. Hornak, Parallel Interface IC's for 120Mb/s Fiber Optic Links, *ISSCC Digest of Technical Papers 30*(February 1987), 190,191.

6. FDDI Physical Layer Protocol (PHY), X3T9/85-39, May 10, 1985. Draft Proposed American National Standard.

7. Crandall, D., S. Hessel, T. Hornak, R. Nordby, K. Springer and R. C. Walker, DC-Free Code for Arbitrary Data Transmission, U.S. Patent Pending, Filed Nov. 1988 .

8. Corsetto, C., T. Hornak, R. Nordby, R. C. Walker and C. Yen, Phase Locked Loop for Clock Extraction in Gigabit Rate Data Comuuncation Links, U.S Patent #4,926,447, Granted May 15, 1990.

9. Syed, K. E. and A. A. Abidi, Gigahertz Voltage Controlled Ring Oscillator, *Electronics Letters 22*, 12 (5th June 1986), 677-679.

10. Walker, R. C., A Fully Integrated High-Speed Voltage Controlled Ring Oscillator, U.S. Patent #4,884,041, Granted Nov. 28, 1989.

11. Yen, C., Z. Fazarinc and R. Wheeler, Time-Domain Skin-Effect Model for Transient Analysis of Lossy Transmission Lines, *Proceedings of the IEEE 70*, 7 (July 1982), 750-757.

12. Skokan, Z. E., EFL Logic Family for LSI, *ISSCC Digest of Technical Papers 16*(February 1973), 162-163.