

A Two-Chip 1.5-GBd Serial Link Interface

Richard C. Walker, Cheryl L. Stout, *Member, IEEE*, Jieh-Tsong Wu, *Member, IEEE*, Benny Lai, *Member, IEEE*, Chu-Sun Yen, *Member, IEEE*, Tom Hornak, *Fellow, IEEE*, and Patrick T. Petruno

Abstract—A silicon bipolar transmitter and receiver chip pair transfers parallel data across a 1.5-GBd serial link. A new “conditional-invert master transition” code and phase-locked loop are described and analyzed that provide adjustment-free clock recovery and frame synchronization. The packaged parts require no external components and operate over a range of 700 to 1500 MHz using an on-chip VCO. The line code and handshake protocol have been accepted by the Serial-HIPPI implementor’s group for serially transmitting 800-Mb/s HIPPI data, an ANSI standard, and by SCI-FI, an IEEE standard for interconnecting cooperating computers.

I. INTRODUCTION

PARALLEL computers, high-resolution graphics, and network backbones are among the many applications that could immediately benefit from inexpensive, compact, and easy-to-use gigabit-rate fiber-optic data links. Serial links have been widely used for telecom applications, however, parallel data interfaces are required for convenient connection to computer equipment.

The use of fiber media for gigabit-rate computer communication has been limited by the lack of low-cost link interface chips. An earlier four-chip chip set [1] established the feasibility of several integrable circuit techniques to achieve these data rates, but was difficult to use because of the high-speed chip interconnections and extra support circuitry required.

In this paper we report a monolithic transmitter (TX) and receiver (RX) chip pair that can be used for the transmission of parallel data, and that requires no external active components. From the user’s viewpoint, this chip set implements a full-duplex “virtual ribbon cable” interface (Fig. 1). For short-distance applications, an on-chip equalizer is provided to allow the use of coaxial cables rather than a more costly fiber link. The chips require no external frequency-determining elements or user adjustments and operate over a range of 700 to 1500 MHz using an on-chip VCO. A state machine controller (SMC) is also implemented on the RX chip to transparently handle a




Fig. 1. Block diagram of full-duplex link built from two chip-set pairs.

start-up handshake protocol. This work is the highest speed link interface chip set reported to date at this level of functionality and integration.

The architecture of the link is largely determined by the line code design, which is discussed next.

II. LINE CODE, CLOCK, AND FRAME SYNCHRONIZATION

Codes used for fiber-optic links are dc-balanced to permit the regulation of laser bias current by simply maintaining a fixed average optical power. Balanced data streams may also be conveniently ac-coupled at the receiver without incurring extra baseline wander or jitter. The “conditional-invert master transition” (CIMT) code used in this chip set transmits the parallel data words in either true or complement form, as needed, to maintain dc balance on the line.

To make the decision of whether or not to invert a data frame, the TX chip uses a majority gate, built from a DAC-like current summing circuit and comparator, to compute the polarity of the incoming frame. The frame polarity is compared against the sign of an up/down counter, which keeps track of the total disparity of transmitted bits. If the two signs agree, the frame is sent inverted. Otherwise it is sent uninverted. As shown in Fig. 2, four extra coding bits create a coding field (C-field) which is then appended to the data field (D-field) during transmission. The chip set is programmable to allow the transmission of either 16 or 20 b of data to produce a 20- or 24-b line code frame. In addition, an extra input FLAG bit is also available as an extra data bit, thereby increasing the data bits to 17 or 21, or can be internally toggled by the transmitter to allow enhanced receiver frame error detection. The 17-b form of the line code has been accepted as the standard code for the IEEE P-1596 Scalable Co-

Manuscript received April 26, 1992; revised July 25, 1992.

R. C. Walker, C. L. Stout, C.-S. Yen, and T. Hornak are with Hewlett-Packard Company, Palo Alto, CA 94304.

J.-T. Wu was with the Microwave Semiconductor Division, Hewlett-Packard Company, San Jose, CA. He is now with the Department of Electronics Engineering, National Chiao-Tung University, Hsin-chu, 300, Taiwan, Republic of China.

B. Lai and P. T. Petruno are with the Communications Components Division, Hewlett-Packard Company, San Jose, CA 95131.

IEEE Log Number 9203486.

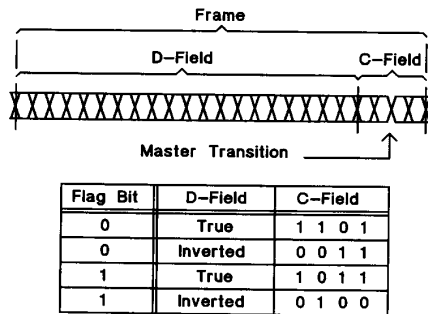


Fig. 2. Line code.

herent Interface Group [2], and the 21-b form of the code has been accepted by the Ad-Hoc High Performance Parallel Interface (HIPPI) Serial Implementors Group [3].

The central pair of bits in the C-field are always complementary and provide a "master transition" phase reference for the receiver phase-locked loop (PLL). This master transition is used by the PLL as the phase reference for both bit and frame clock recovery. The frame clock is used by the demultiplexer for frame alignment. 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 undesirable.

The polarity of the master transition is used to encode the extra FLAG bit. The other two bits in the C-field are used to signal whether a given frame represents inverted data, noninverted data, inverted control, noninverted control, or fill. Control frames are special non-data frames that can be used as packet headers, trailers, and other protocol-specific information. Fill frames have only a single rising edge at the master transition location, and are used as training sequences to provide unambiguous frequency, phase, and frame acquisition during link start-up. There are two logical fill frames: FF0, which is 50% duty-cycle square wave, and FF1L and FF1H, which are 2 bits light and 2 bits heavy, respectively. The two FF1 forms are sent in alternation to maintain dc balance.

Coding schemes that satisfy the needs of clock recovery and dc balance are a trade-off between coder complexity and bandwidth utilization. Simple Manchester coders have an efficiency of only 50%: two symbols are sent for each received bit. Other codes, such as the 4B/5B code used in FDDI [4] and the 8B/10B code proposed for fiber channel [5], are more efficient than Manchester, but are more complex to implement. Our code is only moderately complex to encode, very simple to decode, accommodates variable data widths, and achieves a high efficiency of 21/24, or 87.5%

III. CHIP BLOCK DIAGRAMS

A simplified block diagram of the TX chip is shown in Fig. 3. The user supplies the parallel inputs, D0-D19, a




Fig. 3. Simplified block diagram of TX chip.




Fig. 4. Simplified block diagram of the RX chip.

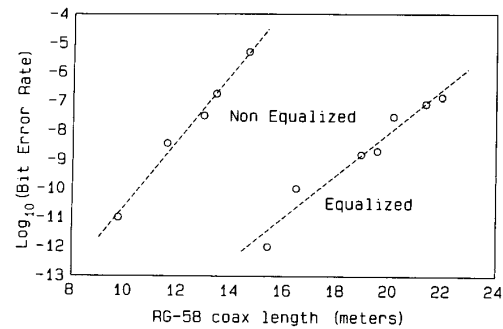


Fig. 5. Bit error rate performance with/without equalizer.

frame rate clock, the DAV and CAV inputs, and the FLAG input (optional). The PLL/clock generator block generates the high-speed serial clock by phase locking onto the incoming low-speed clock, which can be either at the full or one-half frame rate.

Before transmission, the dc balance of each frame is determined by a segmented DAC and a comparator. The long-term dc balance of the previously transmitted data is monitored by an up/down counter that counts up on ONES and down on ZEROS. To determine whether a given frame should be sent in either true or complement form, the dc balance of each frame is compared with the long-term disparity of the previously transmitted bits.

The frame is then serialized with a circuit that multiplexes the parallel inputs into a serial bit stream and performs any necessary frame inversion. The output of this block is driven off-chip to be transmitted across the link.

The receiver block diagram is shown in Fig. 4. The data path consists of an input selector, two input sampling latches, a demultiplexer, a C-field decoder, and a D-field decoder. The input (DIN) is programmable to select data

from either the normal data input, a loopback data input, or an equalized input. The equalized input provides a 3-dB boost at 600 MHz to compensate for skin loss in long coaxial lines [6]. The improvement of link BER with the equalizer used with RG-58 coax is shown in Fig. 5. For a given BER, the equalizer extends the usable link length by over 50%.

IV. PHASE-LOCKED LOOP DESCRIPTION

The on-chip PLL's used in both the TX and RX chips are nearly identical. For simplicity, this section only describes the implementation of the RX PLL.

The incoming data stream is latched by two matched D-latches, one on the rising edge, and the other on the falling edge of the bit-rate VCO clock. When the loop is locked, the rising-edge retiming latch samples the center of each data bit and produces retimed data. The falling-edge phase detector latch samples the transitions between bits. The transition sample corresponding to the master transition is selected for use as a phase error indication. Since the code allows the master transition to be of either polarity, the sample is corrected for transition polarity by being XORed with the immediately preceding data bit to derive a binary-quantized (bang-bang) phase error indication. Because the phase detector and retiming latches are matched, assuming a 50% duty cycle VCO, the retiming clock phase is inherently aligned to the center of the bit cell over both process and temperature variation. In addition, the circuit can operate at the full speed at which a process is capable of building a functioning latch.

Because the internal VCO is capable of operating over nearly a 3:1 range of frequencies, a frequency detector is necessary to avoid false locking problems. The frequency detector operates only when simple square-wave fill frames are being sent. A conventional sequential frequency detector determines the sign of the frequency error. A gating circuit drives the loop filter with the frequency error information whenever the phase error is greater than $\pm 22.5^\circ$. When the phase error is less than this amount, the output of the phase detector is used. Before data are allowed to be sent, the state machine controller disables the frequency detection circuit.

A. Loop Dynamics

The phase detector described is nonlinear, and conventional linear PLL theory is not useful for design or analysis. Precise loop behavior can be simulated efficiently with time-step simulators, but this is cumbersome to use for routine design. Fortunately a simple decomposition of the loop provides accurate closed-form expressions for both loop tracking jitter and loop stability. An outline of this analysis is given in this section.

A simplified version of the clock recovery loop that assumes a fixed, rising, master transition is shown in Fig. 6. The transition samples are decimated by the number of bits per frame, M , to isolate the one sample corresponding to the master transition.




Fig. 6. Simplified clock recovery loop.

If certain assumptions are met, as described in a later section, we can consider the system to be composed of two noninteracting loops. These are the loops labeled “bang-bang branch” and “integral branch” in Fig. 6. The first loop includes the connection of the phase detector to the VCO input through the bang-bang branch of the loop filter, while the second loop includes the integral branch of the loop filter. The binary control, or “bang-bang” loop, can be considered a phase tracking loop, while the integral branch can be viewed as a frequency tracking loop.

The fact that the phase detector output is quantized implies that the loop behavior will be oscillatory. In steady-state conditions, the output of the phase detector (due to inevitable noise and jitter) will be a quasi-random string of ONES and ZEROS, which will program the VCO frequency to switch between two discrete frequencies, causing the VCO to ramp up and down in phase, thereby tracking the incoming signal phase. The phase detector output tends to alternate every frame, so that, other than the dc component, the bulk of the phase detector output spectrum falls outside the effective passband of the integrator branch of the loop, and can be practically neglected.

The integrator branch then operates on just the dc component of the phase detector output. Its job is to servo the center frequency of the VCO so that the two discrete VCO frequencies programmed by the bang-bang input will always bracket the frequency of the incoming data signal. This frequency adjustment occurs so slowly that it does not materially affect the operation of the high frequency bang-bang portion of the loop.

B. The Proportional Branch of the Loop Filter

With a locked loop and assuming that the integrator output changes negligibly during a single phase update, the VCO frequency step programmed by the bang-bang tuning input is

$$F_{\text{step}} = \pm \frac{\beta V_{\phi} K_{\text{vco}}}{2}$$

where $\beta = K_{\text{vco}}/K_{\text{bb}}$, the ratio of the VCO wide range tuning gain to the bang-bang tuning gain. V_{ϕ} is the peak-to-peak voltage from the phase detector, and K_{vco} is the VCO gain constant in hertz per volt.

Assuming a high dc gain in the loop integrator, the steady-state duty cycle from the phase detector output will be very close to 50%, usually alternating between ZERO and ONE with an occasional doubling-up of bits to compensate for leakage in the integrator. The worst-case walk-off in degrees is then given by the phase walk-off of two successive update periods:

$$\text{Jitter}_{\text{peak/peak}} = \frac{360\beta V_{\phi} K_{\text{vco}} M}{F_{\text{vco}}^2}$$

where M is the loop division ratio and F_{vco} is the nominal VCO frequency in hertz. In our loop, the hunting jitter is designed to be below 18 ps rms.

Before turning to the analysis of the integral branch of the loop, we need to derive the dc component, or duty cycle of the phase detector output stream. As already mentioned, in steady state the duty cycle is 50%. Because the loop is phase-locked, the average frequency of the VCO is, on the average, equal to the frequency of the serial data stream. If the incoming frequency is switched from F_{vco} to $F_{\text{vco}} + \Delta F$, with $-F_{\text{step}} \leq \Delta F \leq F_{\text{step}}$, the duty cycle C of the phase detector will necessarily shift such that

$$F_{\text{vco}} + \Delta F = C(F_{\text{vco}} + F_{\text{step}}) + (1 - C)(F_{\text{vco}} - F_{\text{step}}).$$

Solving for the duty cycle

$$C = \frac{\Delta F}{2F_{\text{step}}} + \frac{1}{2}.$$

Unlike a traditional PLL, this result shows that the dc component of the phase detector output is proportional to frequency rather than phase. The effective gain constant of this "virtual" frequency detector, K_f , in volts per hertz is

$$K_f = \frac{V_{\phi}}{2F_{\text{step}}}.$$

Both the binary-quantized phase detector and the bang-bang branch of the loop are replaced by an equivalent linear "virtual" frequency detector with gain constant K_f . Standard linear feedback theory can then be easily used to determine the bandwidth and other salient characteristics of this loop. The unusual result is that the low-frequency loop is only first order.

Because the phase-detector dc component is proportional to frequency rather than phase, an implicit integration does not appear in the loop transfer function. This means that there is no jitter buildup due to the action of the low-frequency integrator. The jitter statistics are simply dominated by the hunting behavior of the high-frequency portion of the loop. However, unlike a normal first-order loop, the behavior of the bang-bang portion of the loop ensures that the average loop phase error remains zero with changes in input data frequency.

C. Loop Stability Criteria

The preceding analysis assumed that the two branches of the loop were essentially noninteracting. For this to be




Fig. 7. VCO delay cell.

true, it is important that the loop be set up so that, between phase sample update times ($t_{\text{update}} = F_{\text{vco}}/M$), the phase walk-off of the bang-bang branch of the loop, $\Phi_{bb}(t)$, must dominate over the phase walk-off of the integral branch, $\Phi_{int.}(t)$.

Taking the ratio of $\Phi_{bb}(t)$ and $\Phi_{int.}(t)$ at the end of one frame update time gives a figure of merit ξ for the loop stability:

$$\xi = \frac{2\beta\tau}{t_{\text{update}}}.$$

ξ must be greater than one for the two loops to be considered noninteracting. In fact, if ξ becomes significantly less than 1, the "bang-bang" portion of loop will no longer stabilize the system and large low-frequency second-order oscillations may occur in the loop.

V. ON-CHIP VOLTAGE-CONTROLLED OSCILLATOR

The VCO is composed of a cascade of three variable delay blocks as shown in Fig. 7. The low-frequency signal from the integral branch of the loop drives the main tuning input which is bandwidth limited to reduce its sensitivity to on-chip noise, and which tunes over a 700–1800-MHz range by interpolating between delay gates. The bang-bang tuning input has a wide tuning bandwidth, but only produces about a $\pm 0.1\%$ variation in VCO center frequency by modulating the base charge in $Q1$ and $Q2$. Fig. 8 shows the measured VCO tuning curve at three different power supply voltages: -4.5 , 5.0 , and -5.5 .

VI. STATE MACHINE CONTROLLER

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

Fig. 9 is a state diagram describing the start-up handshake procedure for a full-duplex link. Both the near and far ends of the link independently follow the state diagram of Fig. 9. 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

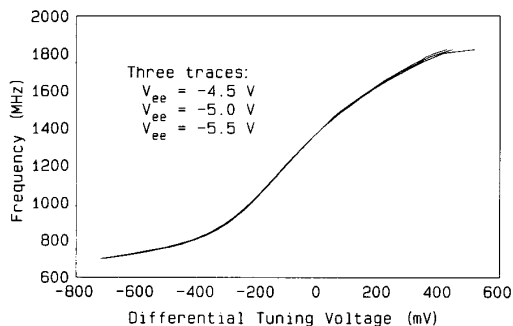
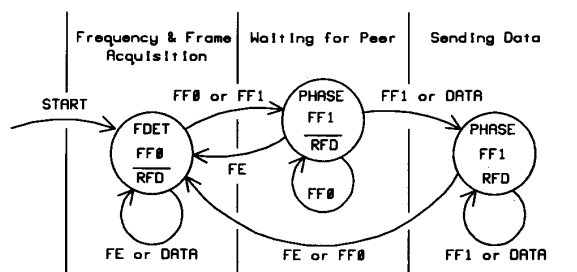
Fig. 8. Measured VCO tuning curve at 4.5-, 5.0-, and 5.5-V V_{ee} .

Fig. 9. Link start-up state machine.

frequency detector has been enabled in the RX chip PLL. 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 frames: FF0 or FF1. The PHASE notation means that the RX 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 frame sent by the node's TX chip. The last notation is the ready for data (RFD) status on the TX chip. When RFD is low, the TX chip signals the user to hold off any incoming data while it is sending fill frames. When RFD is high, data are sent if available, and otherwise fill frames are sent to maintain link synchronization.

The consistent presence of the two master transition bits is monitored by the RX chip to detect a locked condition. If the RX chip detects an unlocked condition, then this is flagged to the start-up state machine as a frame error. The RX chips at both end of the link are able to detect four different conditions: frame error (FE), data (DATA), fill frame 0 (FF0), and fill frame 1 (FF1). Transitions are made from each of the states based on the current status condition received by the RX chip. Each of the subsequent arcs in the diagram is labeled with the relevant state that would cause a transition along that arc.

If either side of the full-duplex link detects a frame error, it will notify the other side by sending FF0. When either side receives FF0, it follows the state machine arcs and reinitiates the handshake process. The user is notified of this action by the deasserting of the RFD signal.

This start-up protocol ensures that no user data is sent




Fig. 10. TX input clock, transmitted frame, and RX recovered clock at 1.5 GBd.




Fig. 11. Phase jitter histogram of RX recovered clock at 1.5 GBd.

until the link connectivity is fully established. The handshaked training sequence eliminates the false lock problem inherent in wide-range random data PLL systems.

For fiber applications, the SMC can also provide laser eye safety, with the addition of an external timer and an OR gate, by pulsing the laser at a low duty cycle when a fiber is broken.

VII. IMPLEMENTATION

The two chips were implemented in a three-level metal, 25-GHz f_t , silicon bipolar process [7] using full-custom differential 4.5-V ECL design. Both chips with their bypass and integrating capacitors are housed in a custom 68-pin surface-mount package. The 1.8-W TX and 2.0-W RX chips are each 3.5×3.5 mm in size and utilize 6100 and 6600 active devices, respectively. Both chips were fully functional at first silicon.

Fig. 10 shows the TX input clock at the parallel word rate, the transmitted frame, and the RX recovered clock at 1.5 GBd. Fig. 11 shows a phase jitter histogram of RX recovered clock at 1.5 GBd demonstrating a loop hunting jitter of 8-ps rms.

ACKNOWLEDGMENT

The authors thank K. Springer, R. Nordby, C. Corsetto, and D. Crandall for their early contributions to this work. We also appreciate the work of D. Yoo, J. Norman, N. McAfee, and L. Lian-Mueller who did the assembly and packaging of the chip set. Special thanks to D. Pettengill, S.-Y. Chiang, and the staff of the Hewlett-Packard Circuit Technology Research and Development group for process development and wafer processing.

REFERENCES

- [1] R. C. Walker, T. Hornak, C. Yen, J. Doernberg, and K. H. Springer, "A 1.5 Gb/s link interface chipset for computer data transmission," *IEEE J. Selected Areas Commun.*, vol. 9, no. 5, pp. 698-703, June 1991.
- [2] *SCI Scalable Coherent Interface (Draft)*, IEEE P1596, Nov. 10, 1990.
- [3] *Serial-HIPPI Specification, Revision 1.0*, May 17, 1991.
- [4] *FDDI Physical Layer Protocol (PHY) (Draft)*, ANSI X3T9/85-39, May 10, 1985.
- [5] *Fiber Channel Physical Layer Medium Dependent (FC-0) (Draft)*, ANSI X3T9.3, Mar. 5, 1990.
- [6] C. Yen, Z. Fazarinc, and R. Wheeler, "Time-domain skin-effect model for transient analysis of lossy transmission lines," *Proc. IEEE*, vol. 70, pp. 750-752, July 1982.
- [7] W. M. Huang *et al.*, "A high-speed bipolar technology featuring self-aligned single-poly base and submicrometer emitter contacts," *IEEE Electron Device Lett.*, vol. 11, no. 9, pp. 412-414, Sept. 1990.



Richard C. Walker was born in San Rafael, CA, in 1960. He received the B.S. degree in engineering and applied science from the California Institute of Technology, Pasadena, in 1982, and the M.S. degree in computer science from the California State University, Chico, in 1992.

He joined Hewlett-Packard Laboratories, Palo Alto, CA, as a Member of the Technical Staff in 1981. Since that time, he has worked in the areas of broad-band cable modem design, solid-state laser characterization, phase-locked-loop theory,


and high-speed circuit design for both Si and GaAs IC processes. He is presently a Principal Project Engineer in the Instrument and Photonics Laboratory. He has authored or co-authored 16 technical papers and holds six patents.



Cheryl L. Stout (S'78-M'83) received the B.S.E.E. degree from San Jose State University, San Jose, CA, in 1979 and the M.S.E.E. degree from the University of California at Berkeley in 1983.

From 1979 to 1982 she was involved in the development of optical communication products at Plantronics, Inc. Since 1983 she has been a Member of the Technical Staff at Hewlett-Packard Laboratories in Palo Alto, CA, where she has designed high-speed silicon and GaAs integrated

circuits for optical communication systems and test instruments.



Jieh-Tsong Wu (S'83-M'87) was born in Taipei, Taiwan, on August 31, 1958. He received the B.S. degree in electronics engineering from National Chiao-Tung University, Taiwan, in 1980, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1983 and 1988, respectively.


From 1980 to 1982 he served in the Chinese Army as a Radar Technical Officer. From 1982 to 1988 at Stanford University, he focused his research on high-speed analog-to-digital conversion in CMOS VLSI. From 1988 to 1992 he was a Member of Technical Staff at Hewlett-Packard Microwave Semiconductor Division in San Jose, CA, and was responsible for several linear and digital gigahertz IC designs. Since 1992 he has been an Associate Professor at National Chiao-Tung University in Hsin-Chu, Taiwan. His research interests include integrated circuits and systems for optical-fiber and microwave telecommunications systems.

Dr. Wu is a member of Phi Tau Phi.



Benny Lai (S'83-M'83) was born in Hong Kong on March 19, 1960. He received the B.S. and M.S. degrees in electrical engineering computer science from the University of California at Berkeley in 1982 and 1983, respectively.

From 1982 to 1991 he was a Member of the Technical Staff at the Microwave Semiconductor Division of Hewlett-Packard Company in San Jose, CA, where he was engaged in the research and development of microwave transistors, wide-band hybrid amplifiers, and digital components for digital communication. Currently, he is the Principal Member of the Technical Staff at the Communications Components Division of Hewlett-Packard in San Jose, CA, where he is engaged in the research and development of high-speed digital data link and lightwave systems. He has authored and coauthored seven papers and received one patent and two patents pending in the area of high-speed circuits for communication systems.



Chu-Sun Yen (S'61-M'62) was born in China in 1933. He received the B.S. degree from the National Taiwan University in 1955, the M.S. degree from the University of Florida, Gainesville, in 1958, and the Ph.D. degree from Stanford University, Stanford, CA, in 1961, all in electrical engineering.


Since 1961 he has been with Hewlett-Packard Laboratories, Palo Alto, CA. He is presently a Project Manager in the Instruments and Photonics Lab.



Tom Hornak (SM'70-F'85) graduated from the Bratislava Slovak Technical University, Czechoslovakia, in 1947, receiving the Dipl. Ing. in electrical engineering. In 1966 he earned the Ph.D. degree from the Czech Technical University in Prague.

From 1947 to 1961 he worked at the Tesla Corporation's Radiotechnical Research Laboratory. Between 1961 and 1968 he was head advisor for electronics R&D at the Computer Research Institute in Prague. He left Czechoslovakia in 1968 and joined Hewlett-Packard's Corporate Research Laboratories the same year. In 1973 he became Lab Department Manager and is presently a Principal Scientist at Hewlett-Packard Laboratories, Palo Alto, CA. In the Corporate Research Laboratory, his research interest is in high-speed data communication links, high-speed analog/digital interfaces, and electronic instrumentation utilizing Hewlett-Packard's advanced Si and GaAs IC processes. He made significant contributions to read-only memories, analog-to-digital converters, charge-coupled devices, and fiber-optic communication systems. His past technical experience also includes work on radar, pulse instrumentation, and magnetic memories. He has published over 40 papers on electronics and holds more than 30 patents.

Dr. Hornak was a lecturer at the Czech Technical University in Prague. He was a Guest Editor in 1978 and Associate Editor from 1986 to 1988 of the *IEEE JOURNAL OF SOLID-STATE CIRCUITS*. From 1979 to 1981 he was Chairman of the IEEE Solid-State Circuits and Technology Committee, and from 1978 to 1982 a member of the IEEE International Solid-State Circuits Conference Program Committee.



Patrick T. Petruno was born in Allentown, PA, on June 4, 1954. He received the B.S.E.E. and M.S.E.E. degrees in electrical engineering from Pennsylvania State University, University Park, in 1976 and 1978, respectively.

From 1978 to 1992 he was employed by the Microwave Semiconductor Division of Hewlett-Packard Company, San Jose, CA, where he was a Member of the Technical Staff and was involved in the research and development of silicon bipolar microwave transistors, wide-band hybrid and in-

tegrated bipolar amplifiers, and high-speed flip-flops and dividers. After 1982 he managed a number of component projects in the high-speed digital transmission, digital cellular radio, and GaAs IC areas which include wide-band variable gain amplifiers, high-speed multiplexer/demultiplexers, monolithic clock extraction circuits, I/Q modulators, and monolithic switches. Currently he is R&D Section Manager for the Communications Components Division of Hewlett-Packard, San Jose, CA, managing component development of high-speed digital data link and lightwave systems. He has coauthored three papers in the area of high-speed circuits for communications systems.