A 160-kb/s Digital Subscriber Loop Transceiver with Memory Compensation Echo Canceller

ROGER P. COLBECK, MEMBER, IEEE, AND PETER B. GILLINGHAM, MEMBER, IEEE

Abstract—**A full-duplex transceiver chip incorporating an adaptive echo cancelling modem and a 2.048-Mb/s serial interface is described. The device provides a full-duplex communication link at 160 or 80 kb/s on up to 4 or 5 km, respectively, of 0.5-mm twisted-pair cable. Full integration is achieved through the use of RAM-based sign-algorithm echo-cancellation, biphase line code, a fixed switched-capacitor equalizer and a digital phase locked loop. The paper emphasizes system design considerations and a chip architecture minimizing power dissipation, silicon area and off-chip components. A double poly 3-μm CMOS technology is used to implement the 5-V 22-pin device which dissipates less that 50 mW and occupies 27.7 mm^2.**

I. Introduction

TO SATISFY the requirements for an Integrated Services Digital Network (ISDN), a cost-effective solution to high speed data communication over existing twisted-pair cable is necessary. Two competing techniques, the time compression multiplexing (TCM) ping-pong system and the full-duplex echo cancelling scheme, have emerged to meet these requirements. Although in theory the latter can achieve longer line lengths, it has been argued that echo cancellation demands greater circuit complexity and higher precision signal processing. Until recently the development of integrated digital echo cancelling hybrids has been restricted to test chips [1], while several fully integrated TCM chips have already been announced [2], [3]. A RAM based echo canceller design [4], [5], which combines an easily implemented algorithm with a simple analog front end, led to the development of a single chip transceiver [6].

The major requirements for a digital transceiver chip are high performance over a variety of line lengths, gauge changes and bridge taps, low-power dissipation to enable remote power feeding, and small die area to minimize cost. This paper describes the design of a CMOS chip which satisfies these conditions, providing a high speed communication link over twisted pair in both the PABX environment and in the medium range loop plant. Section II introduces the transmit side of the chip including the serial telecommunication bus interface, and discusses the choice of line-code. Section III describes the receive path filters and clock recovery circuitry. We discuss the sign algorithm RAM-based echo cancellation technique as it applies to twisted pair transmission and show some simulation results in Section IV. Finally, in Section V, the overall performance of the device on a variety of lines is summarized.

Manuscript received April 23, 1985; revised September 30, 1985.
The authors are with Mitel Corporation, Kanata, Ont., Canada K2K 1X3.
IEEE Log Number 8406484.

II. Transmit Path

Fig. 1 shows a block diagram of the transceiver. A 2.048-Mb/s serial PCM highway supplies the transmit data to the chip. It contains 32 8-bit channels cycling at 8 kHz, only 3 or 4 of which are utilized by the transceiver. If a line rate of 160 kb/s is selected, two 64 kb/s B-channels, one 16 kb/s D-channel, and one 64 kb/s C-channel are extracted from the PCM highway. When the 80 kb/s line rate is enabled only 3 channels of information on the serial PCM highway are used, a single 64 kb/s B-channel, an 8 kb/s D-channel and the C-channel. The B-channels are typically assigned to user voice or data, while the D-channel is normally reserved for signalling information. The C-channel is not transmitted over the line interface, but used internally to configure the chip for various modes of operation. These include selection of 80 or 160 kb/s line rates, accelerated echo canceller convergence, and a suite of test functions spanning from system oriented analog and digital loopback to modes allowing every functional block of the circuit to be tested individually. A separate PCM highway carries the received data stream. The B and D channels are passed through transparently from the line interface while the C-channel holds the internal status of the chip, indicating SYNC and the quality of the received data. For applications not requiring the 2.048-MHz interface the chip can be defaulted into baseband modem operation via external mode select pins. The mode select pins also program the part for SET or CO operation.

The serial transmit data stream is then scrambled before insertion of a SYNC bit through a pseudorandom sequence generator having the polynomial $1+x^{-3}+x^{-5}$. This ensures that data will not emulate the SYNC bit, causing the receiver to lock on out of phase to the true SYNC. Data is scrambled again with polynomial $1+x^{-6}+x^{-7}$ in the

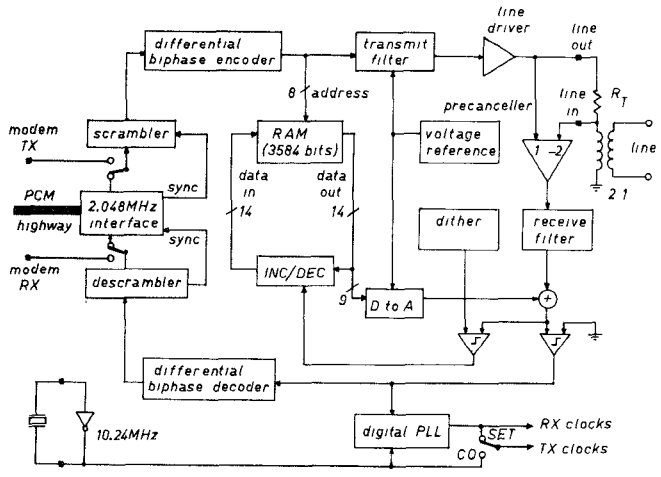

Fig. 1. Transceiver block diagram.

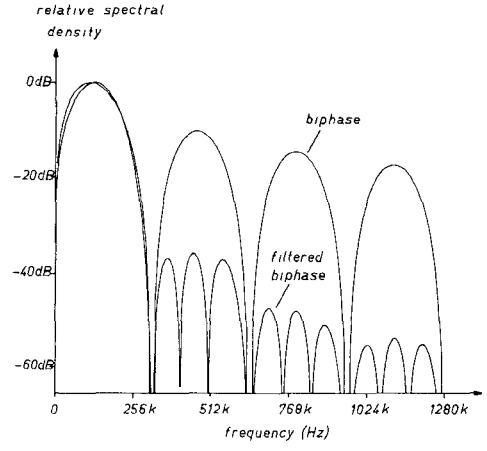

Fig. 2. Simulated TX filter response for 160-kb/s data rate.

SET device or $1+x^{-4}+x^{-9}$ for the CO. Scrambling is necessary for predictable behavior of the echo canceller. Different scrambler polynomials in SET and CO must be used to reduce correlation between the near and far-end signals which could result in the echo canceller converging to far-end data.

Next, the data is differentially biphase encoded. Differential encoding of the baseband serial data makes the receiver polarity independent. The tip and ring leads can be reversed without effect on data recovery. The biphase line-code was chosen for a number of reasons. First, it has the desirable properties of no dc component, necessary for line powered terminals, and a transition within each baud which allows easy timing recovery. Second, the primary spectral lobe of 80 or 160 kb/s biphase falls within the usable bandwidth of twisted pair cable so that a simple fixed equalizer can be used. Finally, biphase has a short impulse response which minimizes the size of the echo canceller RAM.

The differentially encoded biphase is conditioned through a 15-stage switched-capacitor FIR filter to remove unnecessary spectral content in the transmitted output. The filter is clocked at sixteen times the baud-rate (16 f_b), either 1.28 or 2.56 MHz in 80 or 160 kb/s modes, respectively. While it may be argued that this filtering could be performed in the receiver without any change in the overall transfer function there is a number of good reasons for doing it prior to transmission. First, the line driver must drive 800 Ω with a signal around 4 V peak to peak. The difference between the power in a square wave of this magnitude and the filtered signal is actually significant with respect to overall chip power consumption. Second, properties of the digital signal to be transmitted make it much easier to filter on the transmit side. The delay elements of the FIR structure can be flip-flops as opposed to analog sample-and-hold circuits. Finally, as a responsible user of the electromagnetic spectrum, it is undesirable to broadcast wide-band signals which could interfere with other systems.

The transmit filter is a linear-phase symmetrical-coefficient switched-capacitor FIR low-pass filter with zeros located at the secondary spectral peaks of a random biphase signal. Fig. 2 shows simulated response in 160 kb/s mode, where the secondary lobes of the biphase spectrum are attenuated at least 35 dB with respect to the primary lobe without any phase distortion. In the passband a 4-dB pre-emphasis peak with respect to dc is included to offset line attenuation without distorting significantly the eye. Less emphasis would result in larger low frequency components giving a longer transhybrid response and poorer echo cancelling performance. Overemphasis results in gross amplitude differences between the two eye trajectories bringing the low frequency components closer to the noise floor. The 15-tap filter samples the digital input stream at 16 f_b so that the output is based on at most two bauds of information, such that unnecessary intersymbol interference in the transmitted eye is kept to a minimum.

The TX filter and indeed all switched-capacitor filters in the chip (Fig. 3) require the use of a high bandwidth, high slew rate operational amplifier to accomodate passbands on the order of the baud rate and clock frequencies as high as 2.56 MHz. The core amplifier is based on the folded cascode configuration [7] and is compensated by the capacitive load on the output. Internally compensated amplifiers must ensure that the frequency of the parasitic output pole is high enough not to cause instability by virtue of a low-output impedance. This is accomplished at the expense of supply current. In contrast, a transconductance amplifier takes advantage of the natural positions of the internal parasitic pole and dominant load pole with a high-output impedance, and achieves low-power consumption.

The on-chip voltage reference using a vertical n-p-n transistor-based bandgap circuit provides a −1.8-V reference with respect to an internally generated midrail for use in both the echo canceller and TX filter. The peak-to-peak echo which must be covered by the echo canceller D/A converter is a function of the TX filter output level. As both circuits share the same reference, the chip is self-compensating for variations in the voltage reference, reducing the requirement for a high precision trimmable circuit.

External line interface circuitry consists of a 2:1 transformer with a split winding on the line side to permit remote power feed and a line termination network, which

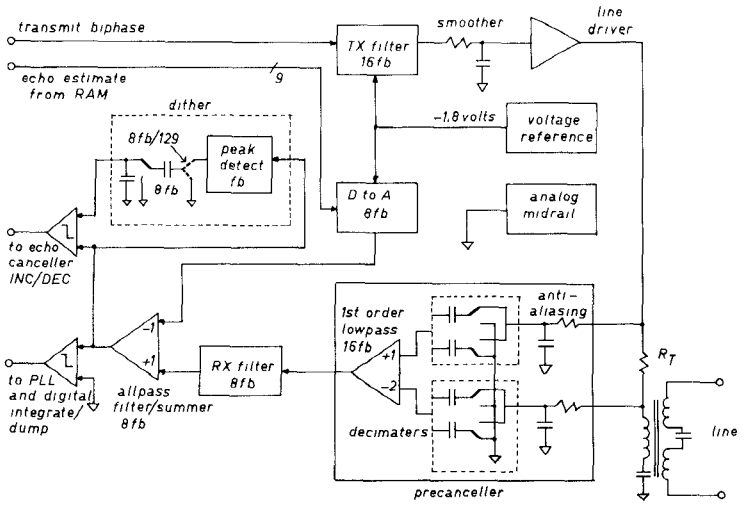

Fig 3 Block diagram of analog circuitry.

can be as simple as a single 400-Ω resistor, connected between the line driver and transformer. The composite line signal at the point between transformer and termination resistor, which contains both near- and far-end biphase signals, is routed back into the chip to the first stage of the receive path.

III. Receive Path

The first stage of the receive path, shown in Fig. 3, is a precanceller realizing a fixed summation of the transmit and line signals to achieve limited cancellation of the TX component. Reducing the amplitude of the near-end signal improves the dynamic range of the receiver and the echo canceller. With a purely resistive termination, a worst-case precancelled signal −9 dB with respect to the line signal was observed over a variety of different lines. Precancellation allows an additional 9 dB of gain without risk of clipping the near-end signal. Clipping of the far-end signal superimposed on the near-end echo is acceptable but if the near-end signal itself were clipped the far-end data would be masked out. Included in this section is an anti-aliasing filter on each of the LIN and LOUT inputs consisting of a continuous-time single-order low-pass filter. These were realized as distributed *RC*'s for fast rolloff having a −3-dB point at 3.2 f_b. Following these are switched capacitor decimators with an effective input sampling rate of 32 f_b for the suppression of the replicated lobe at 16 f_b. Overall anti-aliasing performance is such that frequencies above the Nyquist limit are attenuated at least 20 dB with respect to the passband.

A third-order switched-capacitor bandpass filter serves as a fixed equalizer optimized for best performance at long line lengths. It is comprised of a first order lowpass section, incorporated in the pre-canceller, clocked at 16 f_b and a bilinear bandpass biquad clocked at 8 f_b. The composite response of the receive path, shown in Fig. 4 for a 160-kb/s transmission rate, has a passband between 64 and 256 kHz with a dc zero and a bilinear notch at 640 kHz. When the transceiver is in the 80-kb/s mode the filter responses are compressed in frequency by a factor of two by scaling the

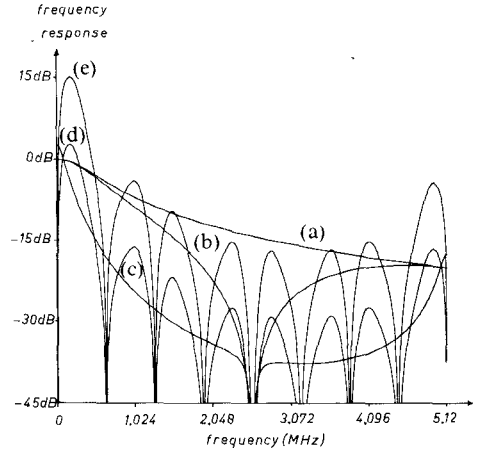

Fig 4 Receive path simulated frequency response. Progressive filtering is shown through: (a) Anti-aliasing filter, (b) Decimator, (c) Precanceller, (d) Bandpass Biquad, (e) All pass.

clocks. The time constant of the *RC* antialiasing filter is doubled by switching in additional resistance and capacitance. The receive filter emphasizes the higher frequency components of the biphase signal to offset the greater attenuation experienced by these components on twisted-pair cables. It also introduces delay which is greater for high frequencies to compensate for the dispersive effects of long lines. Fixed equalization in the receive filter is intended to optimize performance over long lines at the expense of short-line performance. A final filter section just before the point of echo cancellation adjusts short-line performance within acceptable limits.

An all-pass filter realized as a first-order switched-capacitor section clocked at 8fb has a pole–zero location of 0.8fb. Its primary purpose is to adjust the horizontal eye opening for short-line lengths. Fig. 5 shows the horizontal opening as a function of line length at the input and the output of the all pass. The effect for long lines is small while performance at 0 km is greatly improved. The all-pass filter is also the point of echo cancellation, where the analog echo estimate from the D/A converter is subtracted from the filtered composite line signal. Cancellation is

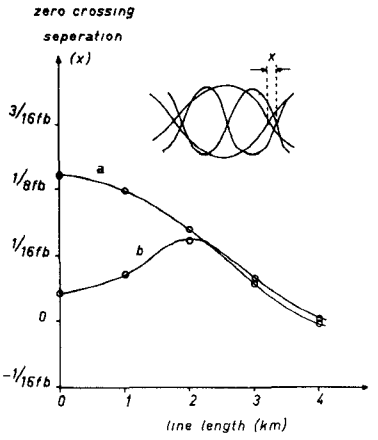

Fig. 5. Horizontal eye closure as a function of line length (24 gauge) for 160 kb/s transmission. (a) Without all pass section. (b) With all pass section.

performed after the receive filter so that the echo canceller can take advantage of the shortened impulse response, keeping the amount of TX data history required to a minimum. After removal of the near-end component, a decision is made as to whether a zero or a one has been received. Each analog sample at the output of the all-pass filter is compared to the analog ground reference with a high-speed comparator and the resulting digital information is made available to a digital decoder. This circuit performs an integrate and dump on 6 of the 8 samples in each baud to achieve an enhanced signal-to-noise ratio, (SNR). The two samples at the mid-baud zero crossing and the band boundary are ignored because the horizontal eye opening is uncertain in these regions.

Either the on-chip crystal oscillator or an external 10.24-MHz clock driving the crystal pin provides the transceiver with a reference timebase. In CO mode the PCM-highway and transmit path timebases are directly derived from this clock. The opposite SET device uses a digital phase locked loop to create transmit and receive timebases locked to the incoming biphase signal. In CO mode the phase locked loop output is used only for the digital part of the receive leg, while the local crystal oscillator supplies clocks to the sampled analog filters. Here, the DPLL accounts only for the varying time delay of the far-end signal over the range of line lengths, since the received signal originating from the SET is locked in frequency to the CO timebase. The digital phase locked loop applies 49 ns corrections to the local crystal oscillator frequency at a rate of 5 or 2.5 kHz for 160 and 80 kb/s, transmission, respectively. Required crystal tolerance is either 60 or 120 ppm.

IV. Echo Cancellation

The echo canceller is based on the memory-compensation principle using a sign algorithm adaptation process. Traditionally echo cancellers have used an adaptive filter, typically an FIR structure, which converges to the impulse response of the echo path. The filter performs a convolution of the transmitted data stream with the estimated impulse response to provide echo estimates which are subtracted from the composite signal. The memory compensation technique circumvents the convolution process by storing echo estimates directly in a RAM which is addressed by the transmitted data stream. A major advantage of the memory compensation scheme is a simple algorithm and an associated reduction in hardware complexity. While a relatively large RAM is needed to store the echo responses of every possible transmit data sequence, the technique need not rely on the linearity of the channel to achieve perfect cancellation. The memory compensation principle provides enhanced performance over the FIR filter approach, by accounting for nonlinear distortion in the echo path.

The sign algorithm adaptation process is chosen for its simplicity and ease of implementation. This method uses the sign of the received signal after echo cancellation to increment or decrement the corresponding echo estimate, as opposed to the stochastic iteration algorithm [8] which adds or subtracts an amount proportional to the magnitude of the received signal. Instead of an A/D converter to digitize the echo cancelled signal, only a comparator is required. The sign algorithm represents the most extreme quantization of the error signal which may result in inaccurate cancellation in the presence of a far-end signal. Extra noise, in the form of a dither signal, is added to the signal before taking the sign, to make the average quantization smoother for accurate echo-cancellation. Without dithering, an echo estimate can experience wander about the correct value, equal in magnitude to the far-end signal, which may prevent convergence completely.

The major disadvantage of the sign algorithm is a relatively long convergence time, particularly when coupled with the memory compensation architecture which is slower than the adaptive filter approach. This can be overcome by an accelerated convergence option which modifies words in RAM by more than 1 LSB, or by providing continuous communication even during on-hook periods.

The echo canceller in the transceiver chip has been designed to accommodate 40 dB of cable attenuation. If we assume a required signal to noise ratio of 15 dB for a bit error rate of 10^{-7} and an analog hybrid attenuation of 10 dB, then we require 45 dB of echo cancellation. The echo canceller achieves this level of cancellation using 14-bit words in the RAM, 5-bit history with 8 samples per baud, a 9-bit D/A converter and an adaptive dither source that tracks the amplitude of the received signal. Holte and Stueflotten [4] have analyzed the noise performance of the sign algorithm memory compensation scheme. The major sources of internal noise are cancellation error, quantization noise, residual echoes and jitter noise. Steady-state cancellation error has been derived by considering the adaptation of a single memory register and is a function of the finite resolution in the register and the dither amplitude. The SNR due to cancellation error with dither amplitude 1.4 times the received signal amplitude and the 13-bit plus sign code equal in magnitude to the echo after 10 dB of precancellation is 25.6 dB [4, eq. (20)]. It should be noted that as an approximation, the received signal is con-

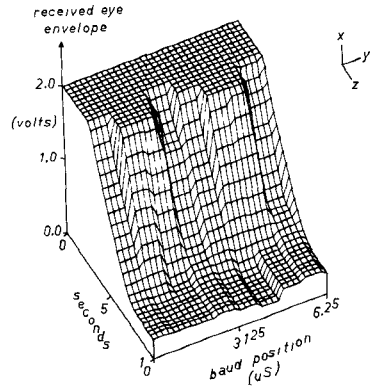

Fig. 6. Simulated echo canceller convergence in the presence of a far-end signal.

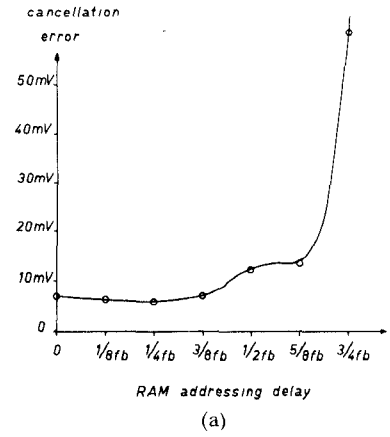

(a)

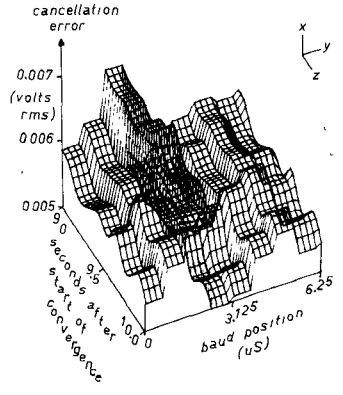

(b)

Fig. 7 (a) Worst-case rms cancellation error versus RAM addressing delay (b) rms cancellation error with RAM addressing delay = 3/8 f_b.

sidered to be a binary signal free of intersymbol interference.

Although a 14-bit echo estimate is stored in RAM, only the nine most significant bits are sent to the D/A converter. This makes the effective D/A step size 32 times the RAM step size. The effect upon the adaptation may be neglected and the noise simply treated as quantization noise. For 40 dB of cable attenuation this results in a signal-to-quantization noise ratio of 28.9 dB [4, eq. (21)].

Another important source of noise in the echo canceller is that of residual echo. The echo path's impulse response has been assumed to be finite but it extends beyond the 5 bauds chosen to address the memory so that earlier pulses will contribute a noise in the form of an uncancelled residual echo. Since the impulse response is a strong function of the line type and the presence of bridge taps or gauge changes, it becomes difficult to predict the level of noise due to residual echo. Simulations and laboratory work with a variety of lines have shown that 5 bauds of echo history provide satisfactory performance.

The necessity of deriving a clock from the data creates further noise due to jitter. The transceiver uses a simple digital phase locked loop with a master clock of 10.24 MHz to provide discrete adjustments of approximatly 49 ns. It has been found that the noise due to jitter is at least 20 dB below the received signal.

Simulations provide a means of verifying an echo-canceller design based on Holte and Stueflotten's analysis. Fig. 6 shows a simulation of the dynamics of echo canceller operation in the presence of a far-end signal. Planes parallel to the $x-y$ axis hold the instantaneous peak envelope of the received eye as it would be seen on an oscilloscope. The three-dimensional surface represents this peak eye envelope as the echo canceller converges. At time zero the signal is completely saturated at the opamp clipping level. Then the eye envelope slopes towards zero with the profile of the as yet uncancelled near end signal in evidence. Finally, after about 7 s, the curve levels off in the form of the highly attenuated far-end signal. Convergence can be speeded up by a factor of 32 with the externally programmed accelerated convergence feature.

The echo canceller cannot account for echo unrelated to the RAM address, which consists of the most recently transmitted 5 bits of data. In order to extend echo coverage as far as possible with the 5-bit prehistory scheme, the address to the RAM was delayed slightly with respect to the TX filter input to account for the delay inherent in the transmit and receive filters. It was determined that an addressing delay of 3/8 baud was optimum. Fig. 7(a) shows RMS cancellation error for 0 km line as a function of this delay. The curve starts to increase after 3/8 baud delay. The rms cancellation error as a function of baud position over time is shown in Fig. 7(b). One baud position shows a greater error which represents a small systematically uncancelled portion of the near-end signal which sneaks through before the RAM is even addressed. The relatively constant noise floor is due to short term wander of RAM coefficients, clock jitter and D/A converter truncation error. Considering the 15-dB receive path gain and 40 dB of line attenuation, the receiver SNR with 7 mV rms noise is better than 15 dB, which is close to what the theoretical analysis leads us to expect.

The D/A converter is an 8-bit plus sign binary weighted switched capacitor structure which uses an on chip reference and a 0.05-pF unit capacitor for a step size of 7 mV. Linearity is not so important in this application as monotonicity and low current consumption. The most significant 9 bits of the 14-bit wide RAM data bus are supplied to the D/A converter at a rate of 8 f_b to cancel the near end component on each sampled output of the receive filter.

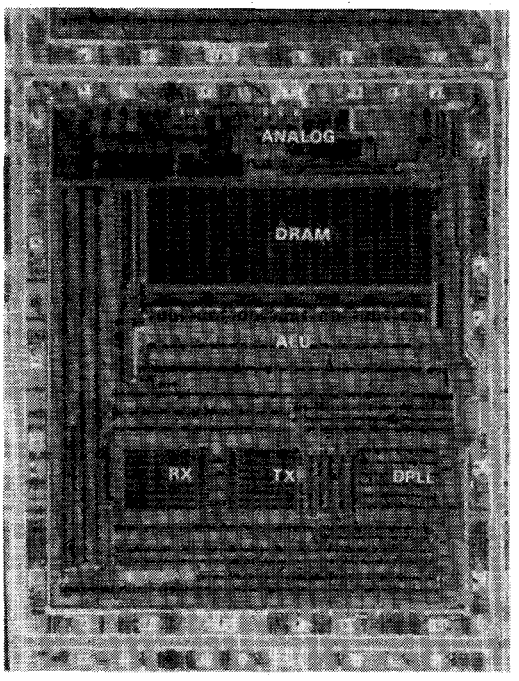

Fig. 8. Chip photomicrograph.

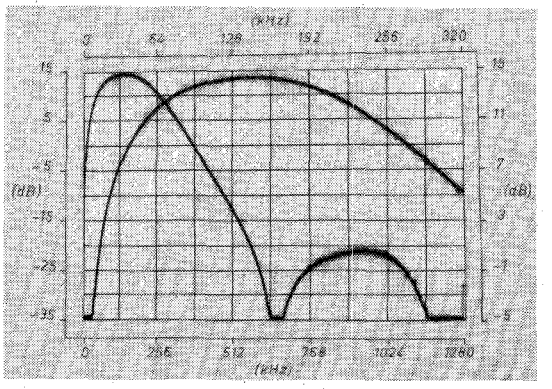

Fig. 9. Receive filter response at 160 kb/s.

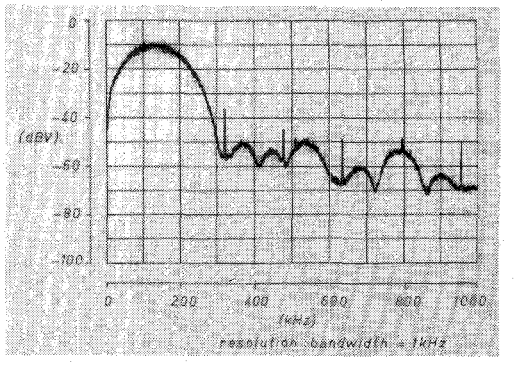

Fig. 10. Transmit signal spectrum at 160 kb/s.

The echo canceller stores estimates of the line response in a 3584-bit CMOS dynamic RAM, which undergoes a READ–MODIFY–WRITE cycle at 8 f_b. A folded bit line architecture is used with each bit stored differentially in complementary memory cells. The refresh is interleaved with normal operation, so that all 32 rows are completely refreshed every 4 bauds. The 8-bit address consists of 5 bits of TX data history plus 3 bits indicating 8 positions within each baud. On each cycle an echo estimate is altered by one LSB depending on the result of a sign comparison between the RX signal and a dither signal having the desired statistical properties for uniform convergence.

The dither circuit generates a triangular waveform centered on analog ground with an amplitude related to the peak received far-end signal. It consists of a switched-capacitor peak detector clocked at the baud rate followed by a passive switched-capacitor ramp generator with an output frequency of 8 $f_b/129$ and a peak amplitude 1.4 times the peak detector output up to a maximum of 1 V. The frequency was chosen so as not to correlate with scrambled data from either end but high enough not to cause significant wander in the RAM coeffecients. On short lines the received signal will sometimes exceed the dither amplitude and will experience wander in the area outside the 1-V peak-to-peak range. This will not affect reliable data recovery.

V. Chip Performance

Fig. 8 shows a photomicrograph of the integrated circuit with the major circuit blocks identified. The die which occupies 27.7 mm^2 (192×224 mils) of silicon is packaged in a 22-pin DIP with 0.4BSC spacing.

The receive filter frequency response and the transmit filter output spectrum are displayed in Figs 9 and 10, respectively. Eye diagrams of the transmit, receive, and composite line signals for 160 kb/s transmission over 2- and 3.2-km 24 gauge lines are shown in Fig. 11. A small far-end component can be seen superimposed on the large near-end component for the 2-km composite line signal while at 3.2 km the far-end is attenuated to the point where it not easily visible on the oscilloscope. The received signal at both line lengths showing no transmit component demonstrates the functionality of the echo canceller. The effect of the fixed equalizer at 2 km is to cause multiple zero crossings which reduces the horizontal eye opening. At 3.2 km the line brings these zero crossings closer together, for improved opening.

Fig. 12 plots the signal-to-noise ratio for a fixed bit error rate of 10^{-6} over a range of 24 gauge plastic insulated cable lengths, at both 80 and 160 kb/s. Low-pass filtered white noise, having a bandwidth of the baud-rate, is injected onto the line at the receive transformer. The SNR is defined as the ratio of far-end signal to total noise, measured at the line input of the transceiver.

For zero line length the SNR is relatively high as a result of the fixed equalization optimized for long lines, and the dither signal which falls short of the far-end signal amplitude. The 160-kb/s curve approaches a minimum of 12 dB at medium length lines and becomes asymptotic after 4 km, indicating that transmission is impossible after this point. At 80 kb/s, SNR performance is marginally worse than that at 160 kb/s. This is primarily due to the line termination components which were optimized for 160-kb/s transmission. The 80-kb/s curve has not yet started to show the

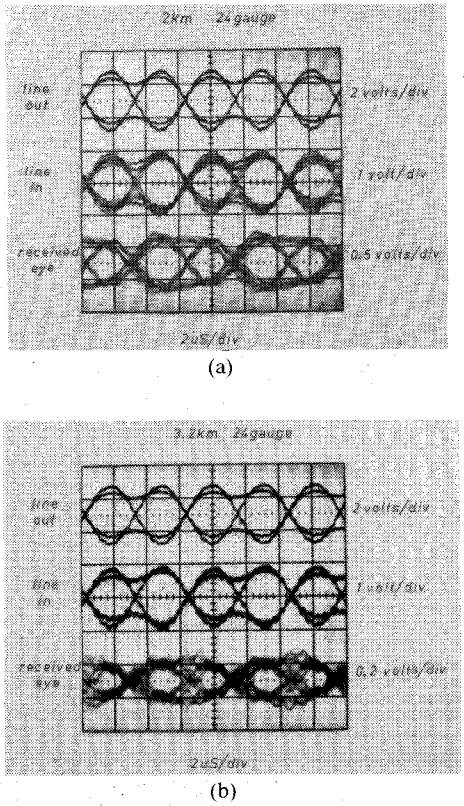

Fig. 11. Transceiver signals. (a) 2-km line, (b) 3.2-km line.

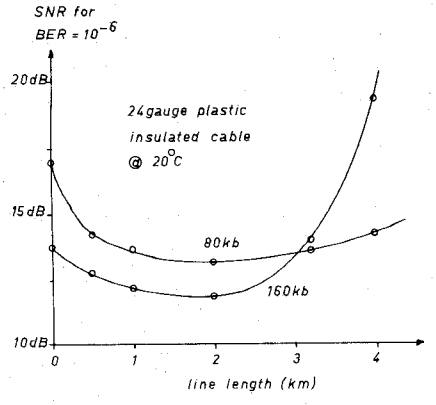

Fig. 12. SNR versus line length.

effects of line attenuation even at 4 km. Reliable transmission on up to 5 km is possible at this data rate before the SNR curve deteriorates.

The signal to noise performance was also measured as a function of bridge tap placement (Fig. 13). A 500-m open-circuit 26 gauge bridge tap was placed along a 3.2-km length of 24 gauge twisted pair cable between the SET and CO devices. When the bridge tap is positioned right at the RX position, where BER is measured, required SNR is high. This is due to gross line mismatch and the resulting clipping within the precanceller. As bridge tap distance from the receiving end increases, the SNR peaks and falls off once again. This phenomena can only be the result of imperfect echo cancellation after the 5-baud echo history, since ISI would be the same no matter where the bridge tap

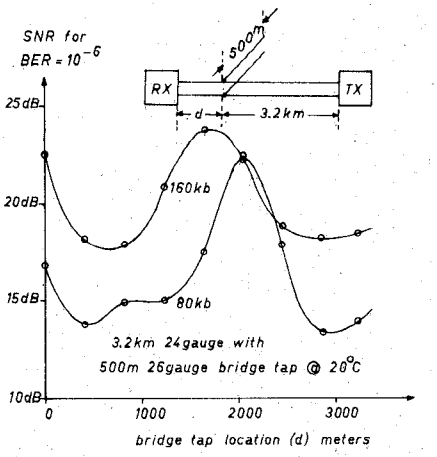

Fig. 13. SNR versus bridge tap location.

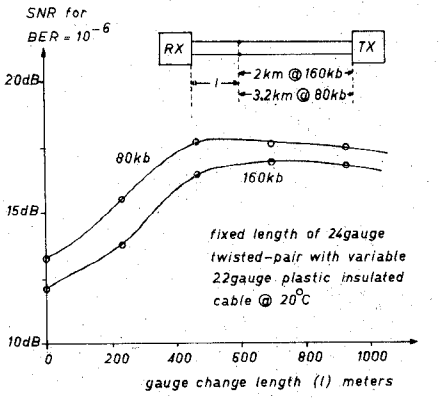

Fig. 14. SNR versus gauge change location.

was located. The peak occurs at the point where the combined effects of round trip echo delay and attenuation to the bridge tap create the largest echo unrelated to the 5-bit history. At no point does the bridge tap curve approach that of the straight line performance. This is due mainly to the ISI introduced into the far-end signal.

A short length of 22 gauge quad installation cable was connected in series with a 2- or 3.2-km 24 gauge twisted pair to examine the effects of gauge changes on performance (Fig. 14). Different lengths of 24 gauge cable were used for the 80- and 160-kb/s measurements so that line attenuation would be similar for the two baud rates. Both curves rise and then level off as a result of two different effects. Increasing the composite line length results in a continually rising curve. Reflections from the gauge change mismatch will cause a peaking curve, so that the resultant SNR levels off after 500 m.

Several key features of the IC provide enhanced performance without penalties in power or area. Biphase line-code tailored to the limited channel bandwidth permits the use of a simple fixed equalizer. RAM-based echo cancellation allows for simple line interfacing and is capable of handling a variety of line impedances, bridged taps, gauge changes, and nonlinearities. Symmetry between SET and CO devices, together with complete control of transceiver functions through external programming, makes the part very flexible.

This chip combines small size, low-power consumption, data transparency, and a high degree of integration in accomplishing reliable 80- or 160-kb/s full duplex data transmission for the PABX environment or the outside plant.

Acknowledgment

The authors wish to thank A/S Elektrisk Bureau of Norway, G. Aasen, P. Beirne, K. Buttle, R. Gervais, G. Reesor, D. Ribner, H. Schafer, and R. White for their substantial technical contribution as well as M. Foster and D. Brown for their support and encouragement in this project.

References

[1] O. Agazzi, D. Hodges, and D. Messerschmitt, "Large scale integration of hybrid-method digital subscriber loops," *IEEE Trans. Commun.*, vol. COM-30, pp. 2074–2082, Sept. 1982.

[2] A. Komori, M. Furukawa, T. Sato, T. Komazaki, "A 200Kb/s burst mode transceiver with two-bridge tap equalizer," in *ISSCC Dig. Tech. Papers*, pp. 236–237, Feb. 1984.

[3] Y. Hino, T. Chujo, N. Ueno, K. Fujita, M. Yamamoto, K. Yamaguchi, and H. Kikuchi, "A burst-mode LSI equalizer with analog building blocks," in *ISSCC Dig. Tech. Papers*, pp. 238–239, Feb. 1984.

[4] N. Holte and S. Stueflotten, "A new digital echo canceller for two-wire subscriber lines," *IEEE Trans. Commun.*, vol. COM-29, pp. 1573–1581 Nov. 1981.

[5] U.S. Patent 4237463.

[6] K. Buttle, G. Aasen, R. Colbeck, R. Gervais, P. Gillingham, D. Ribner, H. Schaefer, and R. White, "A 160kb/s full duplex echo cancelling transceiver," in *ISSCC Dig. Tech. Papers*, pp. 152–153, Feb. 1985.

[7] T. Choi, R. Kaneshiro, R. Broderson, P. Gray, W. Jett, and M. Wilcox, "High-frequency CMOS switched-capacitor filters for communications applications," *IEEE J. Solid-State Circuits*, vol. SC-18, pp. 652–664; Dec. 1983.

[8] N. Verhoeckx, H. van den Elzen, F. Snijders, P. van Gerwen, "Digital echo cancellation for baseband data transmission" *IEEE Trans. Acoust. Speech, Signal Processing*, vol. ASSP-27, pp. 768–781, Dec. 1979.

✠

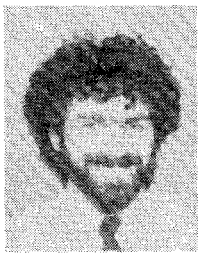

Roger P. Colbeck was born in Sussex, England. He moved to Canada in 1970 and received the B.Eng. degree in electrical engineering, in 1980 from Carleton University, Ottawa, Ont., Canada.

From 1980 to present he has worked at Mitel Corporation, Ottawa, as a designer of analog/digital CMOS telecommunication integrated circuits which included switched-capacitor filters, codecs and modem chips. He is now Project Manager for Data Acquisition and Digital Signal Processing Products. Since 1983 he has also been a graduate student at Carleton University where he is completing a thesis in high speed A/D techniques.

✠

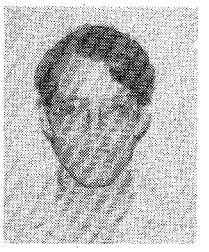

Peter B. Gillingham (S'79–M'82) received the B. Eng. degree in electrical engineering from Carleton University in 1980. He spent the following year as a research assistant at L'Ecole Polytechnique Federale in Lausanne, Switzerland, working in switched-capacitor filters. Returning to Carleton he completed his master's thesis entitled "Switched Capacitor Circuits for VLSI: A Design Study" for which he received the M. Eng. degree in 1983. Since then he has been employed at Mitel Corporation, involved in analog CMOS integrated circuit design and the development of an ISDN family of components.