A 125-MHz Mixed-Signal Echo Canceller for Gigabit Ethernet on Copper Wire

Tai-Cheng Lee and Behzad Razavi, *Member, IEEE*

***Abstract*—A discrete-time analog echo canceller is described that reduces the echo in the front end of Gigabit Ethernet twisted-pair interfaces. Echo cancellation in the analog domain by means of four taps reduces the complexity of the digital echo canceller and crosstalk cancellers. Designed in a 0.4-μm CMOS technology, the circuit employs an LMS algorithm to adapt to the cable length and impedance discontinuities, providing an echo suppression of 10 dB. The design operates at 125 MHz while consuming 43 mW from a 3-V supply.**

***Index Terms*—D/A converters, echo cancellers, Gigabit Ethernet, wireline communication.**

I. Introduction

GIGABIT Ethernet on copper cable is a fast evolving technology enabling 1 Gb/s full-duplex data communication over the existing UTP CAT-5 twisted-pair cables. As depicted in Fig. 1, four pairs of twisted-pair cables and eight transceivers (four at each end) with 250 Mb/s data rate offer 1-Gb/s data communication. Five-level pulse amplitude modulation (PAM-5) reduces the baud rate to 125 MHz.

In full-duplex communication, a hybrid is typically used to isolate the receiver from the transmitted signal. However, cable and connector impedance variation still results in substantial leakage of the large transmitted signal, thereby creating near-end echo. Also, impedance discontinuities along and at the end of the cable produce far-end echo. In addition, the near-end crosstalk (NEXT) between the four cables is significant.

The full digital solution, suggested by the IEEE 802.3ab task force [1], requires at least a 7-bit analog-to-digital converter (ADC) to quantize the attenuated signal as well as the echo and the crosstalk. (In practice, a resolution of 8 bits may be necessary to ensure enough design margin.) Moreover, this solution employs four very long digital adaptive finite-impulse-response (FIR) filters, including one echo canceller and three near-end-crosstalk cancellers, in each transceiver.

This paper proposes a mixed-signal echo canceller to partially cancel the echo in the analog front end, relaxing the complexity in the digital domain. Using a least-mean-squared (LMS) algorithm, the circuit adapts the four taps of a discrete-time analog FIR filter at the startup, thus cancelling the four largest echo signals between two transceivers. Fabricated in a 0.4-μm digital CMOS technology, the circuit reduces the echo power by 10 dB at 125 MHz while consuming 43 mW from a 3-V power supply.

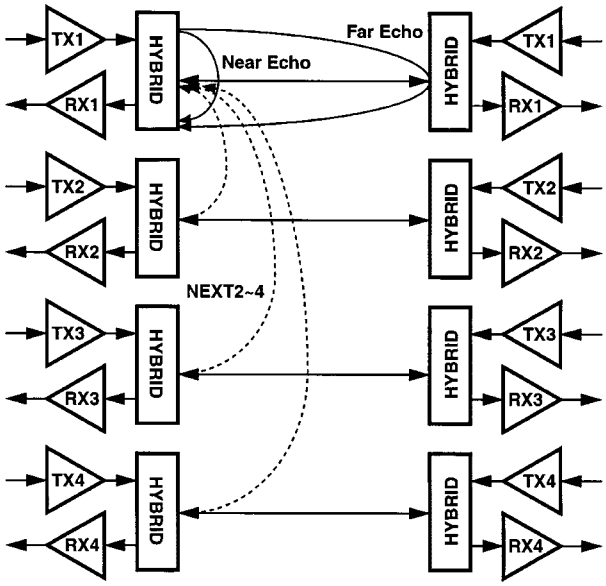

Fig. 1. Gigabit Ethernet over four twisted pairs.

The next section of the paper provides the background for this work, addressing general design considerations. Section III presents the echo canceller architecture and Section IV describes the design of the building blocks. Section IV summarizes the experimental results.

II. Design Considerations

The simultaneous transmission and reception of signals in Gigabit Ethernet requires a hybrid circuit with tightly-controlled parameters. As shown in Fig. 2, a hybrid transformer consists of six windings heavily coupled to one another. The transmitted signal is coupled to W_2 and W_4 and the received signal is picked up from W_3 and W_5.

If $Z_b = Z_0$ and the windings match perfectly, then W_2 and W_4 couple the same amount of the transmitted signal to W_3 and W_5 but with opposite polarity. As a result, W_6 picks up only the received component from W_5.

In practice, mismatches within the hybrid transformer and impedance discontinuities at various points in the transmission medium result in substantial near-end and far-end echo. The echo is quantified by applying an impulse to the transmit port of the hybrid and measuring the response at the receive port. The task of echo cancellation consists of generating another component that, when added to the received signal, cancels the echo.

Manuscript received July 20, 2000; revised October 18, 2000. This work was supported in part by Altima Communications, San Jose, CA.

The authors are with the Electrical Engineering Department, University of California, Los Angeles, CA 90095 USA (e-mail: razavi@ee.ucla.edu).

Publisher Item Identifier S 0018-9200(01)01471-8.

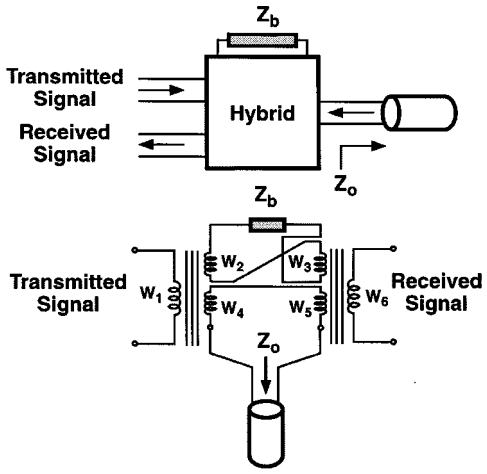

Fig. 2. Hybrid transformer.

Fig. 3 shows a generic Gigabit Ethernet system. The 250-Mb/s binary data stream is first converted to a 125-Mbaud five-level signal. The result is then applied to a partial-response filter that adds 25% of the previous symbol to 75% of the present symbol, thereby reducing the electromagnetic radiation by the twisted pair. Subsequently, a digital-to-analog converter (DAC) drives the hybrid and the cable. In the receive path, an analog interface (e.g., a variable-gain amplifier) couples the signal to the ADC, and the result undergoes digital processing. In this work, an analog echo canceller is designed that partially suppresses the echo before digitization.

Fig. 4(a) illustrates a typical impulse response. While the response contains only a few large reflections, digital echo cancellers commonly require a long FIR filter because the position of the reflections along the cable is not known.

The echo canceller reported here incorporates a "significant echo locator" to identify the temporal location of the four largest reflections in Fig. 4(a). The circuit then applies a sign–sign LMS algorithm [2] to generate an emulated echo and subtracts it from the signal–echo combination, producing a net response as shown in Fig. 4(b).

System simulations indicate that a four-tap canceller reduces the echo by 10–15 dB. Since the near-end crosstalk from other three cables is only 16 dB below the uncorrected echo [1], greater echo cancellation in the analog domain does not improve the performance any more.

System simulations also suggest that the echo canceller reported here can reduce the complexity of the digital echo canceller and crosstalk cancellers. Employing a few more taps, the mixed-signal canceller can even relax the resolution of the ADC and hence the digital processing by one bit.

III. Echo Canceller Architecture

Fig. 5 shows the detailed architecture of the echo canceller. The transmitter consists of a 64×1 b first-in-first-out (FIFO) register and a DAC, which drives the network R_1–R_3 and the transformer. The receiver employs a differential voltage-to-current converter (VIC_1), and the calibration circuit comprises a reflection locator circuit, an LMS machine, and an analog FIR filter. The voltages at nodes X and Y are scaled replicas of the transmitted signal and the transmitted+received signal, respectively. Thus, in the absence of impedance variations and discontinuities, the output of VIC_1 contains only the received signal.

The system operates as follows. At the startup, the transmitter applies a unit step (rather than an impulse) to the hybrid and the cable, generating the uncorrected echo at nodes X and Y (and at the output of VIC_1). The reflection locator circuit then differentiates the result by means of a switched-capacitor circuit, obtaining the impulse response. If the response exhibits an amplitude greater than V_{REF}, the address of the corresponding tap in the FIFO register is stored for actual operation. The value of V_{REF} can be roughly estimated by the following equation:

$$V_{\text{REF}} < V_{\text{transmit}} \times \Gamma \times (\text{Gain of Differentiator}) \quad (1)$$

where $V_{\text{transmit}} = 2$ V, Γ is the reflection coefficient (less than $\pm 15\%$ according to maximum impedance mismatch tolerance), and the gain of the differentiator is around unity. Based on this estimation as well as simulation results, V_{REF} can be set in the vicinity of 50 to 100 mV.

This test is repeated 64 times, identifying four significant taps. Testing the first 64 possible significant echo reflection points can cover the 100-m cable length. For any reflections beyond this distance, the echo becomes less important because of the cable's attenuation. (If the tap locator circuit identifies fewer than four significant taps, the rest of the significant tap addresses are set to the first few taps because their corresponding reflections experience less attenuation.)

Next, the analog FIR filter is adapted so as to emulate the echo. Driven by the four significant taps, the filter produces a current output that is subtracted from the output of VIC_1 and compared to zero. Based on the result, the LMS machine updates the multiplier coefficients m_1–m_4 in the FIR filter, forcing the echo residue to small values. Each multiplier is in fact realized as a DAC whose reference is modulated by the data bits provided by the selection logic. The analog signal path is fully differential in both transmit and receive paths.

The conventional LMS algorithm [2] is based on the following equation:

$$c[n+1] = c[n] + \mu \times err[n] \times data[n] \quad (2)$$

stating that each coefficient is incremented by the product of the error, the data, and the factor μ. This adaptation scheme demands analog multipliers to update coefficients, making the implementation difficult. The sign–sign LMS algorithm, on the other hand, requires only the sign of the error and the sign of data for adaptation [2]:

$$c[n+1] = c[n] + \mu \times \text{sgn}(err[n]) \times \text{sgn}(data[n]) \quad (3)$$

The LMS machine therefore incorporates up/down counters to store the coefficients of the analog multiplier. The counters count up or down according to the product of the error sign, *sgn(err)*, and the data sign, *sgn(data)*.

The performance of the architecture was first analyzed by a behavioral model written in C, including the characteristics of the cable and the sources of echo and crosstalk. Three critical parameters were determined by this analysis.

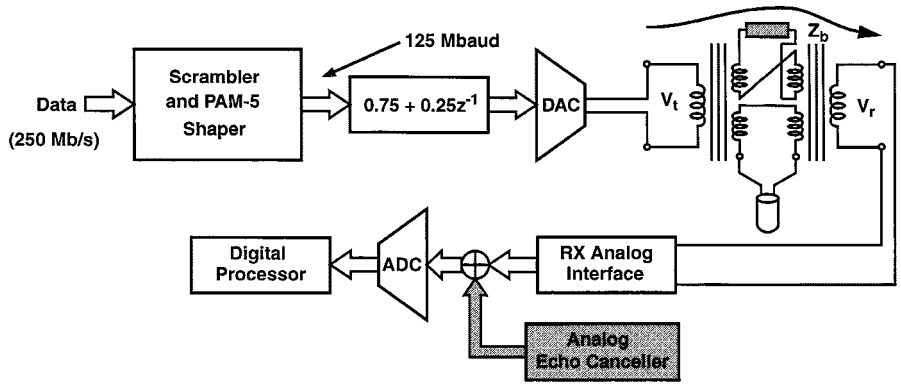

Fig. 3. Block diagram of Gigabit Ethernet transceiver.

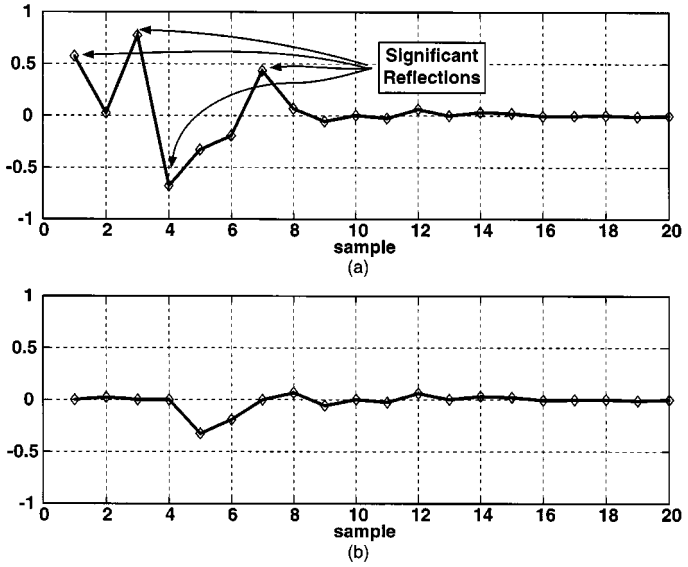

Fig. 4. (a) Typical echo impulse response. (b) Reshaped echo impulse response.

A. Minimum Number of Taps

With sufficiently high precision in the FIR filter coefficients, the echo cancellation becomes proportional to the number of the taps. The simplicity of the sign–sign LMS algorithm further encourages increasing the number of the taps at little cost.

However, the upper bound of the number of the taps is determined by the power dissipation of the analog filter and the adjacent channel crosstalk. In this design, the echo is suppressed in the analog domain to become comparable to the crosstalk. Fig. 6 plots the residual echo as a function of the number of the taps in the filter, indicating that four taps can provide about 14 dB of echo cancellation.

B. Coefficient Accuracy Requirements

The weights associated with the taps in the echo canceller are stored in digital format. To represent the quantization error due to the finite precision of the coefficients in the FIR filter, it can be assumed that echo coefficient has infinite precision but adds some quantization noise. The quantization noise in turn limits the extent to which the echo can be suppressed. Since the required precision of the coefficients determines the resolution of the DACs in Fig. 5, it is important to quantify the tradeoffs between complexity, power dissipation, and the residual echo.

For a given number of taps, the following equation describes the overall noise power (including quantization noise and residual echo):

$$\mathrm{Noise}_{\mathrm{residue}} = \mathrm{Err}_{Q(\mathrm{weights})} + \mathrm{Echo}_{\mathrm{residue}}|_{\mathrm{infinite\ precision}} \tag{4}$$

where $\mathrm{Echo}_{\mathrm{residue}}$ can be obtained by simulations. The quantity $\mathrm{Err}_{Q(\mathrm{weights})}$ is the quantization error resulting from the finite precision of the coefficients and can be obtained by the following derivation:

$$\mathrm{Err}_{Q(\mathrm{weights})} = E_s \cdot H_{\mathrm{echo}} - E_s \cdot \overline{H_{\mathrm{echo}}} \tag{5}$$

where E_s is the power of the input sequence, H_{echo} is the power of the impulse response of the significant echo, and $\overline{H_{\mathrm{echo}}}$ is the power of the quantized impulse response of the significant echo. The power of the quantized error in the coefficient is equal to $(H_{\mathrm{echo}} - 6.02N)$ in decibels, where N denotes the precision of the coefficients. Therefore

$$\mathrm{Err}_{Q(\mathrm{weights})} \approx 10^{(-N \times 6.02\,\mathrm{dB}/10)} \cdot H_{\mathrm{echo}} \cdot E_s \tag{6}$$

The required value of N is first estimated by (6) and then determined by simulations. The loop gain is reduced by using more accurate coefficient bits and results in better echo reduction. However, for high values of N, the echo reduction is limited by the number of the taps in the FIR filter. Fig. 7 plots the echo suppression as a function of the resolution of the multiplier coefficients (with four taps in the filter), suggesting negligible improvement for resolutions greater than nine bits.

Though the digital implementation indicates that a 7–8-bit ADC is sufficient for the signal path, a higher resolution is required for the coefficient path because the mixed-signal echo canceller must adapt in the presence of the echo residue and near-end crosstalk. In fact, the sign–sign LMS algorithm fails to adapt correctly if the coefficient resolution is inadequate. Therefore, we have chosen ten bits (one bit for sign).

C. Maximum Tolerable DAC Nonlinearity

The integral nonlinearity (INL) of the DACs in Fig. 5 impacts the accuracy of the emulated echo, thereby limiting the echo

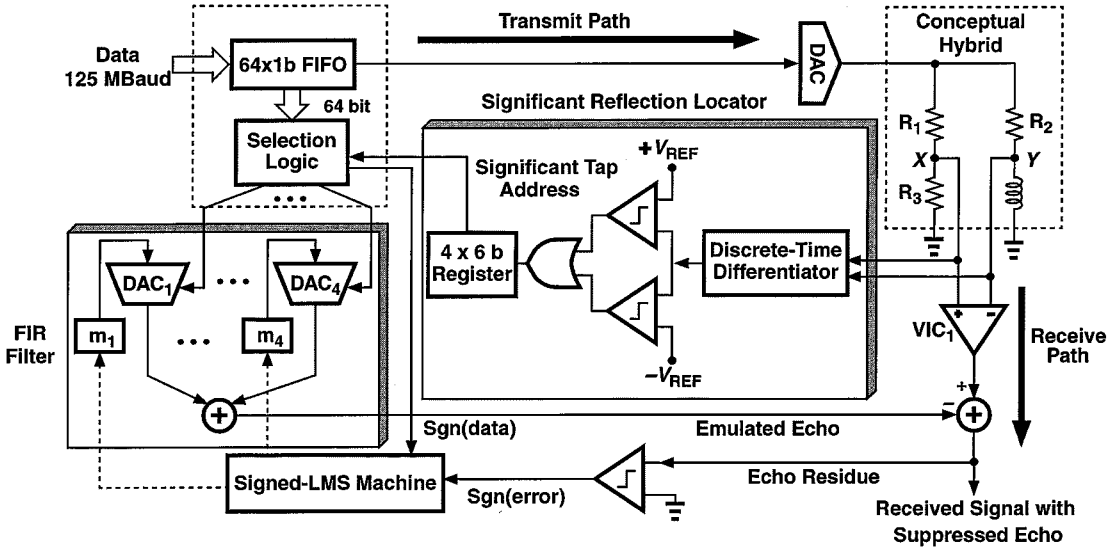

Fig. 5. Echo canceller architecture.

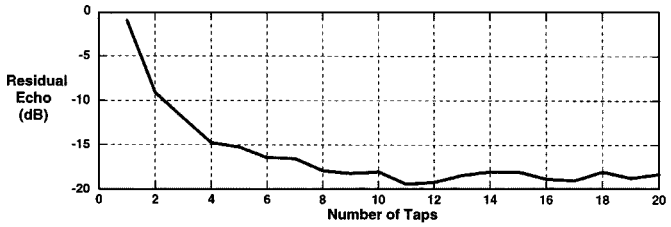

Fig. 6. Performance as a function of number of taps.

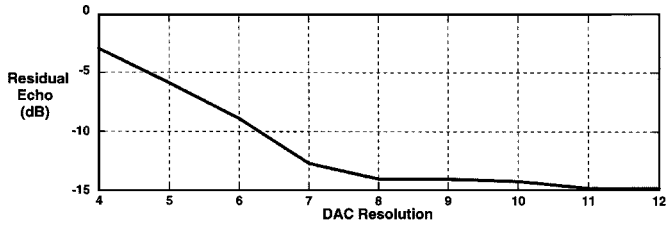

Fig. 7. Performance as a function of DAC resolution.

cancellation. In other words, when the LMS machine updates the digital values of $m_1, \ldots, m_4$, the corresponding analog values suffer from the DAC INL, producing an emulated echo different from the ideal value. Fortunately, so long as the DAC remains monotonic, the LMS machine continues to update the coefficients to reach the optimum point. The tradeoffs between speed, power dissipation, and linearity of DACs make it necessary to quantify the effect of the DAC INL on the residual echo.

Fig. 8 plots the simulated echo cancellation as a function of the DAC nonlinearity (modeled by a third-order polynomial). Here, four taps and a DAC resolution of 10 bits are assumed. These results suggest that a linearity of about 7 or 8 bits proves adequate, greatly simplifying the design of the DACs.

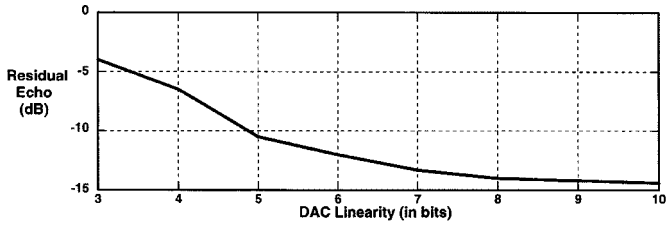

Fig. 8. Performance as a function of DAC linearity.

IV. Building Blocks

A. Tap Locator

Shown in Fig. 9, the tap locator consists of a discrete-time differentiator and peripheral logic. The circuit samples the received data from the cable and produces the impulse response. A threshold detector determines if the impulse response exceeds the threshold, and if the signal indeed falls outside the window, the tap location corresponding to the large reflection is recorded for adaptation and actual operation.

The addresses of the significant taps are recorded as follows. The 6-bit binary counter is turned on when the circuit is initiated while the tap locator circuit samples the uncorrected echo signal. When the threshold detector in the tap locator produces an enable pulse, the register file copies the content of the counter into the register.

The outputs of the register file are used to control the 64-4 MUX to allow the data bits corresponding to the significant reflections to be applied to the mixed-signal FIR filter.

Note that the actual transmitted data need not pass through the FIFO register to drive the transmitter. Thus, the echo canceller introduces no latency in the transmit path.

B. Differentiator

The discrete-time differentiator used in the reflection locator of Fig. 9 must sample the data at 125 Mb/s and subtract it from the previous data. Fig. 10 depicts a compact high-speed realization of this function that requires only two clock phases. In the sampling phase [Fig. 10(a)], A_1 and A_2 are reset while storing

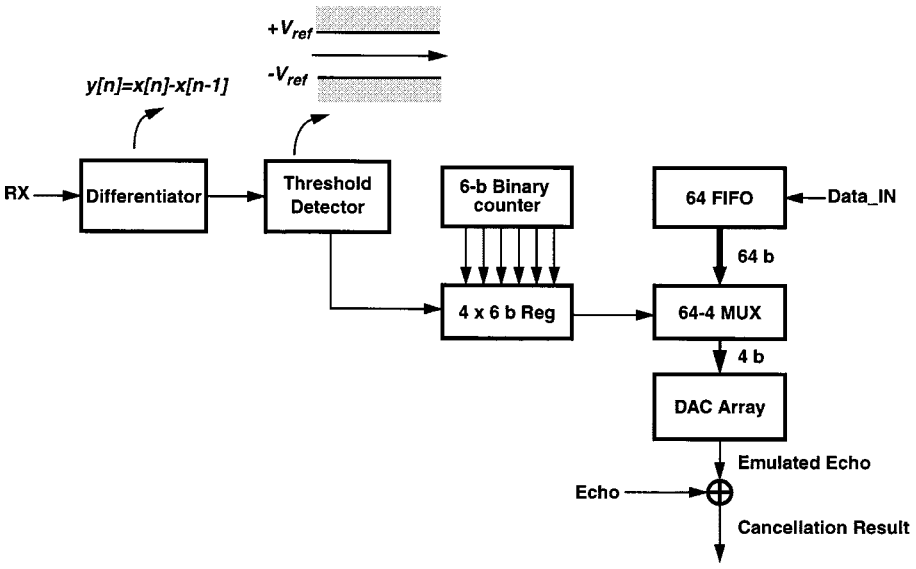

Fig. 9. Tap locator circuits.

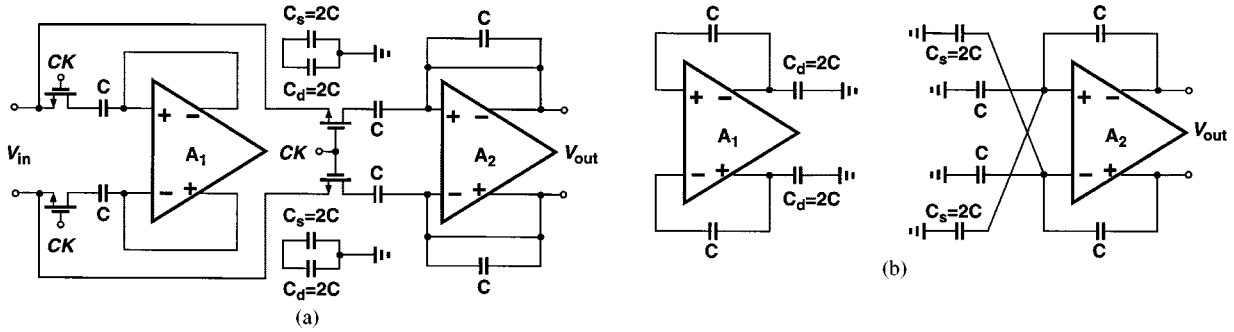

Fig. 10. Differentiator. (a) Sampling mode. (b) Evaluation mode.

$V_{\rm in}$ on their input capacitors. At the same time, the discharged capacitor C_s samples the previous data stored on C_d (while dividing the amplitude by a factor of two). In the evaluation phase [Fig. 10(b)], A_1 holds the sampled data, charging C_d, while A_2 subtracts the sampled data from the value stored on C_s. Note that the virtual ground at the input of A_2 forces the voltage across C_s to zero, thereby preparing it for charge sharing with C_d in the next phase.

In addition to speed, the linearity of the circuit is also important. In particular, the charge sharing between C_s and C_d in the sampling mode is susceptible to the junction capacitance nonlinearity of MOS switches, but differential operation and proper sizing of transistors minimize this effect. Each op amp is implemented by a folded-cascode topology for the reset operation. The capacitors are realized as metal–metal sandwiches. Note that the differentiator is turned off after the significant reflections are identified, saving power dissipation.

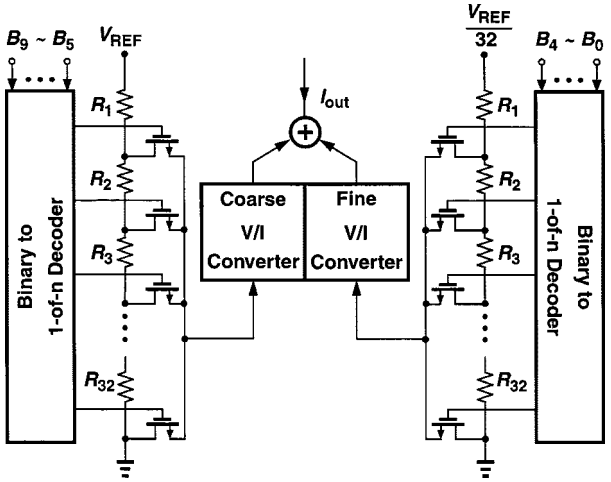

Fig. 11. 10-bit DAC architecture.

C. Analog Multiplier

System simulations indicate that the multiplier coefficients in the mixed-signal FIR filter must have a resolution of ten bits so as to maintain the LMS loop stability and achieve acceptable echo suppression. It is also ascertained that the nonlinearity in the coefficients need not be less than 0.8% (approximately seven bits). Thus, the DACs performing the multiplication must provide 1024 monotonic steps with moderate INL. Since the FIR filter incorporates four such DACs, their complexity and power dissipation must be minimized while allowing operation at 125 MHz.

Fig. 11 illustrates the DAC architecture, which consists of a coarse section driven by the five MSBs and a fine section driven

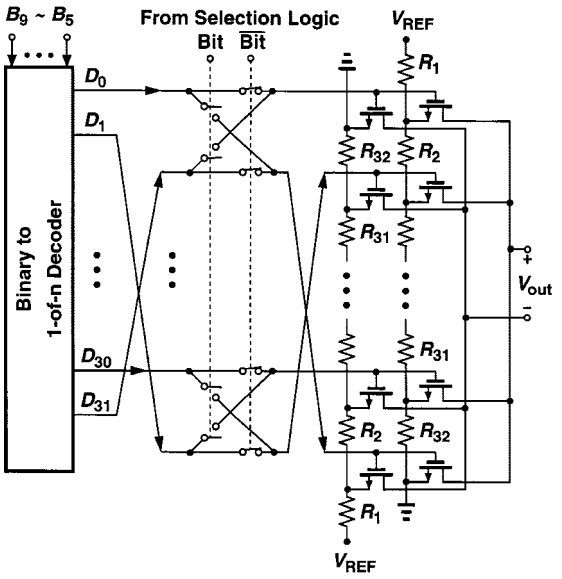

Fig. 12. Actual switching of the differential tap voltages.

by the five LSBs. Each section selects a tap voltage produced by a resistor ladder and converts the voltage to a current. The actual implementation is fully differential. Note that the output must be in the current domain to allow addition to the output of other DACs and eventually the output of VIC_1 in Fig. 5.

An important consideration in this DAC architecture is the required matching between the coarse and fine sections. Interestingly, the matching need not be at the 10-bit level to guarantee monotonicity. The mismatch between the input-referred offset voltages of the two V/I converters simply appears as an offset in the overall DAC characteristic. Furthermore, the gain error mismatch must be only small enough to ensure that the full-scale current of the fine V/I converter is one (fine) LSB less than one (coarse) LSB of the coarse V/I converter. This requirement can be expressed as

$$\sum_{n=1}^{5} 2^{-n} < (1-\alpha) \tag{7}$$

where α denotes the gain error mismatch. Thus, $\alpha < 0.0313$. Such degree of matching is obtained by careful layout and proper choice of device dimensions.

The role of each DAC in Fig. 11 is to multiply the 1-bit signal generated by the selection logic by the 10-bit coefficients m_j. This operation is performed in the DAC by swapping the outputs of each differential ladder according to the 1-bit signal. Fig. 12 illustrates the actual implementation with differential ladders.

D. V/I and I/V Converters

The V/I converters in Fig. 13 must exhibit sufficient linearity with input differential swings of about 2 V. Due to the lack of high-quality resistors in the CMOS technology used here, simple resistive degeneration is not feasible. Thus, as depicted in Fig. 13, the topology proposed in [3] is utilized. The choice $W_i/W_s = 7$ guarantees sufficient linearity for the LMS algorithm.

The currents produced by the four DACs and VIC_1 in Fig. 13 must be added and converted to voltage. As shown

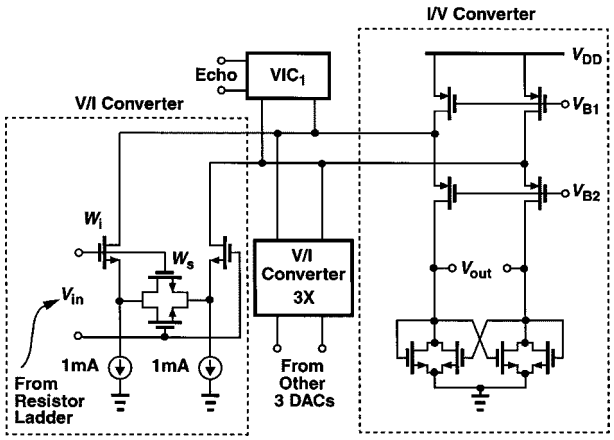

Fig. 13. V/I and I/V converters.

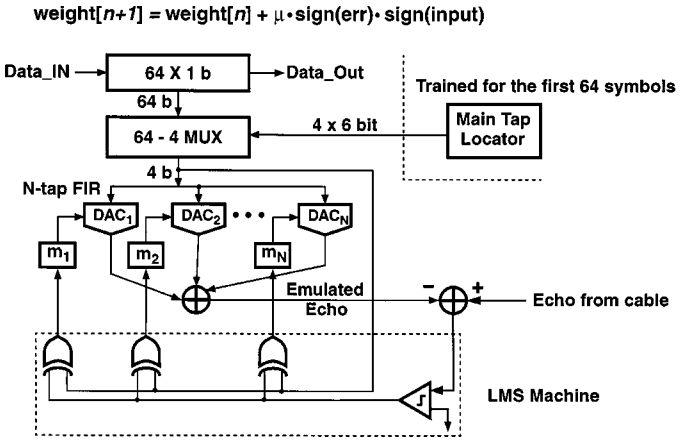

Fig. 14. LMS machine.

in Fig. 13, the addition is performed at the sources of a differential common-gate stage to minimize the resulting time constant. The I/V conversion is then carried out by means of diode-connected devices with a small amount of positive feedback. Interestingly, this type of load partially cancels the nonlinearity of the V/I converter, thus yielding a higher linearity. It also produces a greater load impedance than simple diode-connected devices, thereby increasing the voltage gain, requiring a smaller input swing for a given output swing, and hence improving the linearity.

Simulations indicate that the combination of V/I and I/V converters introduces a nonlinearity of about -40 dB in the signal.

E. LMS Machine

This design uses a sign–sign LMS machine. Shown in Fig. 14, the algorithm is based on (3). Since the updated value is given by the product of the signs of the input and the error, the circuit implementation can be reduced to an analog comparator and a few XOR gates.

With a given set of coefficients, the emulated echo is first subtracted from the echo generated by the cable. The polarity of the difference is then detected by a comparator and multiplied by the sign of the input data, thus producing the new value of the coefficients. The factor m determines the convergence time

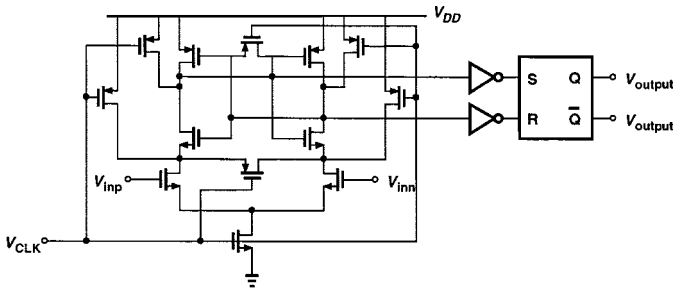

Fig. 15. Comparator.

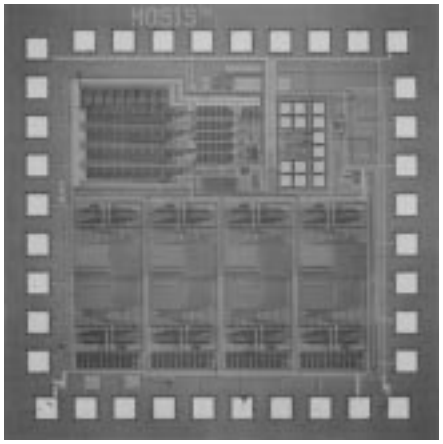

Fig. 16. Die photo.

and the residual echo. The factor is equal to 1 LSB of the DAC in this design.

F. Comparator

The echo canceller of Fig. 5 incorporates three comparators, one for comparing the emulated echo with the cable echo and the other two in the significant reflection locator. The offset of the first comparator is critical because it may raise the echo floor. Depicted in Fig. 15, the comparator topology is based on that in [4]. The device dimensions are chosen to yield an input-referred offset of about 5 mV according to [5], roughly 16 dB below the cancelled echo floor. (Since the comparator has 4 ns to reset in every cycle, it incurs negligible hysteresis.)

V. Experimental Results

The echo canceller has been fabricated in a digital 0.4-μm CMOS technology. The circuit is tested from a 3-V power supply with differential input swings of 2 V_{pp} and a sampling rate of 125 MHz. Shown in Fig. 16, the die has an active area of about 1 mm $\times$ 1 mm.

The test setup is shown in Fig. 17. A Tektronix arbitrary waveform generator, AWG 2020, applies a random sequence to the prototype. An on-chip transmitter drives the twisted pair and generates echo through a hybrid. A CAT-5 cable with intentional impedance mismatches is placed between the hybrid and a far-end load.

The training proceeds in two phases. In phase I, a step is applied and the tap locator circuit calculates the impulse response, identifying the four significant taps. This phase takes 64 symbols, i.e., 512 ns. In phase II, random data is applied and

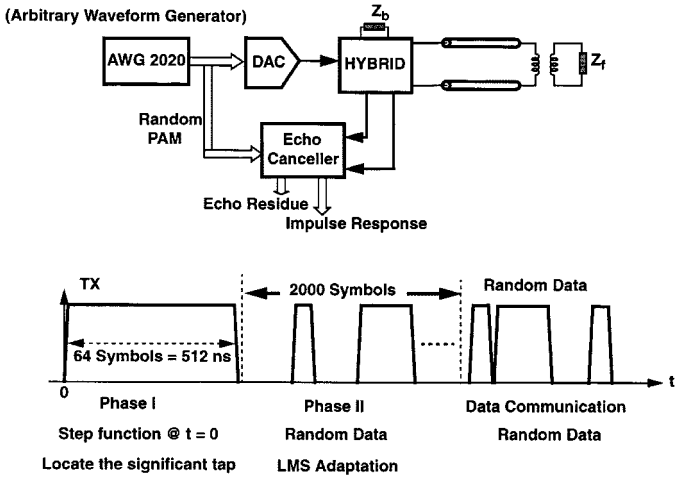

Fig. 17. Testing setup.

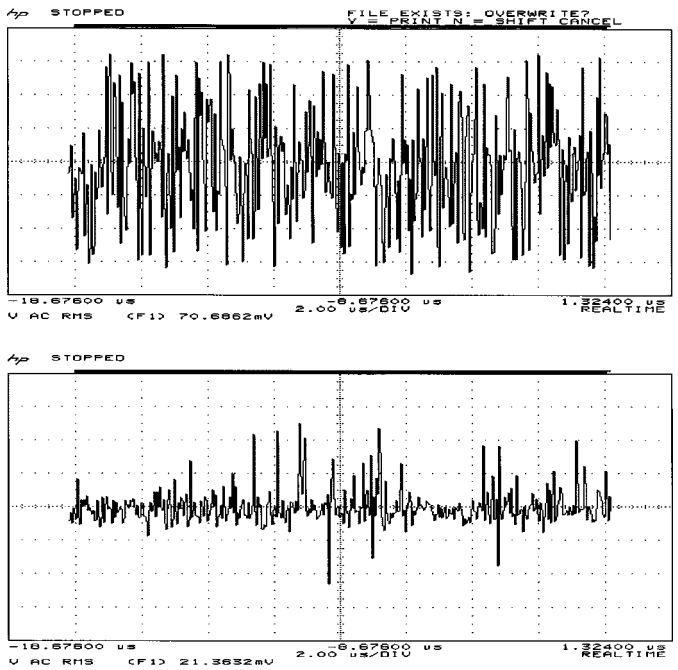

Fig. 18. Measured echo (25 mV/div) in time domain. (top) Before cancellation. (bottom) After cancellation.

the LMS machine adapts the FIR coefficients to minimize the residual echo. After phase II, the circuit is ready for full-duplex data communication.

In Fig. 18, the echo before and after cancellation is shown. Test results show a 10-dB reduction in the power of the echo. This reduction is taken from the average power rather than the instantaneous power. (A similar improvement would be attained by roughly 50 taps in digital echo/crosstalk cancellers.) Note that a few significant reflections are left after cancellation, indicating that more taps are required.

Fig. 19 depicts the convergence behavior of the LMS loop. The worst convergence time is approximately 10 μs—about 1000 symbol periods.

Table I summarizes the overall performance of the prototype. The power consumption is 43 mW from a 3-V supply during actual data communication.

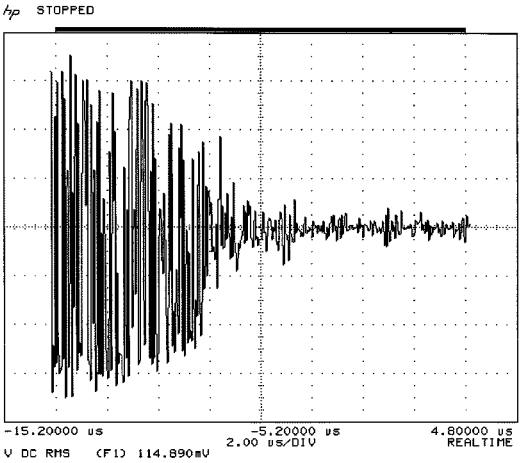

Fig. 19. Convergence behavior during LMS training.

TABLE I
MEASURED PERFORMANCE OF ECHO CANCELLER AT 125 MHz

Clock / Data rate	**125 MHz**
Power Dissipation	**43 mW**
Die Size (Active Area)	**1 mm x 1 mm**
Residual Echo Power	**-10 dB**
Convergence time	**10 μsec**
Supply Voltage	**3 V**
Technology	**0.4-μm digital CMOS**

VI. CONCLUSION

A four-tap 125-MHz CMOS mixed-signal echo canceller has been designed that incorporates a significant-tap locator circuit and mixed-signal adaptive FIR filter. The circuit partially cancels echo in the analog domain to reduce the complexity in the digital domain. A mixed-signal FIR filter with LMS adaptation is introduced that achieves high-speed high-coefficient resolution and low power dissipation. A discrete-time compact differentiator operating at 125 MHz is also presented that requires only two clock phases. A coarse+fine DAC architecture provides high speed with 10-bit resolution and low power dissipation. The performance is comparable with that of all-digital implementations [6].

ACKNOWLEDGMENT

The authors wish to thank S. Wu of Altima Communications for support.

REFERENCES

[1] A Tutorial Presentation for 1000 BASE-T, IEEE 802.3ab [Online]. Available: http://grouper.ieee.org/groups/802/3/ab/public/march98/index.html

[2] E. A. Lee and D. G. Messerschmidt, *Digital Communication.* Norwell, MA: Kluwer, 1994.

[3] F. Krummenacher and N. Joehl, "A 4-MHz CMOS continuous-time filter with on-chip automatic tuning," *IEEE J. Solid-State Circuits*, vol. 23, pp. 750–758, June 1988.

[4] Y.-T. Wang and B. Razavi, "An 8-bit 150-MHz CMOS A/D converter," *IEEE J. Solid-State Circuits*, vol. 35, pp. 308–317, Mar. 2000.

[5] M. Pelgrom, H. Tuinhout, and M. Vertregt, "Transistor matching in analog CMOS applications," in *Proc. IEDM*, San Francisco, CA, 1998, pp. 915–918.

[6] C. J. Nicol, P. Larsson, K. Azadet, and J. H. O'Neil, "A low-power 128-tap digital adaptive equalizer for broadband modems," *IEEE J. Solid-State Circuits*, vol. 32, pp. 1777–1789, Nov. 1997.

Tai-Cheng Lee was born in Taiwan, R.O.C., in 1970. He received the B.S. degree in electrical engineering from National Taiwan University in 1992 and the M.S. degree in electrical engineering from Stanford University, Stanford, CA, in 1994. He is currently working toward the Ph.D. degree in electrical engineering, focusing on the analog front-end of Gigabit Ethernet.

He worked for LSI logic from 1994 to 1997 as a Circuit Design Engineer. His main research interests are in high-speed mixed-signal circuit design, including analog filters, PLL systems, and data converters.

Behzad Razavi (S'87–M'90) received the B.Sc. degree in electrical engineering from Sharif University of Technology, Tehran, Iran, in 1985, and the M.Sc. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1988 and 1992, respectively.

He was with AT&T Bell Laboratories, Holmdel, NJ, and subsequently with Hewlett-Packard Laboratories, Palo Alto, CA. Since September 1996, he has been an Associate Professor of electrical engineering at the University of California, Los Angeles. His current research includes wireless transceivers, frequency synthesizers, phase-locking and clock recovery for high-speed data communications, and data converters. He was an Adjunct Professor at Princeton University, Princeton, NJ, from 1992 to 1994, and at Stanford University in 1995. He is a member of the Technical Program Committees of the Symposium on VLSI Circuits and the International Solid-State Circuits Conference (ISSCC), in which he is the chair of the Analog Subcommittee. He is an IEEE Distinguished Lecturer and the author of *Principles of Data Conversion System Design* (New York: IEEE Press, 1995), *RF Microelectronics* (Englewood Cliffs, NJ: Prentice-Hall, 1998), and *Design of Analog CMOS Integrated Circuits* (New York: McGraw-Hill, 2000), and the editor of *Monolithic Phase-Locked Loops and Clock Recovery Circuits* (New York: IEEE Press, 1996).

Dr. Razavi received the Beatrice Winner Award for Editorial Excellence at the 1994 ISSCC, the Best Paper Award at the 1994 European Solid-State Circuits Conference, the Best Panel Award at the 1995 and 1997 ISSCC, the TRW Innovative Teaching Award in 1997, and the Best Paper Award at the IEEE Custom Integrated Circuits Conference in 1998. He has also served as Guest Editor and Associate Editor of the IEEE JOURNAL OF SOLID-STATE CIRCUITS and IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS, and the *International Journal of High Speed Electronics*.