A Pipelined VLSI NEXT Canceller for Premises Applications

Gi-Hong Im

AT&T Bell Laboratories
200 Laurel Avenue
Middletown, NJ 07748

Naresh R. Shanbhag

AT&T Bell Laboratories
600 Mountain Avenue
Murrary Hill, NJ07974

ABSTRACT - A near-end crosstalk (NEXT) canceller using a fine-grain pipelined architecture is presented. The architecture is derived by the application of the *relaxed look – ahead* technique. This technique is an approximation to the *look – ahead* technique and results in a minimal hardware overhead. Performance of the proposed algorithm is demonstrated for NEXT cancellers used in 125 Mb/s twisted-pair distributed data interface (TPDDI) and the 155 Mb/s asynchronous transfer mode (ATM) local area network (LAN) applications. It is shown that the proposed pipelined architecture can be clocked at a rate of 107 times faster than the serial architecture with a maximum of 2.0 dB loss in the signal-to-noise ratio (SNR) for both the TPDDI and the ATM LAN applications.

I. INTRODUCTION

Applications such as mobile radio, high-definition television (HDTV), asynchronous transfer mode (ATM) networks, have increased the demand for high-speed and high-performance digital signal processing and communications algorithms. Usually, the high-speed and high-performance characteristics of an algorithm also imply a high complexity. Thus implementation, especially VLSI implementation, of such algorithms is made difficult and the transition period from algorithm design to silicon prototyping is increased. Therefore, in addition to the SNR performance, it is being recognized that implementation issues such as area, power dissipation and throughput need to be addressed during the algorithm design phase. One way in which this can be done is to employ *algorithm transformation techniques* [1] such as pipelining. Employed traditionally for high-speed applications, pipelined algorithms have found use in low-power applications as well [2]. Furthermore, by combining pipelining with folding [3], it is possible to trade off area with speed. Thus, it is clear that all the three major parameters of interest in a VLSI implementation, speed, power and area can be optimized by the design of pipelined algorithms.

In this paper, we present a pipelined architecture for a near-end crosstalk (NEXT) canceller and its performance for 125 Mb/s twisted-pair distributed data interface (TPDDI) and 155 Mb/s ATM local area network (LAN) applications. It has been shown in [4] that data rates above 100 Mb/s can be achieved over 100m of unshielded twisted pair (UTP) category 3 cable. In this case NEXT has to be restricted to one single synchronous cyclostationary interferer and the transceiver has to utilize a NEXT canceller. The proposed pipelined architecture for the NEXT canceller is derived via the *relaxed look – ahead* technique [5]. This technique is an approximation of the *look – ahead* technique [1]. The look-ahead technique transforms a serial algorithm into an equivalent (in the sense of input-output behavior) pipelined algorithm. The relaxed look-ahead technique sacrifices this equivalence between the serial and pipelined algorithms at the expense of marginally altered convergence characteristics. Hence, architectures resulting from the application of relaxed look-ahead have a negligible hardware overhead. However, there is a small performance degradation at high speed-ups.

The transmission scheme considered in this paper is carrierless AM/PM (CAP), which is a bandwidth-efficient two-dimensional passband line code [6]. CAP has the same spectral characteristics and provides the same theoretical performance as quadrature amplitude modulation (QAM), but is generally less complex to implement digitally. The 125 Mb/s 32-CAP line code has been proposed to the ANSI X3T9.5 TP/PMD working group as a candidate for the twisted-pair distributed data interface (TPDDI) standard over category 3 cable. Recently, the 16-CAP line code was accepted as ATM LAN standard for transmission at 51.84 Mb/s over UTP-3 [7].

The outline of the paper is as follows. In section II, a transceiver structure with NEXT canceller is described. The proposed pipelined NEXT canceller architecture is presented in section III, while its performance is discussed in section IV.

II. TRANSCEIVER STRUCTURE

In this section, we briefly discuss channel and NEXT models for UTP-3 cables and the CAP transceiver structure.

Channel and NEXT Models for Category 3 Cable

0-7803-1820-X/94 $4.00 © 1994 IEEE

The propagation loss assumed is the worst-case loss given in the EIA/TIA-568 draft standard for category 3 cable [8]. This loss can be approximated by the following expression:

$$L_P(f) = 7.07\sqrt{f} + 0.73f\,. \qquad (1)$$

where the propagation loss $L_P(f)$ is expressed in dB per kilofoot and the frequency f is expressed in MHz. The phase characteristics of the loop's transfer function is computed from $\sqrt{LC}$.

The worst-case NEXT loss model for a single interferer is also given in the EIA/TIA draft standard. The squared magnitude of the NEXT transfer function corresponding to this loss can be expressed in the following way:

$$L_N(f) = 41 - 15\log f\,. \qquad (2)$$

where the frequency f is assumed to be expressed in MHz. The measured pair-to-pair NEXT loss characteristic of category 3 cable has the minima (small loss) and maxima (large loss) in its curve. It has been shown that the NEXT loss corresponding to the minima decreases with increasing frequency and tends to follow the smooth curve of (2), which is the so-called 15 dB per decade model.

Transceiver Structure

The results presented in this paper apply to the two-pair dual-simplex transmission scenario shown in Fig. 1. Each direction of transmission uses a different pair. The two pairs are assumed to be in the same cable so that each pair generates near-end crosstalk (NEXT) and far-end crosstalk (FEXT) in the other pair. In the two-pair dual-simplex approach, NEXT is the only damaging impairment. Thus, FEXT will not be considered any further here.

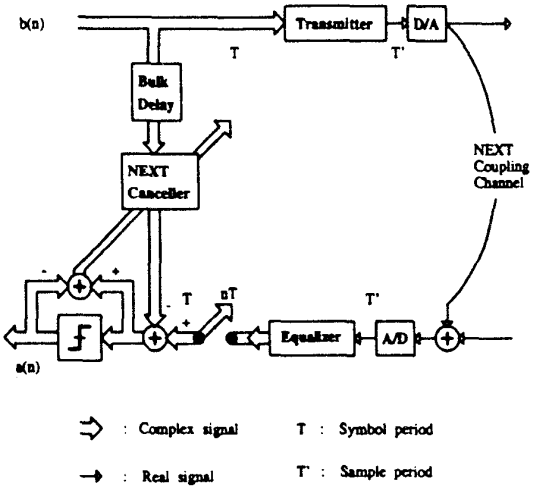

Fig. 1 A Transceiver Structure for Premises Applications

The transceiver incorporates a fractionally spaced linear equalizer (FSLE) and NEXT canceller. The purpose of a NEXT canceller is to generate a replica of the signal that has passed through the NEXT coupling channel. This replica is then subtracted from the incoming signal, thus eliminating the NEXT interferer. The NEXT canceller shown in Fig. 1 uses a so-called cross-coupled symbol-spaced structure. The inputs of the canceller are the symbols $b(n)$ generated by the encoder, and its outputs are subtracted from the real and imaginary signals after the baud sampler at the output of FSLE. The advantage of such a NEXT canceller structure is that all the computations are performed at the symbol rate. An alternative is to do NEXT cancellation immediately after the analog-to-digital (A/D). This require the computations to be done at the sampling rate of the A/D, which is typically 3 to 4 times the baud rate.

The time interval that the NEXT canceller has to span, or memory span, can be obtained from the measured NEXT loss characteristics. In our performance study, we used the worst case measured NEXT whose impulse response including the transmit shaping filter spans about 1 μs. It has been found that the equalizer has little effect on the duration of this impulse response [4]. Thus, the memory span of the NEXT canceller should be in the 1 μs range.

III. PIPELINED NEXT CANCELLER (NC)

In the following, the $a = a_r + ja_i$ denotes a complex variable with a_r and a_i being the real and imaginary parts, respectively. The serial NEXT canceller can be described as follows

$$\mathbf{w}(n) = \mathbf{w}(n-1) + \mu e(n)\mathbf{b}^*(n)\,, \qquad (3)$$

$$e(n) = y(n) - a(n) - \mathbf{w}^t(n-1)\mathbf{b}(n)\,, \qquad (4)$$

where * denotes complex conjugation operation, $\mathbf{w}(n) = \mathbf{w}_r(n) + j\mathbf{w}_i(n)$ is a $N\times 1$ coefficient vector of the NEXT canceller, $y(n) = y_r(n) + jy_i(n)$ is the equalizer output, $a(n) = a_r(n) + ja_i(n)$ is the slicer output, $\mathbf{b}(n) = \mathbf{b}_r(n) + j\mathbf{b}_i(n)$ is the data symbol at the local transmitter and $e(n) = e_r(n) + je_i(n)$ is the error signal. Note that the complex least mean squared (LMS) algorithm has been employed to minimize the error across the slicer.

In Fig. 2, we show a serial NEXT canceller with two complex taps. The convolution of the data symbols and the coefficients (see (4)) is done in the **F** block, while the weight-update (see (3)) is computed in the **WUD** block. Let T_m, T_{a1} and T_{a2} denote the computation times of a multiplier, the adder in the **WUD** block, and the adder in the **F** block, respectively.

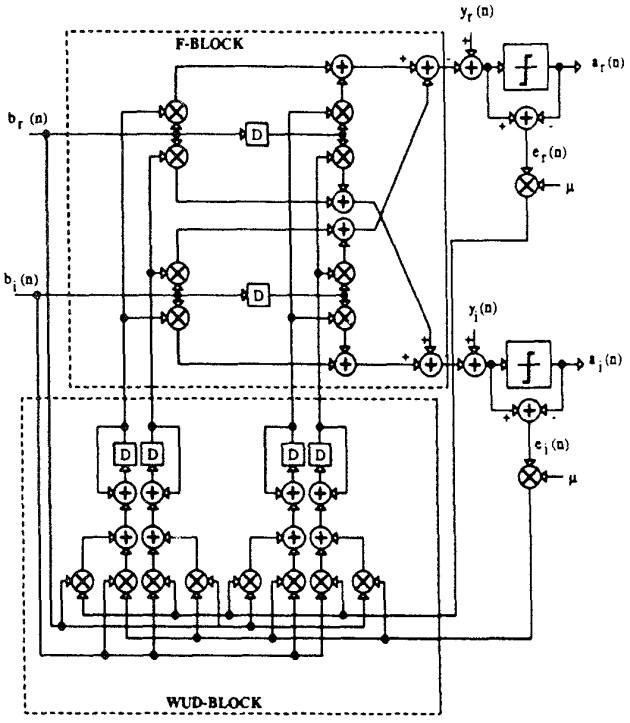

Fig. 2 The Serial NEXT Canceller with $N = 2$.

It is clear (see Fig. 2) that the minimum achievable sample period for the serial NEXT canceller T_{serial} is given by

$$T_{serial} = 3T_m + 2T_{a1} + (N + 2)\, T_{a2}\ , \qquad (5)$$

where N corresponds to the number of complex taps, which is equal to 2 in Fig. 2. Let us assume that the output of a 1-b full adder takes $2ns$ to settle with nominal load. Furthermore, assume an $8b\times8b$ multiplier, $20b$ adder in the **WUD** block and a $10b$ adder in the **F** block. Therefore, reasonable estimates for the computation times are $T_m = 32ns$, $T_{a1} = 40ns$ and $T_{a2} = 20ns$. Substituting these values into (5), we get $T_{serial} = 256ns$. For the NEXT canceller, the input sample rate equals the symbol rate $1/T$. Hence, it is possible that for high symbol rates and large values of N, $T_{serial} > T$. In fact, $T = 40ns$ and $T = 38.7ns$ for the TPDDI and ATM LAN applications, respectively. In such cases, the serial architecture of Fig. 2 cannot meet the sample rate requirements. Therefore, there is a need for a pipelined architecture which can operate at a sample period T_{pipe}, where $T_{pipe} \le T$. Such a pipelined architecture is said to have a speed-up (SU) over the serial architecture, where SU is defined as

$$SU = T_{serial} / T_{pipe}\ . \qquad (6)$$

The pipelined NEXT canceller can be derived via the application of the relaxed look-ahead technique [5]. As (3) and (4) employ the LMS algorithm, the pipelined NEXT canceller architecture is obtained by employing the pipelined LMS algorithm described in [6]. The resulting pipelined NEXT canceller algorithm is described by

$$\mathbf{w}(n) = \mathbf{w}(n-D_2) + \mu \sum_{i=0}^{LA-1} e(n-D_1-i)\,\mathbf{b}^*(n-D_1-i) \qquad (7)$$

$$e(n) = y(n) - a(n) - \mathbf{w}^t(n-D_2)\,\mathbf{b}(n)\ , \qquad (8)$$

where D_1 and D_2 are pipelining latches and the look-ahead factor $LA \le D_2$. Note that the hardware overhead due to relaxed look-ahead are the pipelining latches and $2N(LA-1)$ adders. The introduction of D_1 and D_2 delays results in altered convergence behavior. Convergence analysis of the pipelined LMS [5] indicated that the bounds on step-size μ are slightly tighter than that of the serial algorithm. Furthermore, the convergence speed and adaptation accuracy were also found to be slightly degraded. For most practical applications, the loss in performance due to pipelining is negligible and is overwhelmed by the resulting architectural advantages. This fact is demonstrated in section IV for a 125 Mb/s TPDDI and 155 Mb/s ATM LAN applications.

IV. PERFORMANCE OF THE PIPELINED NC

In this section we investigate the performance of the pipelined NEXT canceller of 125 Mb/s 32-CAP and 155 Mb/s 64-CAP transceivers over category 3 cable. We will assume that one loop is utilized for each direction of transmission, as shown in Fig. 1, and that the same kind of line code is used on each loop. Thus, the NEXT interferer is a data signal that is similar to the disturbed signal. With this model, the inputs to the transmitter on the upper left in Fig. 1 are data symbols $b(n)$, which are assumed to be uncorrelated with the symbols $a(n)$ that are recovered at the output of the slicer on the lower left in the figure. We will also assume that the disturbed and interfering signals have clocks which are synchronized in frequency.

In order to investigate the performance of pipelined NEXT canceller, we used the following start-up procedure which consists of four main steps.

- Step #1: The NEXT interferer appearing at the right adder in Fig. 1 is first set to zero.
- Step #2: The equalizer is then converged to compensate for the linear distortion introduced by the loop; after convergence, the tap coefficients of the equalizer are frozen.
- Step #3: The NEXT interferer is added at the input of the equalizer, as shown in Fig. 1; a number of taps is selected for the NEXT canceller, which is then converged for various values of the bulk delay line until the best bulk delay is identified.
- Step #4: The NEXT canceller is fully converged with the optimum bulk delay and the steady-state SNR at the slicer is computed.

Figure 3 shows the computer simulation results for the performance of the pipelined NEXT canceller with different D_1 and D_2 values. The pipelined canceller with $D_1 = 0$ and $D_2 = 1$ corresponds to the serial NEXT canceller. With $D_1 = 123$ and $D_2 = 5$, the pipelined NEXT canceller can be clocked at a rate of 107 times faster than the serial NEXT canceller. However, there is 1.8 dB degradation in SNR_o and it takes about four times more iterations for this pipelined NEXT canceller to converge to its steady-state. It is noted that the absolute convergence speed with $D_2 = 5$ and $D_1 = 123$ is over 20 times faster than the serial NEXT canceller because of the speed-up factor of the pipelined canceller.

Table 1 summarizes the performance results of the pipelined canceller shown in Fig. 3. In Table 1, α stands for excess bandwidth and SNR_i is the signal-to-noise ratio at the input of the receiver. The first column in Table 1 gives the speed-up factor comparing to the serial NEXT canceller. The last column gives the margin, which is defined as

$$\text{margin} = SNR_o - SNR_{o,ref} \qquad (9)$$

where $SNR_{o,ref}$ is a suitably chosen reference for the SNR at the input of the decision device. In Table 1 we have chosen $SNR_{o,ref} = 27.13$ dB, which corresponds to the value of SNR_o that provides a probability of error of

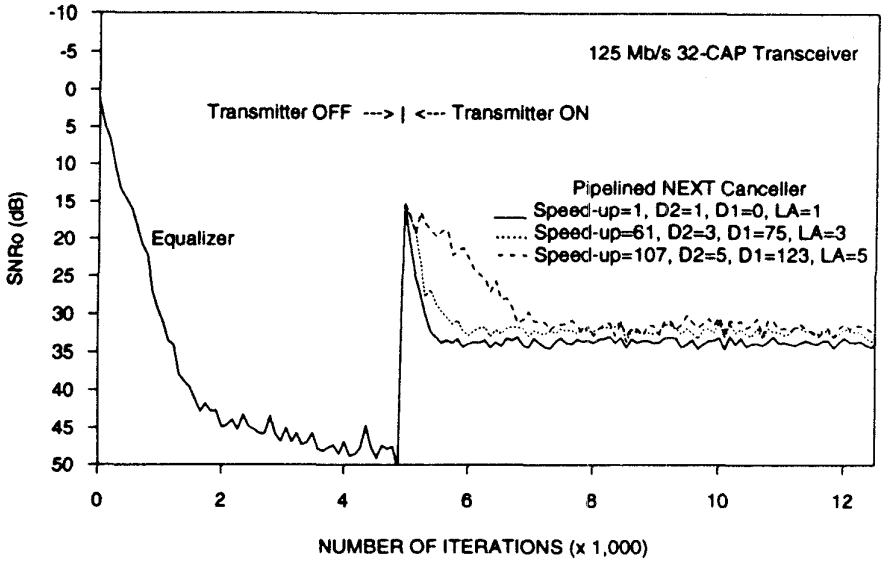

Fig. 3 Performance of the 125 Mb/s 32-CAP Transceiver with Pipelined NEXT Canceller

Table 1 - Performance of Pipelined NEXT Canceller with Relaxed Look-ahead
125 Mb/s 32-CAP TPDDI Transceiver Operating over 100m Category 3 Cable
$\alpha = 0.2$ $T' = T/3$ **Pair (1,3)** $SNR_i = 15.9\ dB$ $P_e = 10^{-12}$ †
Equalizer (42T) NEXT Canceller (32T)

Speed – up	D_2	D_1	LA	SNR_o (dB)	Margin (dB)
1	1	0	1	37.65	10.52
61.14	3	75	3	36.33	9.20
107	5	123	5	35.84	8.71

† With 32-CAP and Gaussian noise, the required SNR_o for this probability of error is 27.13 dB.

10^{-12} for a 32-CAP transceiver. (This assumes that the noise at the slicer is Gaussian, which may be a somewhat pessimistic assumption for the single-interferer case considered here.)

In Fig. 4, we plot the performance of the pipelined NEXT canceller with different values of speed-up *SU* for TPDDI application. The values of D_1 and D_2 for a certain speed-up depend upon the number of taps in the NEXT canceller and the speed of the computational blocks such as the multiplier and adder. For example, suppose the desired speed-up is such that $T_{pipe} > T_{a1}$. In that case, this speed-up can be achieved with $D_2 = 1$ and a sufficiently high value for D_1. Notice that the 125 Mb/s 32-CAP transceiver has still comfortable margins even when the speed-up of the pipelined NEXT canceller is about 100. Table 2 summarizes the performance results of the pipelined NEXT canceller for 155 Mb/s ATM LAN application. It is noted that the speed-up factor of 107 can be achieved with 2 dB loss in the margin.

V. CONCLUSIONS

A hardware efficient pipelined NEXT canceller architecture has been presented. The architecture has been derived via the relaxed look-ahead technique. Performance of the proposed architecture in 125 Mb/s TPDDI and 155 Mb/s ATM LAN environments indicates that substantially high speed-ups can be achieved at the expense of a small SNR degradation. For any given application the speed-up due to pipelining can be traded-off with power and/or area and thereby achieve an efficient VLSI implementation.

VI. ACKNOWLEDGEMENTS

The authors would like to thank J. J. Werner and Jit Kumar for their support of this work.

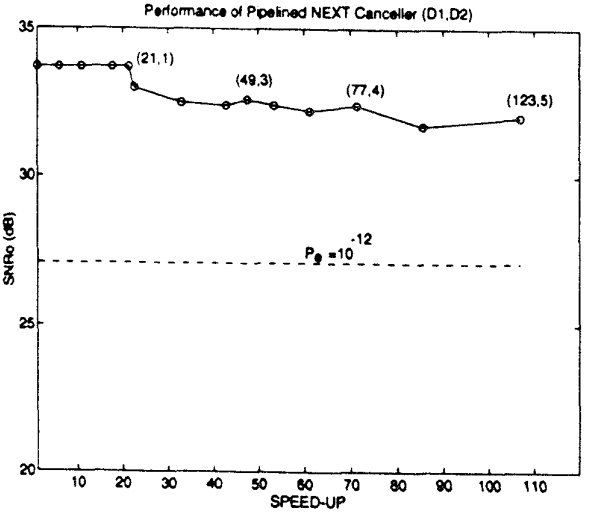

Fig. 4 SNR_o vs. Speed-Up for the Pipelined NEXT Canceller.

Table 2 - Performance of Pipelined NEXT Canceller with Relaxed Look-ahead
155 Mb/s 64-CAP ATM LAN Transceiver Operating over 100m Category 3 Cable
$\alpha = 0.16$ $T' = T/3$ **Pair (1,3)** $SNR_i = 17.0\ dB$ $P_e = 10^{-12}$ †
Equalizer (42T) **NEXT Canceller (32T)**

Speed – up	D2	D1	LA	SNR_o (*dB*)	Margin † (*dB*)
1	1	0	1	36.02	5.69
61.14	3	75	3	34.80	4.47
107	5	123	5	34.03	3.70

† With 64-CAP and Gaussian noise, the required SNR_o for this probability of error is 30.33 dB.

REFERENCES

[1] K. K. Parhi, "Algorithm transformation techniques for concurrent processors", *Proceedings of the IEEE*, vol. 77, pp. 1879-1895, Dec. 1989.

[2] A. P. Chandrakasan, S. Sheng and R. W. Brodersen, "Low power CMOS digital design," *IEEE J. of Solid-State Circuits*, vol. 27, pp. 473-484, April 1992.

[3] K. K. Parhi, C.-Y. Wang, and A. P. Brown, "Synthesis of control circuits in folded pipelined DSP architectures", *IEEE J. of Solid-state Circuits*, vol. 27, no. 1, pp. 29-43, January 1992.

[4] G.H. Im and J.J. Werner, "Bandwidth-efficient digital transmission up to 155 Mb/s over unshielded twisted pair wiring," *ICC 93*, May 1993.

[5] N.R. Shanbhag and K.K. Parhi, "Relaxed look-ahead pipelined LMS adaptive filters and their application to ADPCM coder," *IEEE Trans. on Circuits and Systems*, Dec., 1993.

[6] W.Y. Chen, G.H. Im and J.J. Werner, "Design of digital carrierless AM/PM transceivers," AT&T/Bellcore Contribution T1E1.4/92-149, Aug. 19, 1992.

[7] J.J. Werner, G.H. Im, D.D. Harman, A.V. Mandzik, M-H. Nguyen and R. Townsend, "The 16-CAP Encoding Scheme and Its Application to The ATM UTP-3 UNI at 51.84 Mb/s," UTP-3 Subworking Group of the ATM Forum, May 19, 1993.

[8] *The Commercial Building Telecommunications Wiring Standard*, EIA/TIA-568 Draft Standard, Dec. 1990.