

A Pipelined Adaptive NEXT Canceller

Gi-Hong Im and Naresh R. Shanbhag

Abstract—A near-end crosstalk (NEXT) canceller using a fine-grain pipelined architecture is presented. Performance of the proposed pipelined NEXT canceller is demonstrated in the 125 Mb/s twisted-pair distributed data interface and 155.52 Mb/s asynchronous transfer mode local area network applications. In addition, we analyze the computational complexity of the proposed pipelined NEXT canceller. It is shown that this architecture can be clocked at a rate that is 107 times faster than the serial architecture with a maximum loss of 2.0 dB in signal-to-noise ratio (SNR).

Index Terms—Echo canceller, equalizer, LAN, NEXT, NEXT canceller, pipelining, VLSI implementation.

I. INTRODUCTION

In this correspondence, 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 [1] that data rates above 100 Mb/s can be achieved over 100 m 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 [1]. The proposed pipelined architecture for the NEXT canceller is derived via the *relaxed look-ahead* technique [2], [3]. The transmission scheme considered in this correspondence is carrierless AM/PM (CAP), which is a bandwidth-efficient two-dimensional (2-D) passband line code [4]. The 51.84 Mb/s 16-CAP [5] and the 155.52 Mb/s 64-CAP [1] line codes have been proposed to the PHY subworking group of the ATM Forum as candidates for ATM LAN standard over category 3 cable. Recently, the 16-CAP and 64-CAP (with NEXT canceller) line codes were accepted as ATM LAN standards for transmission at 51.84 Mb/s [6] and 155.52 Mb/s [7] over UTP-3, respectively.

The outline of this correspondence 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. In Section IV, we analyze the computational complexity of the proposed NEXT canceller. Simulation results and discussion with worst-case measured NEXT are presented in Section V.

II. TRANSCIEVER STRUCTURE

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

A. Channel and NEXT Models for Category 3 Cable

The two major causes of performance degradation for transceiver operating over UTP wiring are propagation loss and crosstalk generated between pairs [1]. 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

$$L_P(f) = 7.07\sqrt{f} + 0.73f \quad (1)$$

where the propagation loss $L_P(f)$ is expressed in decibels per kilofoot, and the frequency f is expressed in megahertz. The phase characteristics of the loop's transfer function is computed from \sqrt{LC} , where R , L , G , and C are the primary constants of a cable. 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 as

$$L_N(f) = 41 - 15 \log f \quad (2)$$

where the frequency f is assumed to be expressed in megahertz. The wavy curves in Fig. 1 give the measured pair-to-pair NEXT loss characteristics for three different combinations of twisted pairs in 100-m category 3 cables.

B. CAP Transceiver Structure

In this subsection, we first provide a description of a generic CAP transceiver [5]. We then consider a transceiver with NEXT canceller. The signal at the output of the CAP transmitter can be written as

$$s(t) = \sum_{n=-\infty}^{\infty} [a_r(n)p(t - nT) - a_i(n)\tilde{p}(t - nT)] \quad (3)$$

where T is the symbol period, $a_r(n)$ and $a_i(n)$ are discrete multilevel symbols that are sent in symbol period nT , and $p(t)$ and $\tilde{p}(t)$ are the impulse responses of in-phase and quadrature passband shaping filters, respectively. The passband pulses $p(t)$ and $\tilde{p}(t)$ in (3) can be designed as in

$$p(t) \triangleq g(t) \cos(2\pi f_c t) \quad \tilde{p}(t) \triangleq g(t) \sin(2\pi f_c t) \quad (4)$$

where $g(t)$ is a baseband pulse, and f_c is a frequency that is larger than the largest frequency component in $g(t)$. The two impulse responses in (4) form a so-called Hilbert pair, i.e., their Fourier transforms have the same amplitude characteristics and phase characteristics that differ by 90° .

It is shown in [1] that the usage of NEXT canceller is necessary to achieve data rates above 100 Mb/s over category 3 cable. In this subsection, we consider a transceiver incorporating a fractionally spaced linear equalizer (FSLE) and a NEXT canceller. The purpose of the 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. A NEXT canceller has the same principle of operation as an echo canceller, and all the familiar structure used for echo cancelers can also be used for NEXT cancelers. The NEXT canceller shown in Fig. 2 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 converter (A/D). This requires that the computations be done at the sampling rate of the A/D, which is typically three to four times the symbol 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

Manuscript received January 3, 1997; revised January 29, 1998. The associate editor coordinating the review of this paper and approving it for publication was Dr. Phillip A. Regalia.

G.-H. Im is with the Department of Electrical Engineering, POSTECH, Pohang, Kyungbuk, Korea.

N. R. Shanbhag is with the Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA.

Publisher Item Identifier S 1053-587X(98)05238-6.

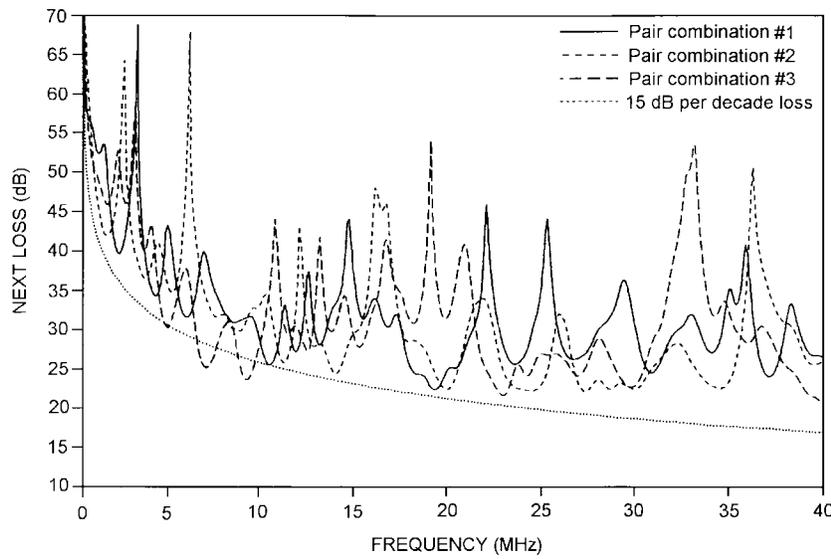


Fig. 1. Measured NEXT loss between pairs of category 3 cables.

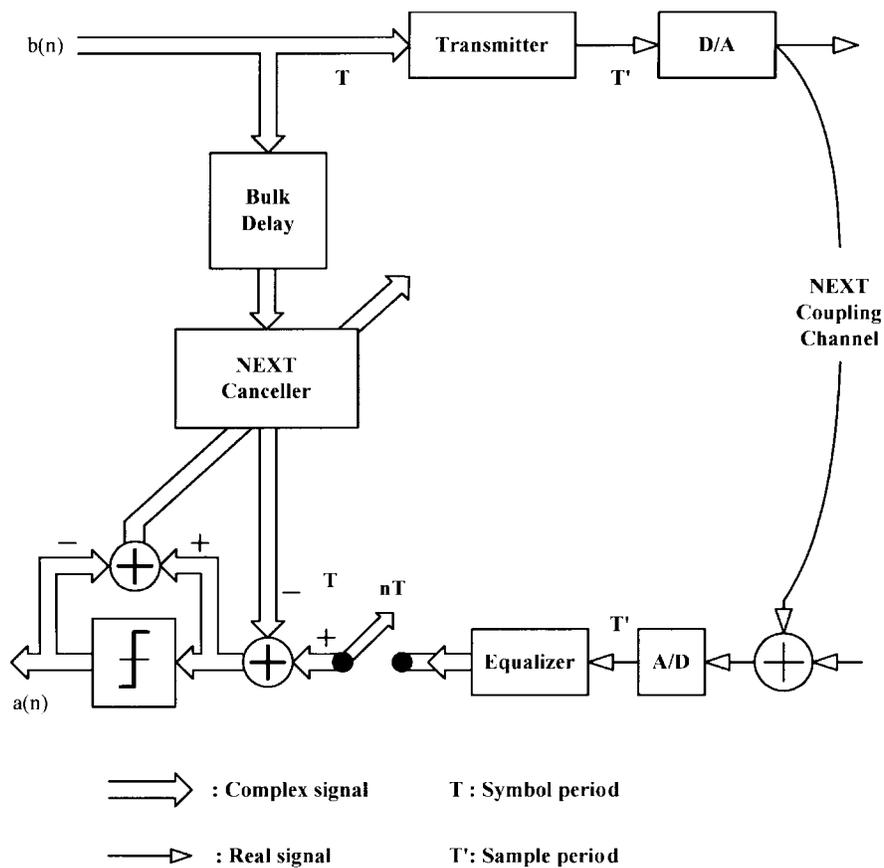


Fig. 2. Transceiver structure for premises applications.

impulse response, including the transmit shaping filter, spans about 1 μ s. It has been found that the equalizer has little effect on the duration of this impulse response [1]. Thus, the memory span of the NEXT canceller should be in the 1- μ s range.

III. PIPELINED NEXT CANCELLER ARCHITECTURE

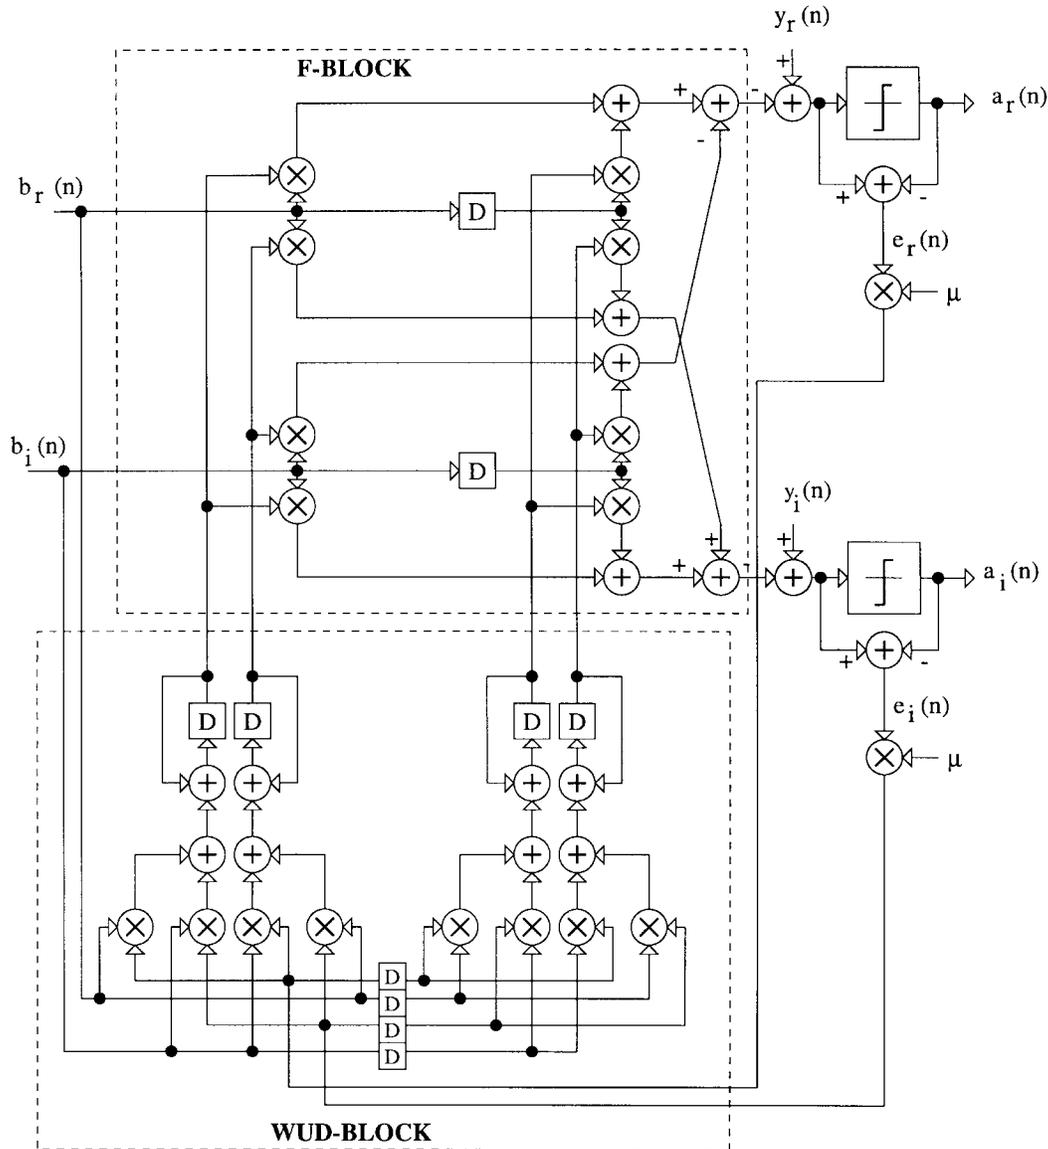
In this section, we propose the pipelined NEXT canceller architecture and briefly summarize its convergence characteristics.

A. Architecture

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

$$\mathbf{w}(n) = \mathbf{w}(n-1) + \mu \epsilon(n) \mathbf{b}^*(n) \tag{5}$$

$$\epsilon(n) = y(n) - a(n) - \mathbf{w}(n)^t \mathbf{b}(n) \tag{6}$$

Fig. 3. Serial NEXT canceller with $N = 2$.

where $*$ and t denote complex conjugation and complex transpose operations, respectively, and $\mathbf{w}(n) = \mathbf{w}_r(n) + j\mathbf{w}_i(n)$ 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 vector at the local transmitter, $e(n) = e_r(n) + je_i(n)$ is the error signal, and μ is the adaptation step size. Note that the complex LMS algorithm has been employed to minimize the error across the slicer.

In Fig. 3, we show a serial NEXT canceller with two complex taps. The convolution of the data symbols and the coefficients [see (6)] is done in the **F** block, whereas the weight update [see (5)] 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. It is clear (see Fig. 3) that the minimum achievable sample period for the serial NEXT canceller T_{serial} is given by

$$T_{serial} = 3T_m + 2T_{a1} + (N + 2)T_{a2} \quad (7)$$

where N corresponds to the number of complex taps, which is equal to 2 in Fig. 3. The critical path employed in computing (7) starts

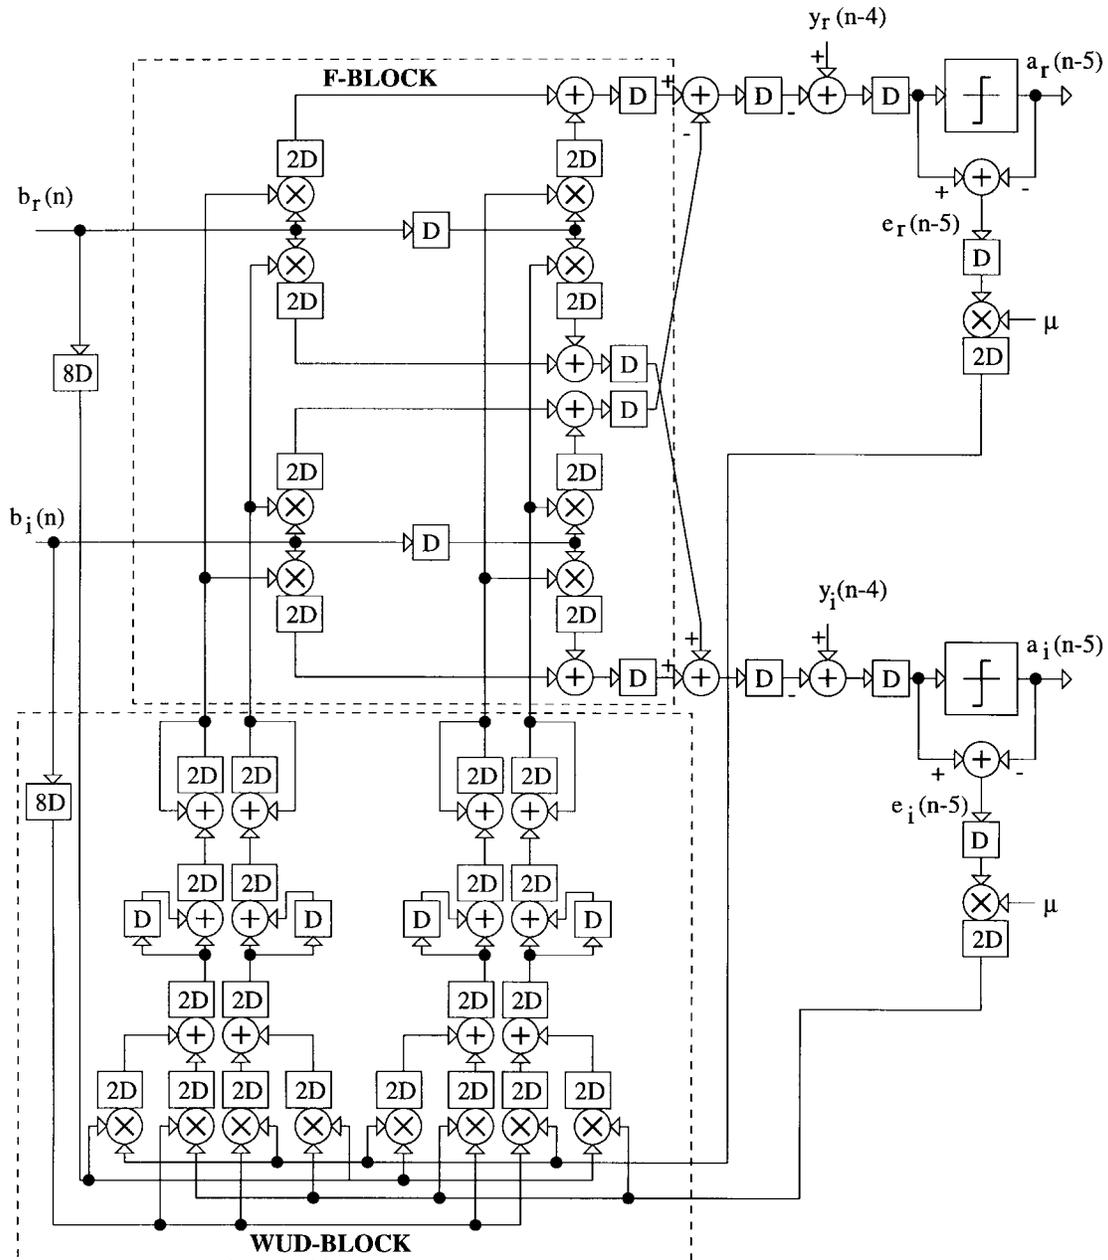
at the primary input $b_r(n)$ [or $b_i(n)$], and it passes through the **F** block before ending at the input to the delay in the inner loop in the **WUD** block. Let us assume that the output of a $1-b$ full adder takes 2 ns to settle with nominal load. Furthermore, assume an $8b \times 8b$ multiplier, $20b$ adder in the **WUD** block and a $10b$ adder in the **F** block. Therefore, reasonable estimates for the computation times are

$$T_m = 32 \text{ ns};$$

$$T_{a1} = 40 \text{ ns};$$

$$T_{a2} = 20 \text{ ns}.$$

Substituting these values into (7), we get $T_{serial} = 256$ ns. 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, $T = 40$ ns and $T = 38.7$ ns for the TPDDI and ATM LAN applications, respectively. In such cases, the serial architecture of Fig. 3 cannot meet the sample rate requirements. Therefore, there is a need for a pipelined architecture that can operate at a sample period T_{pipe} , where $T_{pipe} \leq T$. Such a pipelined architecture is said to have a speedup (SU) over the serial


 Fig. 4. Pipelined NEXT canceller with $N = 2$.

architecture, where SU is defined as

$$SU = T_{serial}/T_{pipe}. \quad (8)$$

The pipelined NEXT canceller can be derived via the application of the relaxed look-ahead technique [2]. As (5) and (6) employ the LMS algorithm, the pipelined NEXT canceller architecture is obtained by employing the pipelined LMS algorithm described in [3]. The resulting pipelined NEXT canceller algorithm is described by

$$\mathbf{w}(n) = \mathbf{w}(n - D_2) + \mu \sum_{i=0}^{LA-1} e(n - D_1 - i) \cdot \mathbf{b}^*(n - D_1 - i) \quad (9)$$

$$e(n) = y(n) - a(n) - \mathbf{w}(n)^t \mathbf{b}(n) \quad (10)$$

where D_1 and D_2 are pipelining latches, and the look-ahead factor $LA \leq D_2$. Note that the hardware overhead due to relaxed look-ahead are the pipelining latches and $2N(LA - 1)$ adders. It can be

shown that the minimum achievable sample period for the pipelined architecture (see Fig. 4, with $N = 2$) T_{pipe} is given by

$$T_{pipe} \approx \max \left(\frac{T_{a1}}{D_2}, \frac{T_{a1}LA + 3T_m + (N + 2)T_{a2}}{D_1} \right) \quad (11)$$

where $D_1 > 0$. The introduction of D_1 and D_2 delays results in altered convergence behavior. Convergence analysis of the pipelined LMS [3] 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 [see (12) and (16)]. 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 V for 125-Mb/s TPDDI and 155-Mb/s ATM LAN applications.

The architecture of the pipelined NEXT canceller with two complex taps is shown in Fig. 4. The latches at the output of the

TABLE I
COMPARISON OF HARDWARE COMPLEXITIES BETWEEN SERIAL AND PIPELINED NEXT CANCELLER

Serial and Pipelined NEXT Cancellor			
Architecture	No. of Multipliers	No. of Adders	No. of Latches
Serial	$8N + 2$	$8N + 2$	$8N - 6$
Pipelined	$8N + 2$	$6N + 2N \cdot LA + 2$	$4N + 2N(D_2 + LA) + 4D_1 - 6$

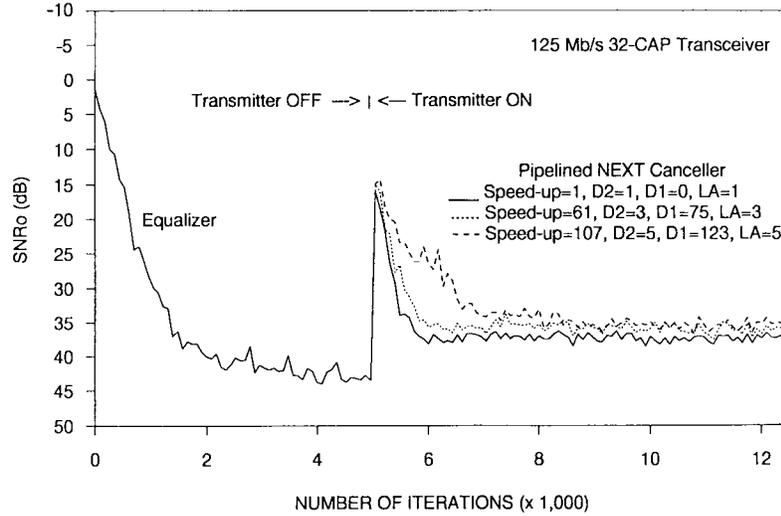


Fig. 5. Performance of the 125 MB/s 32-CAP transceiver with pipelined NEXT canceller.

multipliers and the adders (in the WUD block) and the multipliers (in the \mathbf{F} block) would be employed to pipeline the respective computational blocks. In a practical implementation, the latches D_1 and D_2 would be retimed [9] to pipeline the NEXT canceller. Assuming that the multipliers and adders in Figs. 3 and 4 are identical, it can be shown that the architecture in Fig. 4 (with $D_1 = 14$, $D_2 = 2$, $LA = 2$) has a minimum achievable clock period of $T_{pipe} = 20$ ns. This corresponds to an SU of 12. As mentioned before, the architecture in Fig. 4 can be folded [10] to reduce area at the expense of SU. Furthermore, any additional SU can be traded off with power [11] to obtain a low-power implementation.

B. Convergence Characteristics of Pipelined NEXT Cancellor

The relaxed look-ahead technique alters the input-output behavior of the algorithm to which it is applied. Hence, a convergence analysis of the resulting pipelined algorithm is necessary. In this subsection, we apply the convergence analysis results from [2] to formulate the bounds on the step size μ for convergence in the mean-squared sense and an expression for misadjustment of the pipelined NEXT canceller.

The upper bound on the step size μ to guarantee the convergence of the MSE of the pipelined NEXT canceller is tighter as compared with that of the serial architecture. In particular, exploiting the fact that the input of the pipelined NEXT canceller is uncorrelated, the bound on μ given in [2] is simplified as

$$0 \leq \mu \leq \frac{P + 2K - \sqrt{(P + 2K)^2 - 8K(K + 1)}}{LA \cdot K(K + 1)\sigma^2} \quad (12)$$

where $D_1 = KD_2$ with D_2 being at least unity

$$P = N + \nu - 1, \quad \nu = \frac{\langle |b(n)|^4 \rangle}{(\langle |b(n)|^2 \rangle)^2}. \quad (13)$$

Even though this bound on the step size is tighter than that of the serial algorithm, it does not represent a serious drawback. This is due to the fact that in an actual implementation, the step size μ is much smaller than this upper bound.

The adaptation accuracy of an adaptive algorithm is quantified by its misadjustment, which is defined as

$$M = \frac{\langle |e(\infty)|^2 \rangle - \langle |e_{\min}|^2 \rangle}{\langle |e_{\min}|^2 \rangle} \quad (14)$$

where $\langle |e_{\min}|^2 \rangle$ refers to the minimum mean-squared error, which would be obtained if the filter weight vector $\mathbf{w}(n)$ equaled the Wiener solution \mathbf{w}_o . The misadjustment for the pipelined NEXT canceller with $LA = 1$ can be obtained using the formula in [2] as

$$M = \frac{c \cdot N}{2 - (P + 2K)c + K(K + 1)c^2} \quad (15)$$

where

$$c = \mu \cdot LA\sigma^2, \quad \sigma^2 = \langle |b(n)|^2 \rangle. \quad (16)$$

In (15), as K is increased from unity, the misadjustment would increase. However, in actual practice, the misadjustment does not change substantially as K varies. It will be shown via simulation that very large SU's are possible before the degradation in the adaptation accuracy becomes substantial. The summation in (9) increases the power of the correction term by a factor of LA . This is, however, not equivalent to increasing the step size μ by a factor of LA . This is because the summation in (9) is a lowpass filtered version of the product $e(n)\mathbf{b}^*(n)$, which would be closer to the expected value of the gradient and, thus, would result in a lower misadjustment than the case where the step size μ were increased. This justifies the use of *relaxed look-ahead* and provides an indication of the impact of LA on the convergence behavior.

TABLE II
PERFORMANCE OF PIPELINED NEXT CANCELLER WITH RELAXED LOOKAHEAD

125 Mb/s 32 CAP TPDDI Transceiver Operating over 100m Category 3 Cable
 $\alpha = 0.2$ $T' = T/3$ Pair (1,3) $SNR_0 = 15.9$ dB $P_e = 10^{-12}$ †

Measured Worst-case NEXT

Speed-up	Equalizer (42T)		NEXT Cancellor (32T)		
	D_2	D_1	LA	SNR_0 (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_0 for this probability of error is 27.13 dB.

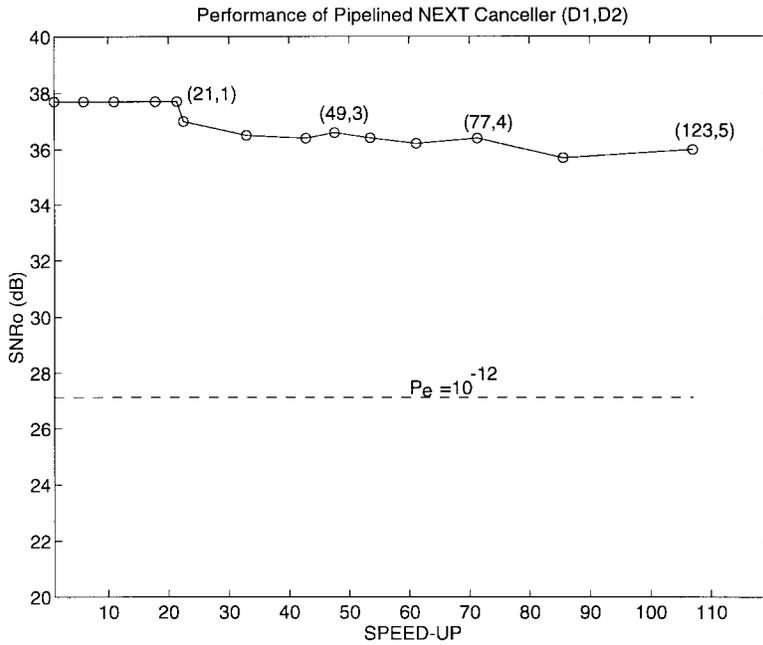


Fig. 6. SNR_0 versus speedup for the pipelined NEXT canceller (D_1, D_2).

IV. HARDWARE COMPLEXITY OF THE PIPELINED NEXT CANCELLER

As mentioned before, the pipelined NEXT canceller achieves substantial SU's with a low hardware overhead. In this section, we compare the hardware complexities of the serial and the pipelined architectures and quantify the overhead due to pipelining.

In Table I, we show the complexity of the serial (see Fig. 3) and the pipelined (see Fig. 4) architectures. It can be seen that the pipelined architecture requires $2N(LA - 1)$ additional adders. These adders are required to compute the summation in (9). In spite of being in the critical path, these adders do not present a throughput bottleneck as they perform a nonadaptive computation and can be realized in an equivalent transpose form. The number of multipliers for both the pipelined and serial architectures are the same, and hence, there is no increase. Comparing the number of *algorithmic latches*, we find from Table I that the pipelined architecture requires $4D_1 + 2N(D_2 + LA) - 4N$ additional latches. In a practical implementation, these algorithmic latches would be retimed [9] to generate *hardware latches*, which would be employed to pipeline various hardware operators. The process of retiming results in the number of hardware latches being higher than the algorithmic latches. However, the complexity of a latch is much smaller than that of a multiplier and an adder. Furthermore, due to the continuous flow of data through the architecture, these latches can be very simple (two

switches and two inverters). Hence, the overhead due to the latches can be considered minimal.

V. SIMULATION RESULTS AND DISCUSSION

In this section, we investigate the performance of the pipelined NEXT canceller of 125-Mb/s 32-CAP, and 155-Mb/s 64-CAP transceiver over category 3 cable. We will assume that one loop is utilized for each direction of transmission, as shown in Fig. 2, 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. 2 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 that are synchronized in frequency.

In order to investigate the performance of the 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. 2 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.

TABLE III
PERFORMANCE OF PIPELINED NEXT CANCELLER WITH RELAXED LOOKAHEAD

155.52 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 $P_e = 10^{-12}$ †

Measured Worst-case NEXT

Speed-up	Equalizer (42T)		NEXT Canceller (32T)		Margin † (dB)
	D_2	D_1	LA	SNR_o (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 30.03 dB.

- Step 3) The NEXT interferer is added at the input of the equalizer, as shown in Fig. 2; 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.

Fig. 5 shows the computer simulation results for the performance of the pipelined NEXT canceller with different D_1 and D_2 values. The solid line in Fig. 5 shows the convergence characteristic of the serial NEXT canceller. The dotted and dashed lines in Fig. 5 give the convergence characteristic of the pipelined NEXT canceller with $D_1 = 75$ and $D_2 = 3$ and $D_1 = 123$ and $D_2 = 5$, respectively. The pipelined canceller with $D_1 = 0$ and $D_2 = 1$ corresponds to the serial NEXT canceller. For all cases, we used the sufficiently small value of μ , which results in $2\mu\sigma^2 \cdot LA = 0.015$. 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. Comparing the performance of the serial and pipelined NEXT canceller, we see that it takes over two times more symbol iterations for the pipelined NEXT canceller ($D_1 = 123$, $D_2 = 5$) to converge to its steady state, and there is 1.8 dB degradation in SNR_o . It should also be noted that the symbol clock of the pipelined NEXT canceller can be clocked at a rate 107 times faster than that of the serial NEXT canceller. Thus, the absolute convergence speed with $D_1 = 123$ and $D_2 = 5$ is over 40 times faster than the serial NEXT canceller because of the SU factor of the pipelined canceller. Table II summarizes the performance results of the pipelined canceller shown in Fig. 5. In Table II, α stands for excess bandwidth, and SNR_i is the SNR at the input of the receiver. The first column in Table II gives the SU factor comparing to the serial NEXT canceller. The last column gives the margin, which is defined as

$$\text{margin} = SNR_o - SNR_{o,ref} \quad (17)$$

where $SNR_{o,ref}$ is a suitably chosen reference for the SNR at the input of the decision device. In Table II, we have chosen $SNR_{o,ref} = 27.13$ dB, which corresponds to the value of SNR_o that provides a probability of error of 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. 6, we plot the performance of the pipelined NEXT canceller with different values of SU for TPDDI application. The values of D_1 and D_2 for a certain SU depend on 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 SU is such that $T_{pipe} > T_{a1}$. In that case, this SU can be achieved with

$D_2 = 1$ and a sufficiently high value for D_1 . Notice also that the 125 Mb/s 32-CAP transceiver has still comfortable margins, even when the SU of the pipelined NEXT canceller is about 100. Table III summarizes the performance results of the pipelined NEXT canceller for 155 Mb/s ATM LAN application. It is noted that the SU factor of 107 can be achieved with 2 dB loss in the margin.

VI. 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 and minimal hardware overhead. 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.

ACKNOWLEDGMENT

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

REFERENCES

- [1] G. H. Im and J. J. Werner, "Bandwidth-efficient digital transmission over unshielded twisted pair wiring," *IEEE J. Select. Areas Commun.*, vol. 13, pp. 1643–1655, Dec. 1995.
- [2] N. R. Shanbhag and K. K. Parhi, "Relaxed look-ahead pipelined LMS adaptive filters and their application to ADPCM coder," *IEEE Trans. Circuits Syst.*, vol. 40, pp. 753–766, Dec. 1993.
- [3] ———, *Pipelined Adaptive Digital Filters*. Boston, MA: Kluwer, 1994.
- [4] 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. 1992.
- [5] G. H. Im, D. D. Harman, G. Huang, A. V. Mandzik, M.-H. Nguyen, and J. J. Werner, "51.84 Mb/s 16-CAP ATM LAN standard," *IEEE J. Select. Areas Commun.*, vol. 13, pp. 620–632, May 1995.
- [6] *af phy-0018,000*, ATM Forum, *Midrange Physical Layer Specification for Category 3 Unshielded Twisted Pair*, Sept. 1994.
- [7] *af phy-0047,000*, ATM Forum, *155.52 Mb/s Physical Layer Specification for Category 3 Unshielded Twisted Pair*, Nov. 1995.
- [8] *Commercial Building Telecommunications Wiring Standard*, EIA/TIA-568 Draft Stand., Dec. 1990.
- [9] C. Leiserson and J. Saxe, "Optimizing synchronous systems," *J. VLSI Comput. Syst.*, vol. 1, pp. 41–47, 1983.
- [10] K. K. Parhi, C.-Y. Wang, and A. P. Brown, "Synthesis of control circuits in folded pipelined DSP architectures," *IEEE J. Solid-State Circuits*, vol. 27, pp. 29–43, Jan. 1992.
- [11] A. P. Chandrakasan, S. Sheng, and R. W. Brodersen, "Low power CMOS digital design," *IEEE J. Solid-State Circuits*, vol. 27, pp. 473–484, Apr. 1992.