# Minimum Mean-Squared Error Echo Cancellation and Equalization for Digital Subscriber Line Transmission: Part II—A Simulation Study

DAVID W. LIN, SENIOR MEMBER, IEEE

*Abstract*—**In designing DSL (digital subscriber line) transceivers to support ISDN (integrated services digital network) Basic Access, a key issue is the choice of proper line codes. This paper reports a study to understand and compare the transmission performance of various line codes in the subscriber line environment and at a line rate around 160 kbits/s. The study is carried out by way of computation. The simulated transmission environment includes transhybrid echo, intersymbol interference, and near-end crosstalk from identical DSL transmission systems. The simulated transceiver incorporates an echo canceller and a decision-feedback equalizer which are optimal in the minimum mean-squared error sense as discussed in Part I [1]. Both the echo canceller and the decision-feedback equalizer have a tap spacing equal to the symbol period. The line codes investigated include five block codes which offer reduced baudrates and five other codes which do not reduce the baudrate. Under the conditions of this study, block codes do show advantage over nonblock codes, with the 2B1Q, a block code, yielding the best performance.**

## I. Introduction

THE design of full-duplex DSL (digital subscriber line) transceivers to support ISDN (integrated services digital network) Basic Access at a data rate of 144 kbits/s has attracted much recent study. Several main noise sources in DSL transmission are transhybrid echo, intersymbol interference (ISI), and near-end crosstalk (NEXT). Echo can be dealt with by an echo canceller (EC) and ISI by a decision-feedback equalizer (DFE), as shown in Fig. 1 which depicts a bare-bones DSL transceiver. Here we assume that the EC and the DFE are made of finite-length tapped-delay-line filters. If long and well-behaved EC's and DFE's can be designed and manufactured economically so as to eliminate virtually all echo and ISI, then NEXT will be left as the only dominant channel noise. For pair-wire channels, we know that the signal transmission loss and the NEXT coupling factor are, roughly speaking, both increasing functions of frequency. Thus, in the NEXT-dominant situation, it seems desirable to keep the signal spectrum to frequencies as low as possible. It is out of this concern that many have pursued the reduction of DSL symbol rate by block line coding.

However, one should note that, in reducing the baudrate, block coding may result in an increased number of signal levels and thereby increase the NEXT disturbance at lower frequencies. It is in transmission over a long enough distance that the reduction in transmission loss, arising from a reduced bandwidth due to block coding, may more than balance out the adverse effect of increased NEXT at lower frequencies to prove in a block code.

Several questions then naturally surface, some regarding the absolute performance of a code in a specific situation, and some the relative performance of different codes. First, within the economical reach of today's technology, how many taps can an EC and a DFE have? For a given line code, a channel and an NEXT level, what kind of transmission performance (in terms of signal-to-noise ratios, say) can we get with such filter lengths? Is such performance acceptable? To compare the performance of different codes, especially that of block codes to that of nonblock codes, we may ask the following questions. With the imperfect cancellation of echo and ISI by a finite-length EC and a finite-length DFE as above, is the residual noise now predominantly NEXT? Do block codes offer any advantage over nonblock codes?

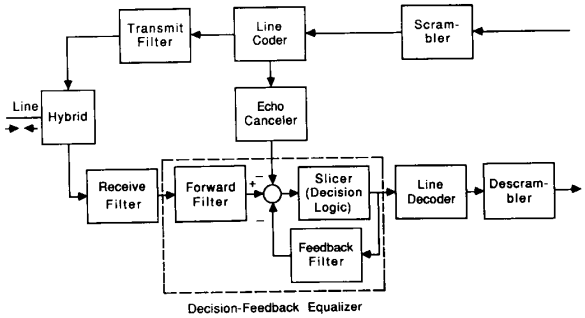


Fig. 1. A DSL transceiver structure with jointly adapted echo canceller and decision-feedback equalizer.

In this work, we try to answer the above questions, at least partially. We do this by computing the performance of MMSE (minimum mean-squared error) echo cancellation and decision-feedback equalization for various channels with various line codes using the theory presented in Part I [1]. From the numerical results we gain an understanding and a comparison of the transmission performance of the various line codes. We shall frequently refer to this study as a simulation. However, it should not be understood as a sample-by-sample simulation. Rather, the results have been calculated using the previously developed theory.

The line codes studied are listed in Table I. The coder inputs are assumed to have a 160 kbits/s rate, except in the case of the 3B2T code, where for divisibility by 3 we let the input rate be 162 kbits/s. The additional 16 or 18 kbits/s over the bare data rate of 144 kbits/s is to accommodate various housekeeping functions for the transmission system. The EC and the DFE in a transceiver are assumed to be optimized jointly and have a tap spacing equal to the symbol period. We also assume that the only significant sources of noise are echo, ISI, and self-NEXT (NEXT from identical DSL transmission systems). Code performance is compared in terms of the nominal SNR (signal-to-noise ratio) at the decision point, subject to the considerations discussed in [1].

Section II describes the conditions of study. Section III considers the case of a finite-length EC and a finite-length DFE. There we study the data from several angles. After that, one may be interested in knowing how much improvement in transmission performance is possible with a complete cancellation of echo, say, by a perfect infinite-length EC. This is considered in Section IV. Section V concludes the paper.

## II. Conditions of Study

From Fig. 1, one can see that there are a number of design and environmental variables in the system, in addition to the line code

Paper approved by the Editor for Channel Equalization of the IEEE Communications Society. Manuscript received November 24, 1987; revised July 4, 1988.

The author is with Bell Communications Research, Red Bank, NJ 07701.

IEEE Log Number 8932072.

0090-6778/90/0100-0039$01.00 

TABLE I
LINE CODES STUDIED IN THIS WORK

| | |
|---|---|
| (N)RZ[2]: | (*N*on-)*R*eturn-to-*Z*ero − coding 1 as +1 and 0 as −1. RZ has < 100% duty cycle; between two symbols the waveform returns to zero value. NRZ has 100% duty cycle. |
| Dicode[3]: | Both precoded and nonprecoded. Precoded dicode has also been called bipolar and AMI (*a*lternate *m*ark *i*nversion). |
| MDB[3]: | *M*odified *Duo-B*inary or class-4 partial-response code; both precoded and nonprecoded. |
| MS43[4,5]: | *M*onitoring-*S*tate 4 binary-to-3 ternary code. |
| MMS43[6]: | *M*odified MS43. |
| DI43[7]: | *D*ual-*I*nterleaved MS43. |
| 3B2T: | 3 *B*inary-to-2 *T*ernary, excluding 00. |
| 2B1Q: | 2 *B*inary-to-1 *Q*uarternary. |

and the lengths of EC and DFE. They include the characteristics of transmit and receive filters, the receive path's sampling time, and the characteristics of NEXT. Further, to minimize the echo power at the decision point, one may want to launch the transmitted signal at a different phase than the receive path's sampling phase. This relative phase between the transmit and the receive operations is thus yet another such variable. (Of course, we have the freedom of selecting this relative phase only for one of the two transceivers at the opposite ends of the same transmission line. The relative phase in the other transceiver is determined by the relative phase in this transceiver and the delay of the transmission line.) Each of these variables can be viewed as one dimension in a multidimensional space. In principle, then, for each line code we can search the space for the point yielding maximum SNR (in some statistical sense) over all subscriber lines. Then we can study and compare the code-wise optimal results for a conclusion. Unfortunately, this approach is obviously not feasible with our limited computing resources. We can only settle for a limited number of points from the space.

Table II describes the "coordinates" of the "points" for which the numerical results are discussed in more detail. Although more points have been examined, these points are chosen to show trends in, and to facilitate an interpretation of, the data obtained. Below we explain in more detail the conditions listed in the table.

The lengths of the simulated EC and DFE, as listed in the table, are fairly arbitrary; but they are at least about the limit one would consider for implementation with present-day technology. Note that the inputs to the EC and the DFE feedback filter are digital signals with only a very limited number of levels, while that to the DFE forward filter is continuous in amplitude. Hence, it is more difficult to implement a long DFE forward filter than to implement a long EC or a long DFE feedback filter, due to the amount of multiplication required with the former.

The coefficients of the two FIR filters (respectively, for the two choices of the combined characteristics of transmit and receive filters as described in the table) are as follows. For the 37-tap filter,

$$h(\ 1) = h(37) = 1.21746e - 04$$
$$h(\ 2) = h(36) = 2.36600e - 04$$
$$h(\ 3) = h(35) = -1.93939e - 08$$
$$h(\ 4) = h(34) = -9.08020e - 04$$
$$h(\ 5) = h(33) = -2.16002e - 03$$
$$h(\ 6) = h(32) = -2.41881e - 03$$
$$h(\ 7) = h(31) = 2.52423e - 09$$
$$h(\ 8) = h(30) = 5.33798e - 03$$
$$h(\ 9) = h(29) = 1.07020e - 02$$
$$h(10) = h(28) = 1.04884e - 02$$

TABLE II
A PARTIAL LIST OF SIMULATED CONDITIONS

- Length of echo canceler:
  1. 40.
  2. ∞ (complete echo cancelation).
- Respective lengths of DFE forward and feedback filters:
  1. 3, 30.
  2. 10, 30.
- Combined characteristics of transmit and receive filters:
  1. A 37-tap linear-phase FIR filter operating at 4 times the baudrate in series with a first-order analog high-pass filter. The digital filter has a 50% excess bandwidth compared to the Nyquist. (See text for filter coefficients.) Because it operates at 4 times the baudrate, symbols are transmitted with a 25% duty cycle. The analog high-pass filter has a pole at $P$ kHz, where $P$ depends on line code and length of DFE forward filter. For a 3-tap DFE forward filter, $P$ is 8 for (N)RZ, 16 for nonprecoded MDB and precoded and nonprecoded dicode, 32 for precoded MDB, 6 for MS43, MMS43, and DI43, 10.8 for 3B2T, and 12 for 2B1Q. For a 10-tap DFE forward filter, $P$ is 8 for precoded and nonprecoded dicode and MDB, 6 for MS43, MMS43, and DI43, 2.7 for 3B2T, and 2 for 2B1Q. In the latter case, there is no high-pass filtering for (N)RZ, or equivalently, $P = 0$ for that code.
  2. A 21-tap linear-phase FIR filter operating at 4 times the baudrate in series with a first-order analog high-pass filter. The digital filter has a 100% excess bandwidth. (See text for filter coefficients.) Symbols are transmitted with a 25% duty cycle. The analog high-pass filter is as described above for the 37-tap filter.
- Receiver sampling phase and cursor location: Obtained by oversampling the received signal by a factor of 16 and picking the instant yielding maximum decision-point SNR.
- Phase difference between transmitted signal and receiver sampling: (immaterial when echo is fully canceled)
  1. 0.
  2. 180° (half the baud interval).
- Power transfer function of NEXT coupling: $Kf^{3/2}$.
  1. High NEXT situation, $K = 10^{-13}$.
  2. Low NEXT situation, $K = 10^{-14}$.
- NEXT crosstalkers: 2-mile 26 gauge PIC lines.

$$h(11) = h(27) = -8.36297e - 10$$
$$h(12) = h(26) = -1.92577e - 02$$
$$h(13) = h(25) = -3.65345e - 02$$
$$h(14) = h(24) = -3.48144e - 02$$
$$h(15) = h(23) = 4.65799e - 10$$
$$h(16) = h(22) = 6.84674e - 02$$
$$h(17) = h(21) = 1.52835e - 01$$
$$h(18) = h(20) = 2.22815e - 01$$
$$h(19) = 2.50000e - 01;$$

and for the 21-tap filter,

$$h(\ 1) = h(21) = 2.44101e - 04$$
$$h(\ 2) = h(20) = 7.30800e - 04$$
$$h(\ 3) = h(19) = -6.90881e - 10$$
$$h(\ 4) = h(18) = -4.79607e - 03$$
$$h(\ 5) = h(17) = -1.37736e - 02$$
$$h(\ 6) = h(16) = -1.80984e - 02$$
$$h(\ 7) = h(15) = 7.31664e - 11$$
$$h(\ 8) = h(14) = 5.46730e - 02$$
$$h(\ 9) = h(13) = 1.38474e - 01$$
$$h(10) = h(12) = 2.17427e - 01$$
$$h(11) = 2.50000e - 01.$$

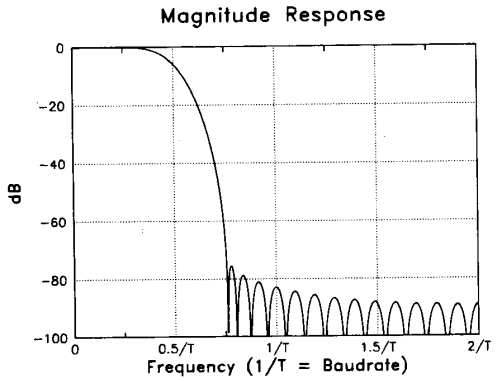


Fig. 2. Magnitude response of the 37-tap transmission filter.

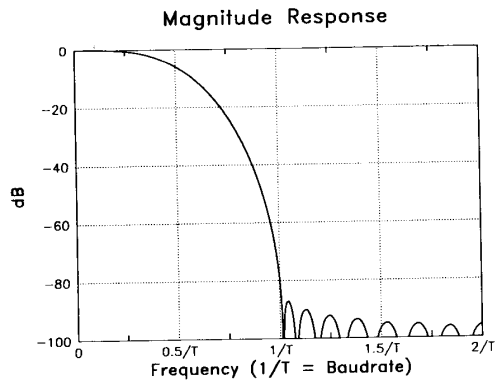


Fig. 3. Magnitude response of the 21-tap transmission filter.

Figs. 2 and 3 depict their respective magnitude responses. The filters were designed to have good stopband attenuation and generate little ISI for transmission over perfect channels. The poles of the first-order analog high-pass filters, as given in the table, were determined by choosing from a set of frequencies the ones yielding a better result where the set consisted of seven frequencies at, respectively, 0, 0.025, 0.05, 0.1, 0.15, 0.2, and 0.25 times the baudrate. (A pole at 0 is equivalent to having a flat unity gain, i.e., no high-pass filtering.) In a practical receiver, there would be an anti-aliasing analog low-pass filter. It was not included in our study because of an inherent low-pass filtering effect in the computation procedure (explained further later).

The reason why we only considered the combined characteristics of the transmit and the receive filters, rather than their characteristics separately, was that we were only concerned with echo, ISI, and self-NEXT. As have been pointed out in Part I [1], these signals all undergo a similar transmission mechanism in the sense that their sources all pass through the same line coding, the same transmit filtering, and the same receiver filtering. The only difference among them is in the different channels that stand between the transmit and receive filter pair. Therefore, only the combined and not the separate characteristics of the transmit and the receive filters need be considered.

The NEXT power transfer function is known [8] to be approximately given by $Kf^{3/2}$. A more exact characterization of it for certain cases is presented in [9]. Our parameter of $K = 10^{-14}$ (as listed in Table II) corresponds to the one-crosstalker curve shown therein and that of $K = 10^{-13}$ the 49-crosstalker curve. Although the example conditions of [9] are different from our simulated ones, we expect these parameter values to have at least an order-of-magnitude significance. The assumption of all crosstalkers being 2-mile 26 gauge PIC (polyethylene-insulated cable) lines, as stated in the table, can

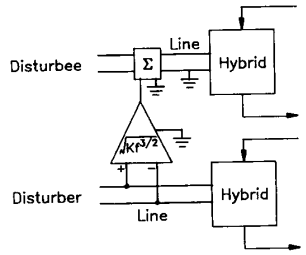


Fig. 4. Circuit diagram for calculating the NEXT spectrum.

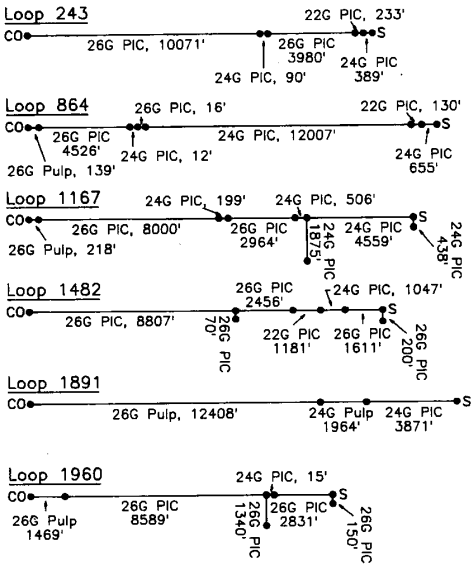


Fig. 5. Composition of simulated loops.

certainly be argued as unrealistic. It was made to provide a uniform operating environment for ease of performance comparison. Fig. 4 explains how the crosstalk spectrum was calculated. The crosstalk disturbance was then simulated by adding the power spectrum of the crosstalk noise to the power spectrum of the signal.

A more realistic simulation of the NEXT effects should consider the phase characteristics of the NEXT paths as well. This is because the self-NEXT disturbance is possibly cyclostationary due to the periodic nature of the digital waveforms in the crosstalking systems. However, such consideration could make the study manyfold more complex. By considering only the magnitude characteristics and not the phase, we effectively randomize the noise phase to make it stationary.

In this study, we considered six loops (subscriber lines) from the I-MATCH.1 database [10]. The loop compositions are shown in Fig. 5. Loosely speaking, these loops are some of the electrically longest ones from the 1983 Bell System Loop Survey and hence present the most difficult transmission conditions. It is this "last mile" that all of the line codes are striving for. Since these loops are, again loosely speaking, electrically similar to 26 gauge PIC lines of lengths about 2.5–3 mi, we also considered two artificial loops of uniform 26 gauge PIC lines with lengths 2.5 and 3 mi, respectively, for comparison. The simulated hybrid network has the configuration shown in Fig. 6 where the RC network, composed of a 1 KΩ resistor in parallel with the serially connected 0.24 μF capacitor and 136 Ω resistor, is a compromise balance network for the transmission line, and we have rather arbitrarily modeled the transformer as two 200 mH perfect inductors coupled together by a 200 mH mutual inductance. The loops were assumed to be at 70° F. Their frequency responses

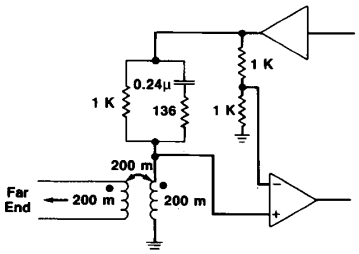


Fig. 6. An example hybrid network.

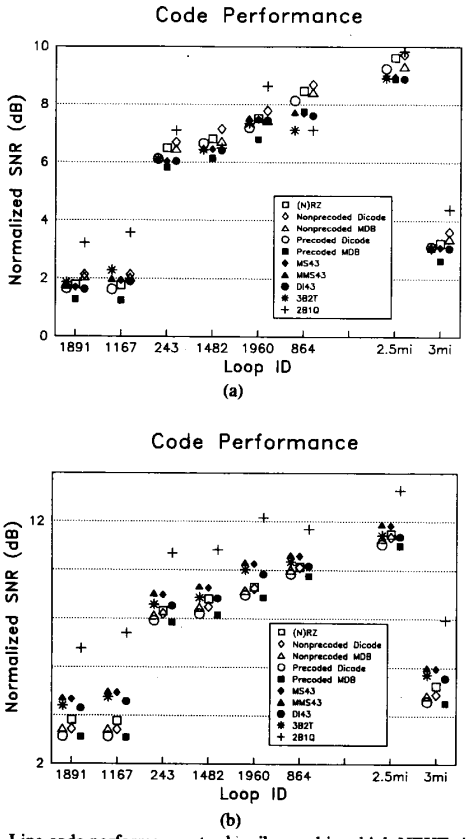


Fig. 7. Line code performance at subscriber end in a high NEXT situation. Transceiver incorporates the 37-tap transmission filter, a 40-tap EC, and a 30-tap DFE feedback filter. (a) Length of DFE forward filter = 3. (b) Length of DFE forward filter = 10.

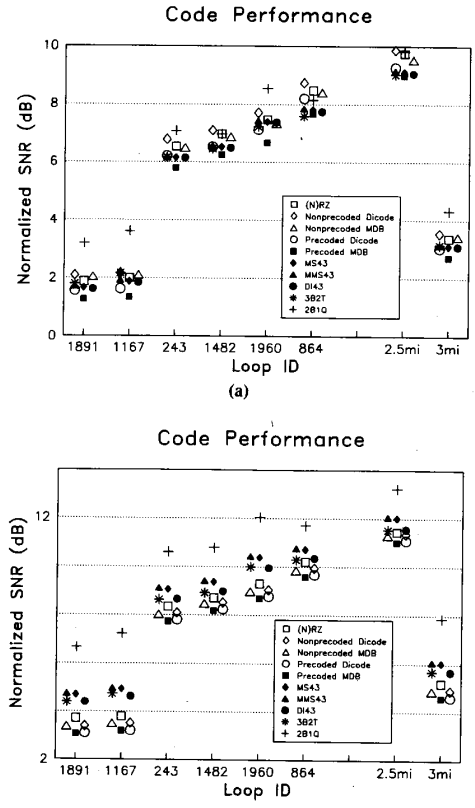


Fig. 8. Line code performance at subscriber end in a high NEXT situation. Transceiver incorporates the 21-tap transmission filter, a 40-tap EC, and a 30-tap DFE feedback filter. (a) Length of DFE forward filter = 3. (b) Length of DFE forward filter = 10.

were calculated by quadratically interpolating the data provided in [11]. The responses were calculated from dc up to twice the baudrate for each different signaling frequency (160, 120, 108, and 80 kHz), in 1025 equally spaced frequency samples. The frequency-sampling rate of 512 points per baudrate translates into 512-sample impulse responses in the time domain. Beyond twice the baudrate the frequency responses were assumed negligible—an assumption equivalent to that of having a brickwall low-pass filter with a cutoff at twice the baudrate. Because of this we did not include an anti-aliasing low-pass filter in the simulation.

## III. Partial Echo Cancellation

Now we turn to the numerical results for the case of partial echo cancellation (length of EC equal to 40). Consider the high NEXT situation first, since it presents the more difficult transmission environment. Figs. 7 and 8 depict some results where for ease of performance comparison we have normalized the SNR values with respect to that required for a $10^{-6}$ symbol error rate. (See Table I in [1].) In other words, we have subtracted 13.5 dB from the (*N*)RZ performance, 16.8 dB from precoded dicode's performance, 17.9 dB from the MS43 performance, and so on; so that, if the noise is signal-independent additive Gaussian, a 0 dB normalized SNR corresponds to a $10^{-6}$ symbol error rate, a 0.75 dB roughly to $10^{-7}$, and a 1.4 dB roughly to $10^{-8}$, and so on. Of course, a real adaptive transceiver will never attain the optimal MMSE performance. In actual designs, one also has to allow some SNR margin for imperfect modeling of the transmission environment. However, in this paper we shall not be concerned with such considerations which are more related to the absolute performance. Rather, we concentrate more on the relative performance. Specifically, four aspects of the data will be examined in more detail. First, we observe the relative performance of the various line codes for each DFE length. Second, we examine the performance improvement as the DFE length increases. Then, after presenting the data for the low NEXT situation, we compare the relative code performance in the two different NEXT environments. And we look into the composition of the residual noise in different instances.

In Figs. 7 and 8, the six loops have been ordered according to ascending SNR performance. The transceivers considered are the ones at the subscriber ends of the loops. For these data, the phase difference between the transmitted signal and the receiver sampling was 0. The SNR performance obtained from assuming a 180° phase difference is practically the same. (We also considered the combination of a shorter EC and a shorter DFE feedback filter, namely, 25 and 15 taps each, and found the above "phase indifference" to be also true in most cases.) It is interesting to note that the 21-tap

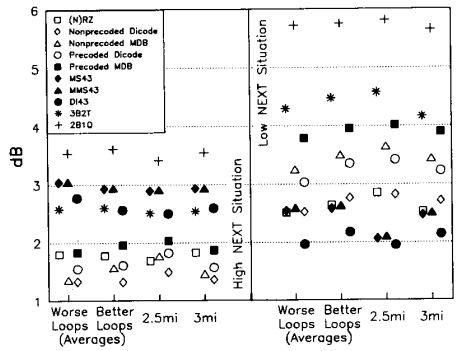


Fig. 9. Improvement of SNR in dB from using a 3-tap to using a 10-tap DFE forward filter, in association with the 21-tap transmission filter, a 40-tap EC, and a 30-tap DFE feedback filter.

filter, in spite of its wider bandwidth, yielded a performance at least comparable to that of the 37-tap filter. This is also found to be quite true for the low NEXT situation. Hence we concentrate our following discussion on the results obtained using the 21-tap filter, unless otherwise noted.

In Fig. 8(a), which shows results from using a 3-tap DFE forward filter, we see that the codes seem to group into several clusters according to their performance. The 2B1Q code usually scores the highest and the precoded MDB the lowest, with all others in between. For the two worst loops (1891 and 1167), 2B1Q has an edge of about 1.5 to 2 dB over the other codes. For the better loops (243, 1482, 1960, and 864), the gap is narrower. A similar phenomenon is observed for the two artificial loops. In summary, the data seem to indicate that, in this case, the reduction in baudrate offered by the 2B1Q code gives it a clear edge over those transmitting at the original bit rate, while that by the 3B2T and the 43 codes (MS43, MMS43, and DI43) does not.

In Fig. 8(b), the case of a 10-tap DFE forward filter, we see that 2B1Q shows an even better relative performance. But now we also see 3B2T and the 43 codes performing better than those providing no baudrate reduction.

Overall, the examples given in Fig. 8 (and 7, too) also seem to indicate that the 3B2T code, despite its lower baudrate, may not offer any advantage over the 43 codes. Moreover, the (*N*)RZ code yields rather comparable performance to the nonprecoded dicode and MDB. This latter phenomenon seems to be attributable to the working of the DFE forward filter: The MMSE DFE forward filter provides a greater capability of equalization than the simple partial-response filtering in nonprecoded dicode and MDB, and thus overrides the function of the latter.

It is instructive to examine the performance improvement, for each code, from using a 3-tap DFE forward filter to using a 10-tap one. The left half of Fig. 9 shows it graphically, where we have considered worse (1891 and 1167) and better (243, 1482, 1960, and 864) loops separately and taken averages of the dB improvement for each loop group. As we see, in this high NEXT situation, the block codes enjoy a greater average improvement than the others, with the 2B1Q at the crown.

In summary, the reduction in baudrate by block coding can yield a significant performance dividend in the high NEXT situation at a transceiver complexity as considered.

Consider now the low NEXT situation, for which Fig. 10 presents some results. In the case of a 3-tap DFE forward filter, the 43 codes are seen to outperform all the other ones, and the 2B1Q and the 3B2T yield no apparent advantage over the codes offering no baudrate reduction. Going to the longer DFE forward filter (10 taps), we see a much greater improvement in the performance of the 2B1Q code in comparison to the others, and the 3B2T code now performs similarly to the 43 codes.

The code-by-code SNR improvement in going from a 3-tap DFE forward filter to a 10-tap one, for the low NEXT situation, has been

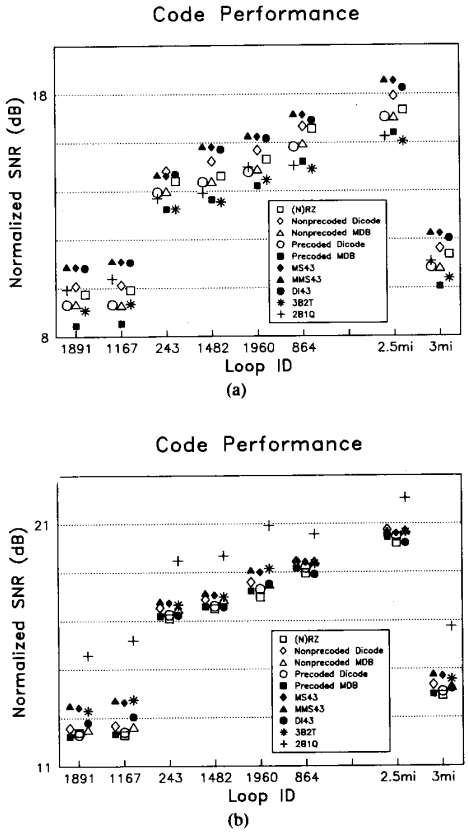


Fig. 10. Line code performance at subscriber end in a low NEXT situation. Transceiver incorporates the 21-tap transmission filter, a 40-tap EC, and a 30-tap DFE feedback filter. (a) Length of DFE forward filter = 3. (b) Length of DFE forward filter = 10.

plotted in Fig. 9. As in the high NEXT situation, 2B1Q again tops. The 3B2T then follows. But the improvement in the 43 codes' performance now lags behind that of the nonblock codes. This last point will be commented on when we consider the difference between high NEXT and low NEXT results next.

We now compare the high NEXT results to the low NEXT ones. At the beginning, we mentioned that the transmission-to-NEXT ratio, roughly speaking, decreases with increasing frequency. Now that block codes yield lower baudrates, one naturally may expect to see an accentuation of block codes' relative performance (compared to nonblock codes) with an increase in NEXT level. Comparison of Figs. 8(a) to 10(a) and 8(b) to 10(b) reveals that this is approximately true, except for the 43 codes in the case of a 3-tap DFE forward filter. This out-of-line behavior of the 43 codes seems to be due (at least) partly to the particular transmission filtering characteristics we used. It also results in the low improvement of the 43 codes' performance from using a 3-tap DFE forward filter to using a 10-tap one in the low NEXT situation, as seen in Fig. 9. For better visualization, we show in Fig. 11 the SNR degradation from operating in a low NEXT environment to operating in a high NEXT environment, for each code. As in Fig. 9, we have taken averages of the dB degradation for the worse and the better loops. A smaller degradation signifies a better relative performance in the high NEXT environment compared to the low NEXT environment. As seen, except for the 43 codes in the case of a 3-tap DFE forward filter, block codes normally have a smaller degradation than nonblock ones.

We noted that, by the fact that the transmission-to-NEXT ratio is a decreasing function of frequency, for transmission over a reasonably long distance it would seem desirable to keep the signal spectrum

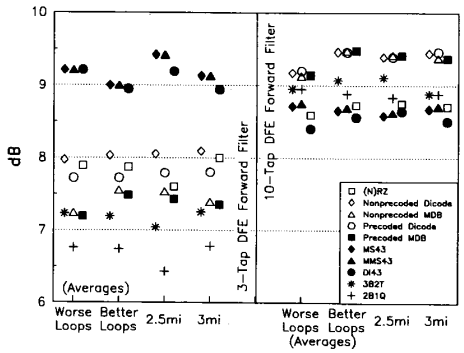


Fig. 11. Degradation of SNR in dB from operating in a low NEXT environment to operating in a high NEXT environment. Transceiver incorporates the 21-tap transmission filter, a 40-tap EC, and a 30-tap DFE feedback filter.

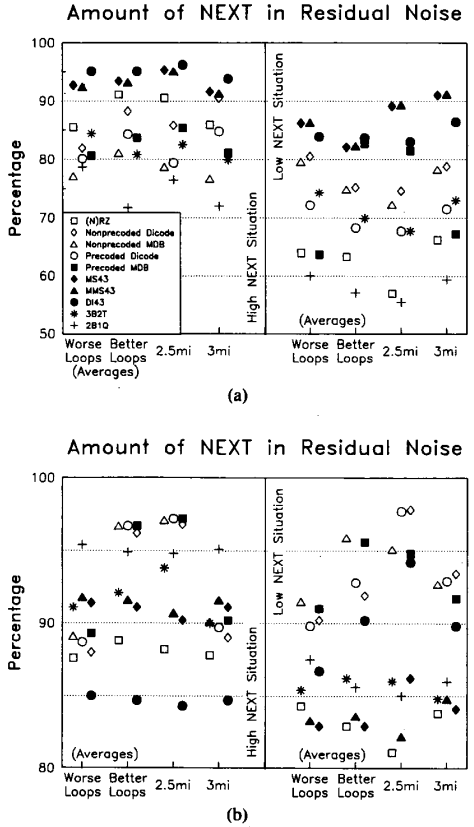


Fig. 12. Amount of NEXT in total residual noise. Transceiver incorporates the 21-tap transmission filter, a 40-tap EC, and a 30-tap DFE feedback filter. (a) Length of DFE forward filter = 3. (b) Length of DFE forward filter = 10.

to as low frequencies as possible with block coding, should NEXT dominate the transmission noise. In the simulation results presented above, we do see that block codes perform on a par with or better than nonblock ones in most instances. Thus, it would be interesting to examine to what percentage NEXT constitutes a portion of the total decision-point noise. The numbers would also show how successful the EC and the DFE have been dealing with echo and ISI, relative to NEXT. Fig. 12 summarizes the result. We see that, in all cases, NEXT constitutes a major part of the total residual noise. However,

TABLE III
AVERAGE PERCENTAGE OF RESIDUAL ECHO IN TOTAL RESIDUAL NOISE. TRANSCEIVER INCORPORATES THE 21-TAP TRANSMISSION FILTER, A 40-TAP EC, AND A 30-TAP DFE FEEDBACK FILTER. (a) WORSE LOOPS (1891 AND 1167). (b) BETTER LOOPS (243, 1482, 1960, AND 864)

| Condition → | 3-Tap DFE Forward Filter | | | | 10-Tap DFE Forward Filter | | | |
|---|---|---|---|---|---|---|---|---|
| | High NEXT | | Low NEXT | | High NEXT | | Low NEXT | |
| ↓ Code | (a) | (b) | (a) | (b) | (a) | (b) | (a) | (b) |
| (N)RZ | 2.7 | 3.0 | 12.3 | 11.5 | 2.6 | 3.4 | 6.3 | 7.8 |
| Nonprecoded Dicode | 0.9 | 0.7 | 2.6 | 2.3 | 0.5 | 0.5 | 1.5 | 1.7 |
| Nonprecoded MDB | 1.7 | 1.4 | 2.8 | 2.4 | 0.5 | 0.6 | 1.9 | 1.8 |
| Precoded Dicode | 0.8 | 0.6 | 2.0 | 1.6 | 0.3 | 0.4 | 1.6 | 1.7 |
| Precoded MDB | 0.8 | 0.7 | 3.1 | 1.1 | 0.4 | 0.6 | 2.0 | 1.9 |
| MS43 | 0.9 | 0.8 | 1.5 | 1.8 | 3.0 | 4.0 | 8.1 | 8.5 |
| MMS43 | 0.9 | 0.8 | 1.5 | 1.7 | 2.8 | 3.8 | 7.9 | 8.2 |
| DI43 | 0.9 | 1.0 | 3.3 | 3.5 | 6.1 | 6.9 | 5.5 | 4.4 |
| 3B2T | 4.5 | 4.0 | 8.0 | 7.4 | 2.4 | 3.1 | 8.1 | 9.4 |
| 2B1Q | 5.7 | 5.0 | 10.0 | 9.5 | 1.7 | 2.3 | 5.3 | 7.1 |

TABLE IV
AVERAGE IMPROVEMENT OF SNR IN dB FROM PARTIALLY CANCELLING THE ECHO BY A 40-TAP EC TO COMPLETELY CANCELLING THE ECHO TRANSCEIVER INCORPORATES THE 21-TAP TRANSMISSION FILTER AND A 30-TAP DFE FEEDBACK FILTER. (a) WORSE LOOPS (1891 AND 1167). (b) BETTER LOOPS (243, 1492, 1960 AND 864)

| Condition → | 3-Tap DFE Forward Filter | | | | 10-Tap DFE Forward Filter | | | |
|---|---|---|---|---|---|---|---|---|
| | High NEXT | | Low NEXT | | High NEXT | | Low NEXT | |
| ↓ Code | (a) | (b) | (a) | (b) | (a) | (b) | (a) | (b) |
| (N)RZ | 0.2 | 0.2 | 0.7 | 0.8 | 0.3 | 0.3 | 0.7 | 0.8 |
| Nonprecoded Dicode | 0.0 | 0.0 | 0.1 | 0.1 | 0.0 | 0.0 | 0.1 | 0.1 |
| Nonprecoded MDB | 0.1 | 0.1 | 0.1 | 0.1 | 0.0 | 0.0 | 0.1 | 0.1 |
| Precoded Dicode | 0.0 | 0.0 | 0.1 | 0.1 | 0.0 | 0.0 | 0.1 | 0.1 |
| Precoded MDB | 0.0 | 0.0 | 0.1 | 0.1 | 0.0 | 0.0 | 0.1 | 0.1 |
| MS43 | 0.0 | 0.0 | 0.1 | 0.1 | 0.2 | 0.2 | 0.4 | 0.5 |
| MMS43 | 0.0 | 0.0 | 0.1 | 0.1 | 0.2 | 0.2 | 0.4 | 0.5 |
| DI43 | 0.0 | 0.0 | 0.2 | 0.2 | 0.3 | 0.4 | 0.4 | 0.3 |
| 3B2T | 0.3 | 0.3 | 0.8 | 0.5 | 0.2 | 0.2 | 0.6 | 0.6 |
| 2B1Q | 0.3 | 0.3 | 0.5 | 0.5 | 0.2 | 0.2 | 0.7 | 0.9 |

we do not seem to be able to draw any nontrivial conclusion from the data, such as a relation between the relative performance of codes and the percentage contribution of NEXT in the residual noise.

In conclusion, within the accuracy of our study, the 2B1Q code clearly stands out to be the winner among all investigated. Other block codes (3B2T, MS43, MMS43, and DI43) do perform on a par with or better than nonblock codes in most instances, but generally not as well as the 2B1Q. The three 43 codes have always performed very similarly. And the 3B2T code, despite its lower baudrate, does not seem to offer any edge over the 43 codes. For the nonblock codes, it does not seem that we have to go beyond the simplistic (*N*)RZ, if the MMSE performance is the only concern.

## IV. COMPLETE ECHO CANCELLATION

In the previous section, we have presented results from a case of partial echo cancellation with a finite-length EC. One may be interested in knowing how much improvement in transmission performance is possible if we could cancel the echo fully, say, by a perfect and possibly infinite-length EC. This we address in this section.

To start, let us look at the percentage contribution of residual echo in the total residual noise for some examples discussed in the previous section. Table III lists the figures. As seen, residual echo constitutes a rather insignificant portion of the total residual noise in most cases. It can thus be expected that a complete cancellation of echo will do little benefit to the SNR performance in these cases. This is confirmed by Table IV, which gives the SNR improvement in going from a 40-tap EC to an infinitely long one in various cases. Before examining Table IV in more detail, note that Table III also shows that, compared to the residual echo in the low NEXT situation, the residual echo in the high NEXT situation constitutes a smaller

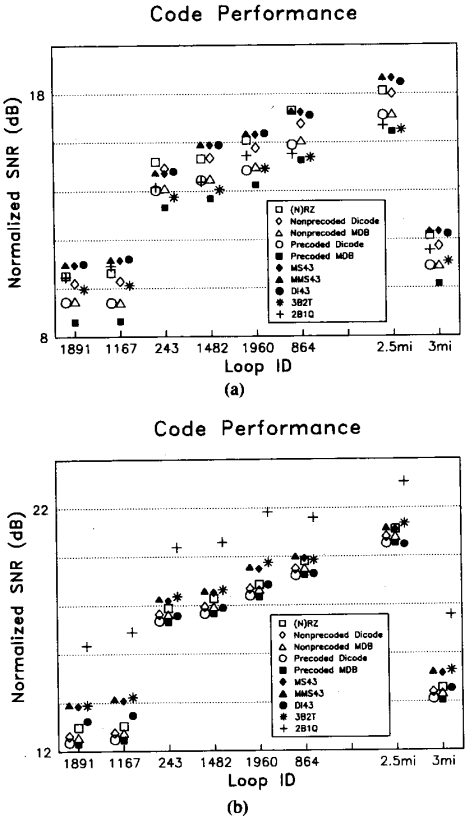


Fig. 13. Line code performance at subscriber end in a low NEXT situation, with complete echo cancellation. Transceiver incorporates the 21-tap transmission filter and a 30-tap DFE feedback filter. (a) Length of DFE forward filter = 3. (b) Length of DFE forward filter = 10.

portion of the total residual noise in most cases. This is intuitively reasonable. The only exception is in the case of the DI43 code with a 10-tap DFE forward filter, which is a result of both the code characteristics and the characteristics of the transmission filtering considered.

Turning back to Table IV, we consider the high NEXT situation first. By the rather small SNR improvement shown in the table, we conclude that, for this environment, a 40-tap EC is close to being as good as a perfect one. Further, block codes have improved somewhat more than all nonblock ones except the nonpartial-response (*N*)RZ, thus giving the former somewhat more upper hand over the latter. The observations and conclusions made in the previous section regarding the high NEXT situation stay substantially the same.

Now consider the low NEXT situation, where full echo cancellation leads to more performance improvement. Here, again, block codes and (*N*)RZ benefit more. It is interesting to note that the greatest improvement goes to the codes with a flat power spectrum (the (*N*)RZ, the 3B2T, and the 2B1Q). Fig. 13 shows the low NEXT, full echo cancellation, results graphically.

## V. Conclusion

We have reported some results from a computational study of the DSL transmission performance of ten line codes. Of all the codes investigated, the 2B1Q stands out to be the best performer considering all the example conditions examined. The four other block codes (3B2T, MS43, MMS43, and DI43) have also performed on a par with or better than the five nonblock codes in most instances. The three 43 codes have performed very similarly much of the time. The 3B2T code, despite its lower baudrate compared to the 43 codes, does not seem to offer any advantage over the latter in MMSE performance. Among the five nonblock codes, the simplistic (*N*)RZ turns out to be at least about as good as any other where the others are some partial-response codes, both precoded and nonprecoded.

As discussed in Section II, we have tried out only a very limited subset of the whole transceiver design space and the possible operating environments. In addition, we discussed in [1] that the measurement of transmission performance by SNR is most meaningful when signal and noise are mutually statistically independent. Yet in the case of precoded partial-response codes and many block codes (e.g., the 43 codes), because of the correlation among successive code symbols, we do have some amount of correlation between the signal and the ISI component of the noise. Thus the interpretation and application of these results must be done with care.

## References

[1] D. W. Lin, "Minimum mean-squared error echo cancellation and equalization for digital subscriber line transmission: Part I—Theory and computation," *IEEE Trans. Commun.*, vol. 38, pp. 31–38, Jan. 1990.

[2] J. C. Bellamy, *Digital Telephony*. New York: Wiley, 1982.

[3] P. Kabal and S. Pasupathy, "Partial-response signaling," *IEEE Trans. Commun.*, vol. COM-23, pp. 921–934, Sep. 1975.

[4] P. A. Franaszek, "Sequence-state coding for digital transmission," *Bell Syst. Tech. J.*, vol. 47, pp. 143–157, Jan. 1968.

[5] G. L. Cariolaro and G. P. Tronca, "Spectra of block coded digital signals," *IEEE Trans. Commun.*, vol. COM-22, pp. 1555–1564, Oct. 1974.

[6] ITT Telecom, Raleigh, North Carolina, USA, and Standard Elektrik Lorenz AG, Stuttgart, W. Germany, "Digital subscriber loop carrier network interface specification," ANSI/ECSA Telecommunications Working Group T1D1.3 Contribution 85-061, Apr. 1985.

[7] P. E. Fleischer and L. Wu, "Novel block codes for DSL applications," *IEEE Global Telecommun. Conf. Rec.*, 1985, pp. 1329–1334.

[8] A. J. Gibbs and R. Addie, "The covariance of near end crosstalk and its application to PCM system engineering in multipair cable," *IEEE Trans. Commun.*, vol. COM-27, pp. 469–477, Feb. 1979.

[9] Bell Communications Research, Inc., "ISDN basic access digital subscriber lines," Tech. Ref. TR-TSY-000393, Issue 1, May 1988.

[10] Bell Communications Research, Inc., "I-MATCH.1 loop characterization data base," Special Rep. SR-TSY-000231, Issue 1, June 1985.

[11] "Digital data system channel interface specification," Bell Syst. Tech. Ref., PUB 62310, AT&T, Sep. 1983. The publication seems to lack a table for 19 gauge PIC cable at 0°F, which can be found in "Circuit switched digital capability network access interface specifications," Tech. Ref. TR-880-22135-84-01, Issue 1, Bell Commun. Res., Inc., July 1984. However, certain other tables in the latter reference contain errors or missing digits.

**David W. Lin** (S'78–M'82–SM'88), for a photograph and biography, see this issue p. 38.