

Design Issues in the ISDN U-Interface Transceiver

DAVID G. MESSERSCHMITT, FELLOW, IEEE

Abstract—The design challenges faced in the integrated circuit realization of the basic customer access U-interface transceiver at 144 kbits/s in the integrated services digital network (ISDN) are summarized. Given the cost and performance objectives, this represents a very challenging design problem from an algorithmic and technology point of view. This survey paper describes the alternative design approaches concentrating on algorithmic issues as opposed to circuit design issues, in the context of the echo cancellation (EC) method of full duplex data transmission. Particular emphasis is given to the areas of echo cancellation, equalization, line code selection, and timing recovery.

I. INTRODUCTION

THE Integrated Services Digital Network (ISDN) provides telephony customers with a set of standard interfaces that will support a variety of services including voice, video, and various types of data services. The most common interface to the individual customer will be the basic interface, that provides for a 144 kbit/s full duplex bit stream divided into two “B channels,” each at a rate of 64 kbits/s, and a “D channel” at a rate of 16 kbits/s.

The basic interface is described in Section II, including the U-interface that provides the full duplex bit stream over a two-wire subscriber loop. Sections III–IX summarize the design approaches available and the environmental conditions that must be overcome in the U-interface. Finally, Section X attempts to place these issues in perspective, and gives the author’s opinion as to the best combination of design approaches available.

II. THE S- AND U-INTERFACES

There are several interfaces that provide the ISDN basic access at the physical layer. Of these, the two most important are the S-interface for transmission within the customer premises (residential, business, or from a private-branch exchange to the telephone instrument), and the U-interface for connection of the serving center to the customer premises. These two interfaces are illustrated in Fig. 1 for a typical residential application. The U-interface describes the full duplex data signal on the two-wire (single wire pair) subscriber loop between central office line termination (LT) and customer premise network termination (NT). The S-interface uses separate wire pairs

for the two directions of transmission between the NT and terminal equipment (TE).

The U-interface is more challenging to implement, due to the problem of separating the two directions of transmission and the environmental conditions. Therefore, in the sequel we concentrate on the U-interface design.

III. ENVIRONMENT FOR THE U-INTERFACE TRANSCEIVER

The U-interface transceiver must contend with the non-idealities of the existing subscriber loop plant, including gauge changes and bridged taps, as well as accommodate the maximum distance to the central office in the presence of impairments like impulse noise and crosstalk [1]. It must also be inexpensive, requiring a monolithic realization without external precision components or trimming to improve component accuracy.

The frequency-dependent loss of the cable is dependent on the gauge mix of the cable, the length of the cable, and the presence in some administrations of open-circuited cable pairs bridged onto the main cable pair, called *bridged taps*. The circuit configuration corresponding to a main cable pair and a single bridged tap is shown in Fig. 2. It is common to have two or more bridged taps of unknown length.

The attenuation introduced by the circuit configuration of Fig. 2 is shown in Fig. 3 for a long main cable without a bridged tap and with single bridged taps of different lengths. (Figs. 3–7 were generated by a transmission line simulation program [2].) In the absence of any bridged taps the attenuation increases rapidly with frequency, and, of course, also increases rapidly near dc due to the transformer. The response of the cable (to an isolated transmitted pulse of duration equal to 50 percent of a baud interval at 160 kbaud/s) for several cable lengths is shown in Fig. 4. The cable acts like a low-pass filter with pass-band attenuation, attenuating and dispersing the pulse as cable length increases. The effect of the transformer is to add a long negative tail to the response, which is necessary in order for the response pulse to have a net area of zero.

Returning to Fig. 3, one effect of a bridged tap is to load the cable, providing an overall additional attenuation. For example, a very long bridged tap would have an input impedance equal to its characteristic impedance, which will increase the attenuation by 3.5 dB at all frequencies. However, to minimize the effect of bridged taps at voiceband frequencies, the design of the network has generally placed a limit on the total length of all bridged taps (this limit is 6 kft or 1.8 km in the U.S.). These

Manuscript received October 4, 1985; revised April 5, 1986. This work was supported by Advanced Micro Devices, Harris Semiconductor, Intel Corporation, National Semiconductor, Fairchild Semiconductor, and Bell Communications Research, with a matching grant from the University of California Microelectronics and Computer Opportunities (MICRO) Program.

The author is with the Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA 94720.

IEEE Log Number 8609376.

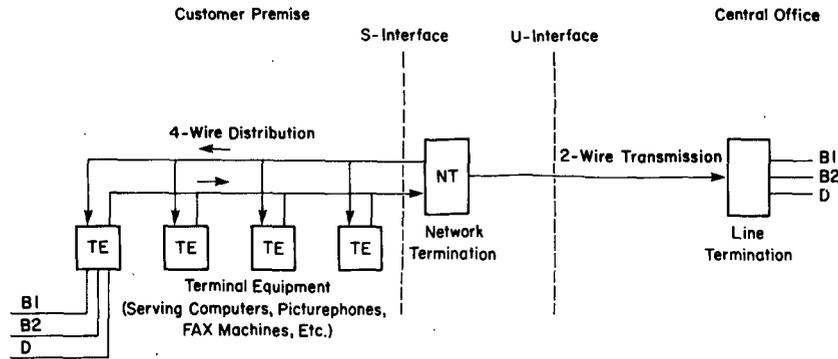


Fig. 1. Typical application of the S-interface and U-interface transceivers on an ISDN digital subscriber loop. The S-interface is in the customer premises, and the U-interface provides transport to the serving center or central office.

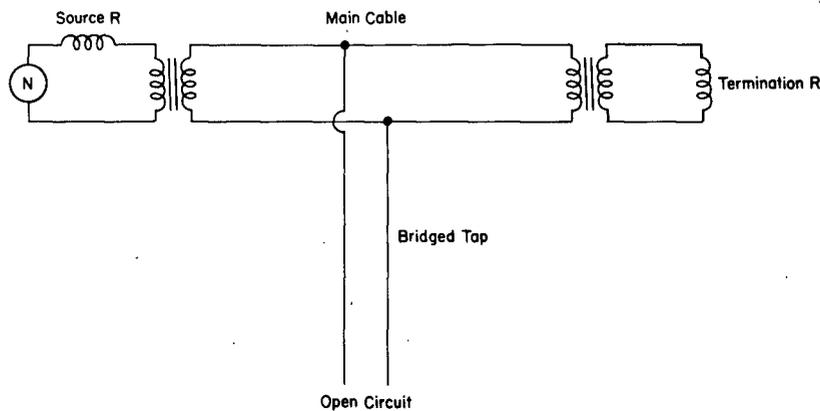


Fig. 2. Circuit configuration of a cable and a single bridged tap.

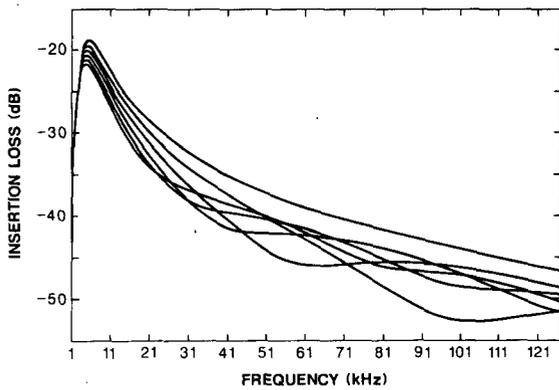


Fig. 3. 135 Ω insertion loss of an 18 kft (5.5 km) 24 AWG cable with one 22 AWG bridged tap with lengths 0, 1.5, 2.5, 3.5, 4.5, and 5.5 kft. Each end has a 5 mH transformer.

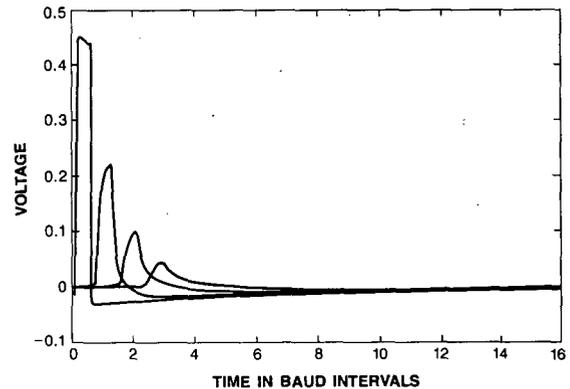


Fig. 4. Response of cable to a transmitted 50 percent duty cycle 1 V pulse (baud rate 160 kHz) for 24 AWG cable with lengths 0.1, 1.0, 2.0, and 3.0 km.

shorter bridged taps have a more complicated effect than just attenuation, since some of the signal energy traveling down the main cable pair and diverted into the bridged tap is reflected back to the main cable pair with an attenuation and delay that depends on the length of the tap. At some frequencies, this delayed and attenuated reflection causes destructive interference, resulting in an attenuation greater than 3.5 dB, and at other frequencies interferes in a constructive manner, reducing the attenuation. The frequency-selective dips in the overall response are evident in

Fig. 3. Shorter bridged taps result in a dip at higher frequencies, whereas the longest allowed bridged tap (6 kft) results in a dip at about 27 kHz. Generally, the additional attenuation introduced by the bridged tap is greater for the shorter bridged taps (higher frequency dips) since these taps will have a lower attenuation to the reflected energy.

In the time domain the bridged tap can result in "ghost pulses" that are delayed relative to the main received pulse and result in postcursor intersymbol interference. The pulse responses similar to Fig. 4 in the presence of

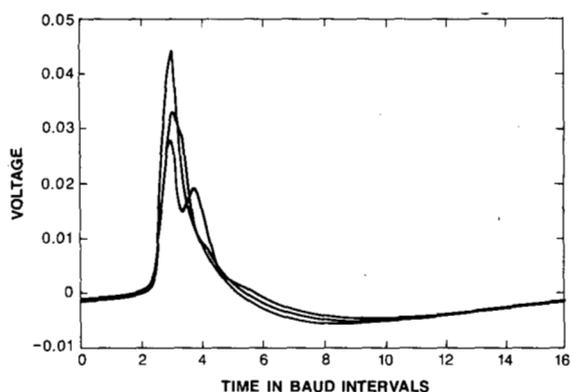


Fig. 5. Response of cable to a transmitted 50 percent duty cycle 1 V pulse (baud rate 160 kHz) for 3 km 24 AWG cable with bridged taps of length 0.025, and 0.5 km.

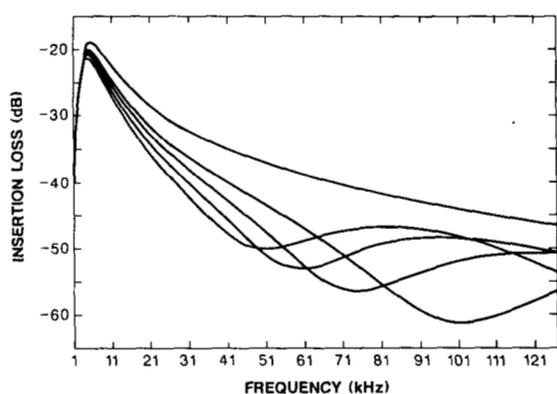
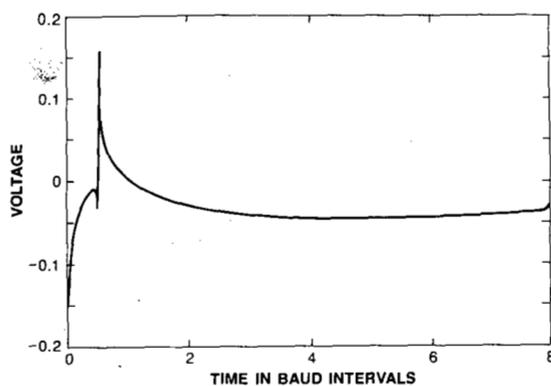


Fig. 6. 135 Ω insertion loss of an 18 kft (5.5 km) 24 AWG cable with two bridged taps of equal lengths 0, 1.5, 2, 2.5, and 3 kft. The two bridged taps are spaced at 6 kft and 12 kft on the cable.

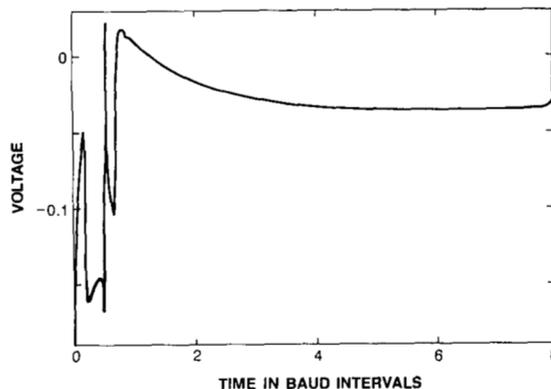
bridged taps of several lengths are shown in Fig. 5. As the bridged tap length increases, it causes an increasing attenuation of the maximum amplitude of the pulse as well as broadening the pulse. For the longest bridged tap shown, a second distinct "ghost pulse" becomes evident.

When two or more bridged taps are present, the effects are approximately additive. However, since the total length of bridged taps is limited, multiple bridged taps must be correspondingly shorter, placing their dips at higher frequencies. An example of two bridged taps of the same length (resulting in the maximum attenuation at the dip frequency) is shown in Fig. 6. The lowest possible dip frequency for this case is at about 50 kHz. The additional attenuation due to the bridged taps is greater than for the single bridged tap case, but the effect is limited to higher frequencies.

The characteristic impedance of the cable is approximately 100 Ω resistive at high frequencies. However, at the frequencies of interest in the U-transceiver, the impedance varies fairly rapidly with frequency, and is considerably larger than 100 Ω at low frequencies. The loss of a hybrid used to separate the two directions of transmission will depend on the accuracy with which a chosen balance impedance matches the input impedance of the cable. A resistive hybrid termination is typically



(a)



(b)

Fig. 7. Response of a hybrid with termination impedance 120 ohms to a transmitted 50% duty cycle one volt pulse (baud rate 160 kHz). Cable is 24 AWG 18 kft (5.5 km). (a) Without a bridged tap. (b) With a 22 AWG bridged tap with length 5 km at a distance 0.1 km from the hybrid.

used, but this can be improved upon with a more complicated termination impedance. The cable input impedance is also affected materially by the presence of bridged taps. With a fixed compromised balance impedance, the lowest transhybrid loss that is encountered is on the order of 10 dB. The nature of the echo response through the hybrid is important to the design of an echo canceler (EC). Shown in Fig. 7(a) and (b) are the echo pulse responses for a cable without and with a bridged tap near the end of the cable. While the details of the responses are different (the bridged tap makes the response much more complicated), in both cases they consist of a short active region confined to a few baud intervals together with a long tail corresponding to the collapse of the inductively and capacitively stored energy in the transformer and cable. The total active length of the response is on the order of 40 baud intervals. Some U-transceiver designs have taken advantage of the characteristics of this response to reduce the EC complexity as described in Section VII-B.

The range (distance from LT to NT) that can be achieved in a U-transceiver is generally limited by the high-frequency gain that must be inserted into the receiver equalization to compensate for cable attenuation. This gain amplifies noise, crosstalk, and interference signals that may be present, causing the error rate to deteriorate as the range increases. The most important noise and interference signals are impulse noise (caused for example

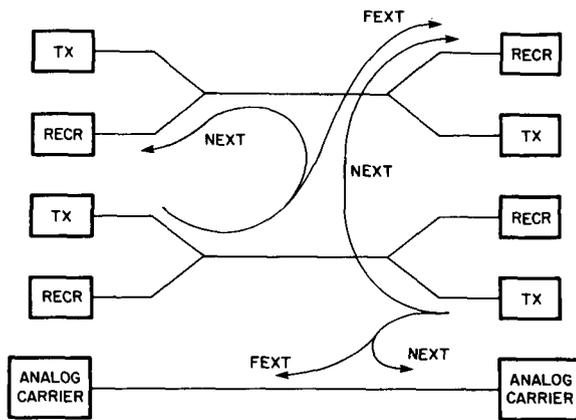


Fig. 8. Illustration of near-end crosstalk (NEXT) and far-end crosstalk (FEXT) for two digital subscriber loops and an analog carrier system sharing the same cable.

by dial-pulse signaling in the same central office) and crosstalk. There are two basic crosstalk mechanisms, near-end crosstalk (NEXT) and far-end crosstalk (FEXT), illustrated in Fig. 8. NEXT represents a crosstalk of a local transmitter into a local receiver, and experiences an attenuation that is accurately modeled by [1]

$$|H_{\text{NEXT}}(f)|^2 = K_{\text{NEXT}}|f|^{3/2}$$

where $H(f)$ is the transfer function experienced by the crosstalk. FEXT represents a crosstalk of a local transmitter into a remote receiver, with an attenuation given by

$$|H_{\text{FEXT}}(f)| = K_{\text{FEXT}}|C(f)|^2|f|^2$$

where $C(f)$ is the transfer function of the cable. Where present, NEXT will dominate FEXT because FEXT experiences the loss of the full length of the cable (in addition to the crosstalk coupling loss) and NEXT does not. Both forms of crosstalk experience less attenuation as frequency increases, and hence, it is advantageous to minimize the bandwidth required for transmission in a crosstalk-limited environment.

IV. METHODS OF REALIZING FULL DUPLEX TRANSMISSION

A significant challenge in the U-transceiver design is to derive a full duplex data stream from a single wire pair, for which there have been three methods proposed [3]: frequency division multiplexing (FDM), time compression multiplexing (TCM), and echo cancellation (EC). FDM has not been considered because of the stringent filtering requirements. EC generally has greater range than TCM [4], [5] due to its lower bandwidth, lower cable attenuation, and reduced NEXT crosstalk coupling. However, in Japan they are willing to synchronize the bursts of data in TCM, which eliminates NEXT, and have concluded that TCM has a greater range in that environment [6]. The EC method is gaining acceptance in most parts of the world as the method of choice, while Japan is pursuing the TCM method.

The EC method is shown in Fig. 9, and allows the data

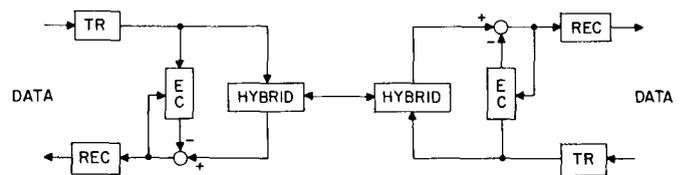


Fig. 9. Configuration for the EC method of separating the two directions of data transmission.

signals in the two directions to be transmitted at the same time within the same bandwidth. The impulse response of the undesired echo of the local transmit signal into the receiver (an example of this echo was shown in Fig. 7) is learned by an adaptive echo canceler (EC), which also generates an accurate replica of this echo and subtracts the replica to cancel the interference.

Many approaches combine two of the three methods. They generally increase the transmitted signal bandwidth and, therefore, reduce the ultimate range as dictated by cable attenuation, but exchange this for a simplification of some portion of the design. An example of such an approach is a combination of EC and FDM in which the frequency bands in the two directions are only partially overlapped. To the extent that the two directions are non-overlapping, the requirement on the EC accuracy is reduced and the effect of NEXT is reduced. Another very interesting combination, discovered independently in Italy [7], [8] and Japan [9], includes both TCM and EC. Following [7], [8], we will call this echo cancellation burst mode (ECBM). In ECBM, the baud rate is increased to about 20 percent above the bit rate (from 160 to 192 kbaud/s), thereby leaving dead-times during which there is no transmission. By increasing the interval between these dead-times, the baud rate can be reduced to be closer to 160 kbaud/s [9]. The dead times are adjusted so that there are brief periods during which only one direction or the other is transmitting. The timing recovery is simplified, since there is a period during which timing recovery can be performed without interference from the near-end echo signal and without the need for prior EC. The adaptation of the EC is also eased, since a period is available during which there is no far-end data signal to interfere with the adaptation. As the interval between dead-times is increased to reduce the baud rate, the opportunities for timing recovery and EC adaptation become less frequent, slowing down the initial acquisition. Therefore, there is a tradeoff between baud rate and acquisition performance.

In the sequel, we concentrate on the significant design choices that must be made in the design of a U-transceiver based on the EC method [3], [10]. A block diagram of a typical EC transceiver is shown in Fig. 10. The transmitted data is scrambled to ensure sufficient timing energy and also to ensure that the data signals in the two directions are uncorrelated (correlation would bias the EC adaptation). The line coding determines the manner in which the scrambled data bits are mapped into the transmitted data symbols. If the line coding is linear, the data bits at the input to the line coder can be fed directly to the EC,

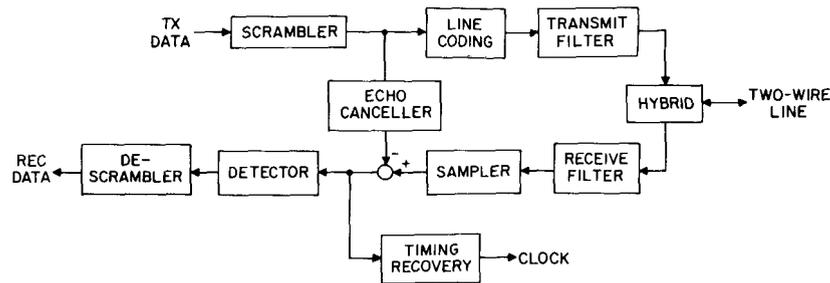


Fig. 10. Block diagram of an EC transceiver at one end.

and this binary input signal simplifies the implementation of the EC since no multiplications are required. The transmit filter minimizes high-frequency energy to reduce the radio frequency interference (RFI) and crosstalk. In the receive direction, the receive filter prevents aliasing in the subsequent sampling. The echo replica is subtracted at this point, and the resulting signal is free of interference from the near-end transmitter. The detector, which may include equalization to compensate for high-frequency attenuation and the effect of bridged taps, recovers the data symbols, and the timing recovery circuit recovers the proper phase of the receive sampling in relation to the received data symbols.

The most critical issues in the design of the transceiver are covered in the sequel. We generally discuss only the algorithmic issues, and not the difficult circuit design issues involved in the design of the line driver and data converters.

V. LINE CODE

The line code specifies the mapping between transmitted data stream and the actual pulses transmitted on the line. This is a critical choice in the system design because it affects both performance and implementation complexity. Among the performance ramifications are the following.

- 1) The probability of error for a given noise and interference level is affected by the spacing between received signal levels. The latter is influenced by the peak and average transmitted power, which are constrained by crosstalk considerations, and by the line code, which determines the number of received signal levels.
- 2) The line code can achieve, at a given bit rate, a lower baud rate at the expense of more transmitted and received levels. The baud rate relates directly to the speed of the hardware and also impacts the level of crosstalk interference and RFI.
- 3) The pulse shape as defined by the line code affects the power spectrum of the transmitted signal and, hence, the crosstalk into foreign systems, RFI, and required equalization. The line code can also introduce correlation into transmitted data symbols, thereby influencing the transmitted power spectrum, crosstalk, and RFI.
- 4) The details of the line code affect many aspects of the implementation, such as the complexity of the equalization, detection, EC, and timing recovery circuitry. For

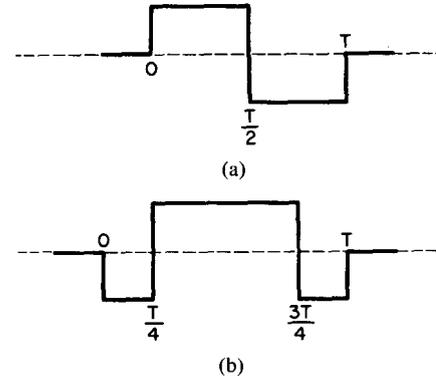


Fig. 11. Two transmitted pulses with no dc content. Normally the pulse or its negative is transmitted (binary antipodal signaling). (a) The biphase pulse. (b) The Walz pulse.

example, a line code with more low-frequency energy will generally require more EC taps.

- 5) The line code can be chosen to achieve specific objectives. For example, the cable is almost always transformer coupled, in which case the line code can prevent baseline wander if it ensures that there is no dc content to the transmitted data signal.

We will now briefly enumerate some of the popular choices for line codes in U-transceivers. The simplest line codes transmit one of two pulse shapes for a zero and one bit, where normally one pulse is the negative of the other (known as binary antipodal signaling). If the signal is to have no dc content for all possible transmitted data sequences, the basic pulse must have zero area. The transmitted pulses for two such binary antipodal line codes, the biphase and Walz, are shown in Fig. 11 [11]. These two line codes are often chosen for their “self-equalizing” properties, meaning that often a single compromise equalizer will suffice for a wide range of line lengths (in the absence of bridged taps). In addition, the effective echo impulse response is much shorter (on the order of 8–16 baud intervals) than that shown in Fig. 7, since the long tails due to the positive and negative pulses approximately cancel. The price that is paid for a transmitted pulse with extra zero crossings in the middle of the baud interval is a relatively large high-frequency content, which increases crosstalk and RFI.

The alternative mark inversion (AMI) code is a more conservative choice in terms of minimizing crosstalk, but requires adaptive equalization for even modest variations

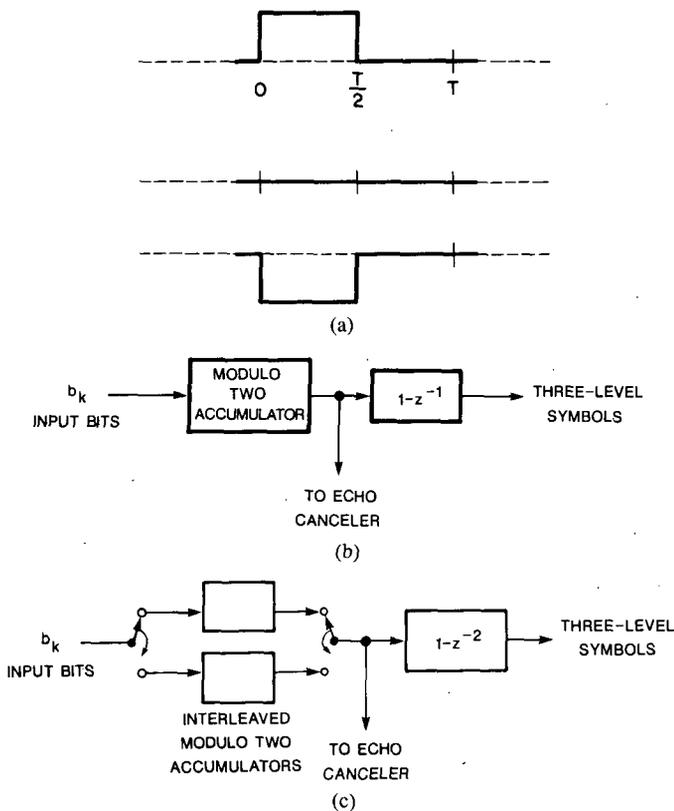


Fig. 12. The alternative mark inversion (AMI) line code. (a) The three possible transmitted pulses for one baud interval. Each transmitted nonzero pulse is the opposite polarity to the last transmitted nonzero pulse. (b) A realization of AMI as dicode partial response. (c) A modification to AMI as modified duobinary partial response.

in range. This code transmits a three-level signal, where the three transmitted pulses are shown in Fig. 12(a): a positive, negative, and zero pulse (the duty cycle shown is 50 percent, which can be varied). The pulses have a dc content, but the transmitted data signal can be constrained to have no average dc content by use of the redundancy available in mapping a single bit into three levels. In particular, zero-bits are transmitted as the zero pulse, and one-bits are transmitted as alternating positive and negative pulses. Since each nonzero pulse is balanced by a subsequent pulse of opposite polarity, a net absence of dc is assured.

The AMI code is equivalent to a form of partial response called "dicode partial response" and illustrated in Fig. 12(b). In this form, the AMI encoder is implemented as a modulo-two accumulator (toggle flip-flop) together with a first-difference operator ($1 - z^{-1}$). This realization would almost always be used in practice because the EC input can be connected to the binary (rather than three-level) signal at the accumulator output, and the EC can take care of the linear distortion introduced by the first-difference. More importantly, the first-difference reduces the number of required EC taps (as in biphase), due to the cancellation of the echo response tail.

Modified duobinary (class IV partial response) in Fig. 12(c) is a simple extension of dicode partial response in which the first difference is replaced by $(1 - z^{-2})$ and the

modulo-two accumulator is replaced by two such accumulators, each operating on half-rate bit streams that are then interleaved. This has the advantage of placing a zero in the spectrum at half the baud rate, thereby reducing RFI and crosstalk. With the exception of RFI, it is beneficial to place a portion of this response $(1 + z^{-1})$, in the receiver after EC. This does not affect crosstalk, and results in fewer EC taps and reduced sensitivity to additive noise.

The AMI and modified duobinary codes make inefficient use of the signal set, since only one bit out of $\log_2(3) = 1.58$ bits available in each baud interval is used. By using the three-level signal more efficiently, for a given bit rate the baud rate can be reduced. This is done in the 4B3T line code, in which four bits (16 alternatives) are mapped onto three ternary transmitted digits (27 alternatives) [12]. This mapping is performed in such a way that there is no dc although this requires a more complicated scheme of running digital sums. The resulting number of transmitted bits per baud is 1.33, closer to the ultimate of 1.58 for a three-level signal, and the baud rate is reduced relative to AMI by 25 percent (120 kbaud/s rather than 160). The reduction in baud rate in 4B3T is desirable for crosstalk and RFI. It also reduces the speed of the circuitry, but complicates it due to a finite state machine that must be implemented, as well as a large number of EC taps because of the lost beneficial effect of the first-difference operator. The baud rate can be further reduced by using a 3B2T code, or by transmitting four or more levels [13]. The benefit of this depends on the details of the frequency dependence of cable attenuation, the dominant source of noise and interference, and implementation details. For example, in a crosstalk-limited situation, reducing the baud rate with multilevel transmission reduces the crosstalk but also results in reduced spacing between levels. Recent work has indicated that in a NEXT-limited environment five levels is optimum, and four levels are almost as good [14]. It is also interesting to consider further reducing the susceptibility to noise and crosstalk using trellis coding [15], as has been done in voiceband data transmission.

VI. HYBRID DESIGN

The required accuracy of EC is directly related to the maximum cable attenuation to be encountered and the transhybrid attenuation in the echo path. For example, if the cable attenuation is 45 dB and the transhybrid loss is 10 dB (typical worst case numbers), then the signal-to-echo ratio at the hybrid output is minus 35 dB. The EC must then reduce the echo by an additional 55 dB to achieve a 20 dB signal-to-echo ratio, a typical number required to support nearly error-free data transmission [16].

In [9] it is reported that when the best of seven balancing networks is selected, the worst case transhybrid attenuation is improved by 18 dB relative to a single compromise termination. This improvement would reduce the required EC accuracy by 18 dB and reduce the required

data converter precision from approximately 12 bits to 9. The appropriate network is easy to select, since the choice is only critical on long cables where the hybrid output is dominated by the undesired echo. The balance termination can be chosen to minimize the size of this echo.

VII. NONLINEAR ECHO CANCELLATION

Effects like nonlinearity in data converters, transmitted pulse asymmetry, and timing jitter can compromise the EC performance objectives, unless special measures are taken in either the EC design or the design of the remaining circuitry.

There are therefore some fundamental tradeoffs that the transceiver designer faces.

1) Is the EC accuracy requirement to be minimized by use of a selective termination for the hybrid? We described a technique to achieve this in Section VI. This reduces the EC complexity, but requires additional analog circuitry in the hybrid.

2) Are the data converters and line driver designed with sufficient linearity to allow a linear EC, possibly at the expense of additional die area (Section VII-C)?

3) Are we going to implement an EC that can tolerate nonlinear distortion in the data converters, line driver, and other parts of the transceiver, again at the expense of die area?

4) Are we going to build an EC that is tolerant of timing jitter as described in Section IX?

In this section we discuss techniques for realizing an EC that is tolerant of nonlinear distortion in the transmitter or the data converter.

The EC in its simplest form is realized as a transversal filter. If the EC is realized digitally, the output can be converted to analog using a D/A converter prior to cancellation, or the hybrid output can be converted to digital by an A/D converter prior to cancellation in the digital domain [17]. In either case the required precision of the data converter (A/D or D/A) is around 12 bits for a 60 dB cancellation accuracy objective. Large word lengths are also required in the EC for the adaptation of the tap coefficients (greater than 20 bits), due to the long time constants of adaptation required to keep the asymptotic fluctuation small enough to meet the accuracy objectives.

Often it is desired to have a sampling rate for purposes of timing recovery that is at least twice the baud rate, since a baud rate sampling results in aliasing distortion. This requires that the EC have a sampling rate at its output that is a multiple of its input sampling rate. This can be accomplished with an interleaved EC, as shown in Fig. 13 for the case of an output sampling rate four times the input sampling rate [3]. Each of the four EC's has an input and output sampling rate equal to the baud rate, and each takes care of one phase of received sampling. The complexity of the EC is directly proportional to the sampling rate at its output. This leads to an interest in minimizing that sampling rate through special techniques for timing recovery as discussed in Section IX.

In the following subsections we discuss three types of

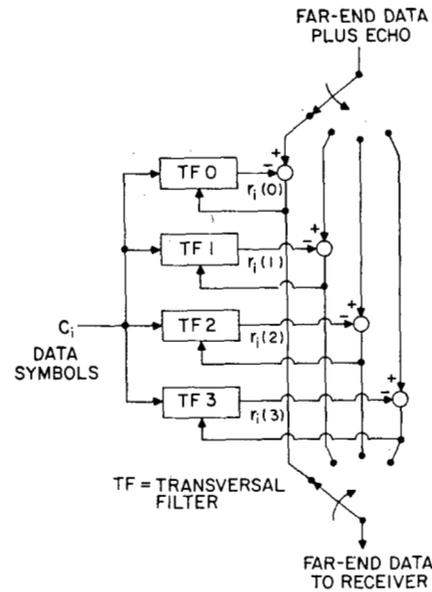


Fig. 13. Interleaved EC's to achieve a sampling rate at the output four times as great as at the input.

nonlinear EC's: the paired memory technique, the memory compensation EC, and the Volterra EC.

A. Paired Memory Canceler

The simplest form of nonlinear echo generation is actually due to nonlinearities in the transmitter that result in errors in the spacing of transmit levels. For example, for the binary case it is difficult to generate positive and negative pulses that are identical (within an accuracy of 60 dB) except for sign. This "transmitted pulse asymmetry" cannot be compensated by a linear EC. To see this, consider the simple case of binary antipodal signaling, for which the transmitted pulse is either $g(t)$ or $h(t)$. The case of pulse symmetry corresponds to $g(t) = -h(t)$. The transmitted pulse for the k th transmitted bit can be written in the form

$$\begin{aligned} p_k(t) &= \frac{1}{2}(1 + B_k) g(t) + \frac{1}{2}(1 - B_k) h(t) \\ &= b_k g(t) + \bar{b}_k h(t) \end{aligned} \quad (7.1)$$

where the k th transmitted bit is b_k , assuming the values 1 and 0, and the transmitted data symbol B_k assumes the values +1 and -1. The case of pulse symmetry results in a transmitted pulse, from (7.1).

$$p_k(t) = B_k g(t) \quad (7.2)$$

for which a linear EC will suffice. A two-coefficient linear EC for this case is shown in Fig. 14(a), in which the echo replica \hat{e}_k is a linear combination of B_k and B_{k-1} . When the transmitted pulses are asymmetric, then the representation of (7.1) can be used as shown in Fig. 14(b). There are now four tap coefficients rather than two, since positive and negative transmitted pulses are assigned different tap coefficients. The number of additions, however, is only one, the same as in Fig. 14(a), since the transmitted bit b_k actually selects which of the two coefficients is ap-

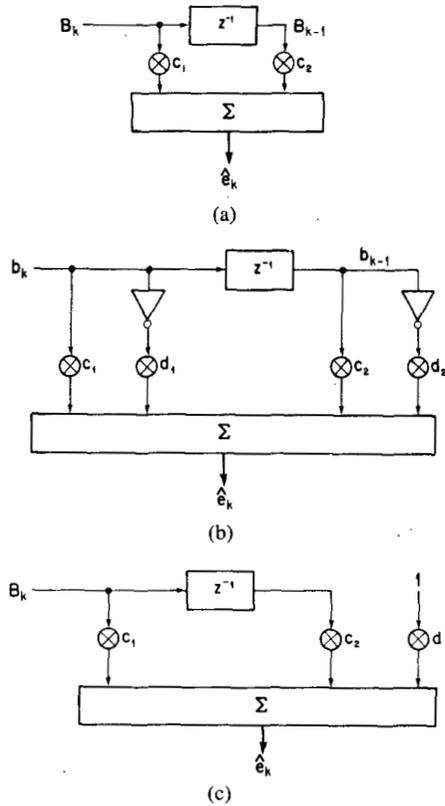


Fig. 14. EC's with baud rate sampling at the output and that cancel echo from the current and one past transmitted data symbol B_k corresponding to bit b_k . (a) A linear EC. (b) One form of nonlinear EC that compensates for transmitted pulse asymmetry by storing and adapting two sets of coefficients, one for positive and one for negative transmitted pulses. (c) A second form of nonlinear EC that compensates for transmitted pulse asymmetry by adding a dc tap coefficient.

appropriate at each delay. The structure of Fig. 14(b) can be adapted in the same manner as that of Fig. 14(a); namely, each coefficient is incremented by an amount proportional to the product of the input of the coefficient multiplier and the cancellation error. Since in Fig. 14(b) one tap coefficient multiplier at each delay has a zero input, in fact only one coefficient at each delay is updated by the adaptation algorithm—the one that is used in the generation of the echo replica. This structure will adapt automatically to the two polarities of transmitted pulses, whatever they may be.

An alternative to Fig. 14(b), shown in Fig. 14(c), follows from (7.1),

$$p_k(t) = \frac{1}{2}(g(t) + h(t)) + \frac{1}{2}B_k(g(t) - h(t)). \quad (7.3)$$

The first term in (7.4) corresponds to a dc offset pulse that is transmitted for every data symbol, and this offset can be compensated by adding a single tap coefficient to the filter. For a large number of delays this form of the filter is clearly superior, since it only requires memory for one more coefficient rather than doubling the coefficient memory.

B. Cancelers for General Nonlinearity

Other sources of nonlinearity are the data converters, line driver, and transformer. To implement EC's that

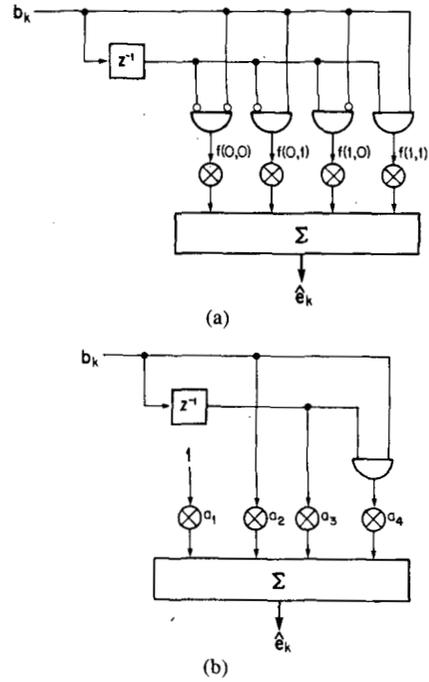


Fig. 15. EC's with baud rate sampling at the output and that cancel echo from the current and one past transmitted data symbol B_k corresponding to bit b_k . Both EC's will compensate for arbitrary time-invariant nonlinear echo mechanisms. (a) A memory compensation EC that stores a value of the echo response for each possible sequence of transmitted bits. (b) A Volterra series EC that adds nonlinear taps to a linear EC.

adaptively compensate for general nonlinearities, subject to the limitation that the nonlinearity have finite memory, assume that the echo path generation results in an echo of the form

$$e_k = f(b_k, b_{k-1}, \dots, b_{k-M}) \quad (7.4)$$

where $f(\cdot)$ is an arbitrary nonlinear function and again b_k is the k th transmitted bit. The *only* assumption that is made in (7.4) is that the current and past M transmitted data symbols contribute to the current echo sample in the same way at each sample time. The nature of the nonlinearity need not be memoryless (effects of, for example, hysteresis can be modeled).

Two forms of a nonlinear EC that generate an echo replica in the form of (7.4) have been proposed. Taking again the simple case of $M = 1$, (7.4) can be expanded in the form

$$e_k = \bar{b}_k \bar{b}_{k-1} f(0, 0) + \bar{b}_k b_{k-1} f(0, 1) + b_k \bar{b}_{k-1} f(1, 0) + b_k b_{k-1} f(1, 1) \quad (7.5)$$

resulting in an EC as shown in Fig. 15(a). This EC has been called a "memory compensation EC" [18], and stores the value of the echo replica for each possible transmitted data sequence (with finite memory, of course). While Fig. 15(a) emphasizes the similarity of this structure to a transversal filter, in actuality only one tap coefficient is used at each baud interval. Therefore, it is natural to store the echo replicas in an RAM, and the (b_k, b_{k-1}) can be used as an address for this RAM to select the proper coefficient. Only the coefficient that is actually

used is adapted in each baud interval. This structure generalizes to the general M case, requiring 2^{M+1} total coefficients.

An alternate expansion of (7.4), given by

$$e_k = a_1 + a_2 b_k + a_3 b_{k-1} + a_4 b_k b_{k-1} \quad (7.6)$$

corresponds to the structure of Fig. 15(b). This approach [19] approximates the nonlinear echo replica by a Volterra-like expansion, which can be viewed as adding extra nonlinear taps to a linear transversal filter EC, in this case a dc tap and a tap with input $b_k b_{k-1}$. The approach generalizes to general M and requires 2^{M+1} tap coefficients. While the complexity of the memory compensation EC is 2^{M+1} regardless of the severity of the nonlinearity that is being canceled, the Volterra approach can be simplified by eliminating negligibly small tap coefficients when the actual echo response is nearly linear.

The complexity of even the memory compensation EC can be reduced under some circumstances, such as for the echo response as shown in Fig. 7. The nature of the nonlinearities usually encountered is such that only signals that exercise a significant portion of the dynamic range evidence significant nonlinear effects. The long tail in Fig. 7, even when superimposed from many transmitted pulses, causes a relatively small excursion in total signal level. Therefore, it has been reported that the combination of a nonlinear EC (such as memory compensation) for short echo delays with a linear transversal filter for long delays yields adequate EC accuracy [7]. Another variant is to use two memory compensation EC's, one for short delays and one for long delays, and add their outputs [5]. In fact, this idea can be extended to a multiplicity of memories, with the addition of their outputs [20]. The simple approach to canceling transmitted pulse asymmetry is another illustration of a simplification of the nonlinear EC.

C. Avoiding Nonlinear Echo Cancellation

There are some ways in which the need for nonlinear EC can be reduced through careful design of other parts of the transceiver. For example, the impact of data converter nonlinearity can be greatly reduced if the D/A converter is used to convert the tap coefficients to analog, and analog switched capacitor circuitry is used to implement the transversal filter [17] as shown in Fig. 16. The effect of the nonlinearity in the conversion of digital tap weights to analog can be compensated by the adaptation, and additional nonlinear taps can be added to compensate for external nonlinearities. Recent work indicates it may be possible to implement the entire EC, including adaptation circuitry, in analog circuitry [21], thereby eliminating data converters altogether (along with their nonlinearity) and possibly resulting in a substantial reduction in die area.

It is feasible to design data converters, line drivers, etc., with sufficient linearity to allow a linear EC, although this generally requires greater die area and may or may not be the best overall tradeoff. For example, self-calibrating data converters [22], [23] can achieve adequate linearity,

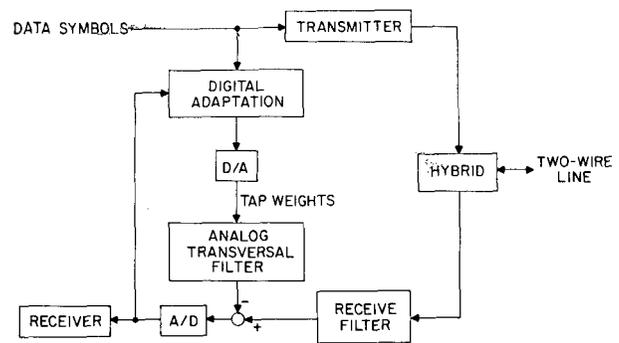


Fig. 16. Implementation of an EC insensitive to D/A converter nonlinearities.

as can a very high sampling rate followed by decimation. With a sufficiently high sampling rate a single bit converter can be used, with minimal accuracy requirements, followed by digital signal processing that can achieve sufficient accuracy simply by the choice of word lengths. For example, in [24], [12] a 15.36 MHz sampling rate is used with a delta-sigma encoder using a switched capacitor integrator.

VIII. EQUALIZATION

The design of the equalization must take into account two factors: the unknown cable length and the presence in some administrations of bridged taps. The basic design choices that must be made in the design of the equalizer are as follows.

- 1) Is the equalizer to be placed before or after the EC?
- 2) Is the equalizer realized digitally, or is some part realized using switched capacitor filters?
- 3) What degree of adaptation is required in the equalizer?
- 4) How does the equalizer take account of bridged taps?

In the absence of bridged taps, there is only one parameter of the cable that must be adapted to—its length. In actuality, the cable may consist of more than one gauge, and in any event the gauge is unknown to the transceiver. However, if the loss of the cable is known at one frequency, then the entire loss versus frequency curve can be predicted with sufficient accuracy to perform equalization. Therefore, the adaptive equalizer can be parameterized by this single parameter, the cable loss at a particular frequency, and this parameter can be automatically adjusted. Such an equalizer is often called a \sqrt{f} equalizer, since the loss of the cable in decibels at high frequencies is roughly proportional to the square root of frequency independent of the gauge or gauge mix.

A common method to indirectly measure the loss of the cable is to monitor the peak received data signal at the equalizer output, and adjust that peak to a predetermined level. For AMI coding, for example, this peak signal is directly related to the cable loss at half the baud rate, and this strategy is therefore equivalent to setting the equalizer gain to be appropriate at the half baud rate. In the older T-carrier systems, the single parameter equalizer was called an ALBO (automatic line build-out), and was

usually implemented as a single-zero active filter with variable zero frequency. More suitable for MOS implementation, however, is a switched capacitor realization of the \sqrt{f} equalizer [7]–[9], [25]–[27], adjusting the equalizer gain in discrete steps (by switching in appropriate capacitors), where the steps are typically about 10 dB apart at the half baud rate.

The switched capacitor equalizer requires a higher sampling rate (four times the baud rate or higher) than is generally practical to realize in the EC, and is therefore usually placed in front of the EC. Any gain in the equalizer at high frequencies will also simplify the echo signal at those frequencies, although the effect is minimized by the fact that the hybrid balance is usually the most accurate at high frequencies. On the other hand, the equalizer in the echo path before the EC can be beneficial if it includes a high-pass filter characteristics, since this will minimize the effect of poor hybrid balance at low frequencies and, in particular, will reduce the required number of EC tap coefficients.

In one interesting design [9], the EC has been split into two parts. The first part operates at a sampling rate eight times the baud rate, and is followed by a switched capacitor equalizer with reduced dynamic range requirements. The second part, which follows the equalizer, operates at the baud rate and removes the remaining echo. It is particularly important that the second part be a nonlinear EC, since it is compensating for the nonlinear residual granularity of the first EC.

For a more digital implementation, a conventional adaptive transversal filter equalizer can be used [28]. This transversal filter can operate at a sampling rate equal to the baud rate, in which case the equalizer noise enhancement is somewhat sensitive to the sampling phase. An alternative is the "fractionally spaced equalizer," in which the sampling rate is twice the baud rate, reducing the sensitivity to sampling phase [29]. The transversal filter equalizer can be placed either before or after the EC, since the sampling rate requirements are modest.

When bridged taps are present, they add postcursor intersymbol interference, which is easily compensated adaptively using a decision-feedback equalizer (DFE) [7], [9], [25]–[27] that uses past decisions to generate a replica of the interference from past symbols. Like the EC, the DFE does not require multiplications since the transversal filter input consists of data symbols; in fact, the DFE is naturally implemented using the same processing engine as the EC.

In principle, a linear transversal filter equalizer could compensate for both cable attenuation and bridged taps. However, the simple implementation of the DFE makes it more attractive to use a DFE for bridged tap equalization even when the cable attenuation is compensated by a linear transversal filter equalizer as in [12], [24]. In current technology it is very desirable to reduce the number of multiplications required for the equalization function. Where switched capacitor filters are also avoided, it is common to do the bulk of the equalization using the DFE.

In this case the DFE compensates for \sqrt{f} dispersion as well as bridged taps. Since the cable output generally has a slow risetime and, hence, has some precursor intersymbol interference (which also depends on timing recovery), it is generally not possible to rely totally on the DFE. However, it is possible to limit the number of forward linear equalizer taps to just a few.

IX. TIMING RECOVERY

Timing recovery is a critical function in the transceiver design [30] because of the need to meet three objectives.

1) Timing recovery should use a discrete-time representation of the far-end data signal. If the timing recovery is performed after EC, then the available representation of the far-end data signal is discrete-time since the EC is a discrete-time device.

2) The number of interleaved EC's (which relates directly to the implementation complexity) is proportional to the sampling rate; therefore, if timing recovery is performed after EC, the sampling rate should be minimized. (Techniques have been proposed for doing timing recovery before EC, as described in Section IX-B. In this case the sampling rate for the timing recovery technique is not critical.)

3) The residual jitter after timing recovery can affect the accuracy of EC; therefore, either this jitter should be minimized or the EC should be designed to accommodate this jitter without compromising accuracy. The nature of this problem will be explained in Section IX-A.

A. Loop Timing

The presence of loop timing in the NT transceiver and the interaction of timing recovery and EC lead to several subtle problems. Therefore, we will look first at the overall system configuration in Fig. 17, which illustrates the source of timing at the LT and NT. At the LT the switch master clock governs the transmit data stream, and in the receiver in the NT this clock is derived by the timing recovery circuit and used to govern the NT transmit data stream. This "loop timing" ensures that the stream returning to the switch is synchronous with the internal clock of the switch, although the incoming phase of the data stream will be indeterminate (reflecting the round trip propagation delay).

In summary, in the NT the frequency and phase of the incoming data stream must be tracked, whereas at the LT the frequency is accurately known but the phase must be estimated. At the NT there is no jitter on the incoming data stream to be tracked (in contrast to voiceband data modems). In the LT there may be a small data pattern-dependent jitter introduced by the fact that the NT is loop timed, but this jitter is so small that there is no advantage to tracking it. Therefore, in the LT the only phase tracking capability required is to compensate for variations in phase due to temperature changes.

In contrast to T-carrier digital transmission systems, the accumulation of jitter in long chains of repeaters is of no

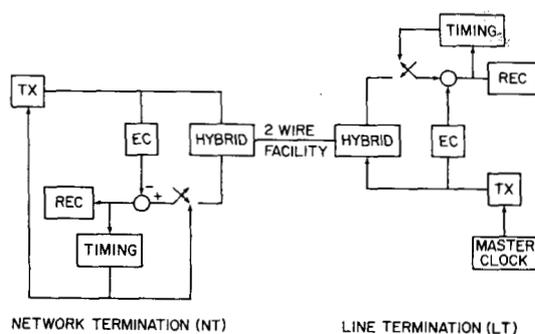


Fig. 17. Illustration of the source of timing in the LT and NT transceivers for a loop timed NT.

concern in a U-transceiver. Save for one critical factor, then, the inherent timing jitter in the U-transceiver would be of little concern. That one factor is the effect of timing jitter on the accuracy of EC in the NT due to the loop timing. Higher frequency jitter components in the derived timing will cause a relative jitter between past transmitted pulses and the current received sampling phase. This relative jitter implies that the effective impulse response of the echo response is changing faster than the EC can track.

At the LT, since the frequency of the incoming data stream is known to be synchronous with the master clock, the receive sampling can be derived from an appropriate phase of the master clock, again using a countdown circuit and a digital PLL. Since there is no appreciable phase jitter that needs to be tracked (other than the initial phase and perhaps long-term temperature changes), a constant phase of receive sampling relative to the master clock can be used. This ensures that there is no relative phase jitter between transmit and receive clock to compromise EC accuracy. However, if there is no high-speed master clock at the LT, a PLL that derives such a clock can introduce relative jitter.

B. PLL Timing Recovery

For monolithic implementation of the NT, it is generally advantageous to derive timing using a phase-locked loop (PLL). The basic elements of the PLL are a variable frequency oscillator (VCO) to generate the sampling clock and a phase detector to estimate the phase difference between that sampling clock and the optimum sampling phase for the received data stream.

There are a couple of phase detectors that have been suggested for deriving the received timing phase in discrete time. The wave difference method [31] is the simplest, but requires a sampling rate equal to at least twice the baud rate, which doubles the complexity of the EC. The baud rate timing recovery method [32], [21] reduces the EC complexity, but care must be exercised to avoid dependence of the recovered timing phase on the configuration of the cable. A particular concern is the effect of pulse distortion due to bridged taps, which can be compensated by using an appropriately designed line code [32], [21].

Another attractive method for baud rate timing recovery is "mean-square timing recovery," in which the timing phase is chosen directly to minimize the mean-square error after EC and equalization [33], [34]. The mean-square error can be estimated in a decision-directed fashion by time averaging the square of the difference between input and output of the slicer. This quantity is required at two different closely spaced phases, which can be achieved without increasing the sampling rate by using lead-lag received sampling [33].

Modifications to the frame format can simplify the timing recovery problem. One technique is to insert in the frame a deterministic synchronization burst, which can be used exclusively for timing synchronization. For example, in [24], [12] an 11-bit Barker code is inserted into the frame. This still requires that EC be done prior to timing recovery, but the synchronization burst may reduce the susceptibility to bridged taps with baud rate sampling. Another example is the ECBM technique [7]–[9], in which a portion of the frame is available where there is a received data signal but no echo from the near-end transmitter. If timing recovery is restricted to this interval, timing recovery can be performed before EC, and therefore, the sampling rate for timing recovery can be chosen for circuit simplicity without impacting the EC complexity.

There are several approaches to limiting the effect of timing jitter on EC accuracy. One solution is to design a timing recovery circuit that has a very narrow bandwidth to jitter, since low-frequency jitter will not result in a relative phase shift between transmit and receive timing phases within the memory of the EC [35], [36]. This requires a relatively difficult analog phase-locked loop design, since a digital PLL will introduce discrete phase jumps with an unacceptable reduction in cancellation accuracy. Most designs, however, use a digital PLL, and somehow eliminate the effect of phase jumps on the EC. Note that the impact of these phase jumps on EC accuracy is transient, since only the relative transmit and receive phases within the memory of the EC are relevant. One approach is, therefore, to force all slips to occur at the beginning of a time period where EC accuracy is not critical. For example, in ECBM [7], [8] the timing slips can be forced to occur at the beginning of the period in each frame where there is no received data signal. Alternatively, if a synchronization burst is included in the frame [24], [12], then the slip can be forced to occur at the beginning of that burst.

Insensitivity to phase jumps can also be achieved without any compromise in the design of the line code or frame format, although at the expense of increased EC complexity. One approach is to store tap sets of tap coefficients in the EC, one for each of two relative transmit-receive phases [37], with lead-lag sampling [33] used to adapt both phases. Another method which increases the computational overhead is to estimate the coefficients at the new sampling phase using a gradient estimation technique [38], [39].

X. PERSPECTIVE

In this section we express opinions as to the best design approaches available for the U-transceiver. There are two separable but important issues in the design of the U-transceiver. The first is the choice of a frame format and line code, and the second is the choice of design approaches within the transceiver receiver to achieve the performance objectives, given that frame format and line code. The worldwide standardization of the frame format and line code is desirable, since manufacturers can focus their efforts on a high-volume worldwide market. Individual manufacturers can choose receiver design approaches as they please, and are distinguished in the marketplace by performance and price. Premature standardization can stifle innovation, but it appears that almost complete penetration of the subscriber loop plant can be achieved within the context of line codes that have been proposed, and therefore, standardization is appropriate. Innovation in design techniques can proceed after standardization, and new modulation formats and line codes can still be explored as candidates for future higher rate ISDN interfaces (much as the voiceband data modems have evolved to ever higher rates over the years).

A very divisive issue in the present standards activities is the choice between a relatively low cost and simple line code and a more sophisticated and better performing design. The advocates of low cost contend that the ISDN will never penetrate the marketplace unless the cost is kept low, while others argue that high performance will reduce the installation costs enough to offset the higher equipment cost. Furthermore, the introduction of large numbers of U-transceivers into the subscriber loop plant has sobering implications in terms of crosstalk and radio-frequency interference, and line codes that minimize these problems also tend to have the greatest range capability and unsusceptibility to bridged taps, since they limit the transmitted spectrum to the low frequencies where cable attenuation is the lowest. The standardization activities are made even more difficult by disagreement over the quantitative relationship between cost and performance, and how this relationship will evolve with future technologies.

The best performance is achieved if high frequencies are limited by transmitting continuously (which rules out techniques like TCM and ECBM) and minimizing the baud rate by using multiple levels efficiently. While ECBM gives us a number of design simplifications for only a modest increase in bandwidth, it is also becoming clear that excellent performance can be achieved without these simplifications, albeit with an increase in circuit complexity. Some variation on the partial response or block codes provide a desirable combination of complexity and performance. At the other end of the range of possibilities, the use of biphase with TCM would be an unfortunate choice.

On the issue of which algorithmic approaches are best, much depends on the technological assumptions. Those who favor mostly-digital implementations, due to present

or evolutionary technological assumptions, will likely gravitate towards much different algorithmic approaches than those willing to use switched-capacitor and analog circuitry more liberally. Therefore, we will express opinions as to the best algorithmic approaches for both camps. In the case of the analog technology, we would choose the following implementation approaches:

1) Switched capacitor transmit filter to limit high-frequency power and switched capacitor band-limiting filter in the receiver.

2) Automatic selection of hybrid termination to limit the required EC accuracy and dynamic range of the receive filters.

3) Switched capacitor \sqrt{f} equalization with associated AGC to compensate for cable attenuation. This filter could also include a highpass characteristic to limit the required memory in the EC.

4) A digitally implemented EC with digital adaptation using the RAM-based memory compensation algorithm for the first few baud intervals and a transversal filter for the remainder.

5) Cancellation of echo in the analog domain preceded by a D/A convert at a baud rate. Due to the hybrid termination selection the D/A precision would be about 9 or 10 bits.

6) A DPLL for timing recovery in which a crystal oscillator at about 16 times the baud rate is counted down to yield the transmit and receive sampling phases.

7) Lead-lag sampling at the EC output with double EC memory for the two relative transmit-receive phases. This eliminates the effect of DPLL phase jumps on the echo cancellation accuracy.

8) Baud rate timing recovery, with a mean-square error phase detector using the lead-lag sampling to detect the appropriate direction of DPLL phase jump.

9) A DFE equalizer to compensate for the high-pass filtering and bridged taps.

This design would be a conservative one, allowing a frame format including no dead time or added sequences for timing recovery. If it was desired to simplify the implementation, then the following modifications would be desirable.

1) Use the ECBM technique, introducing a dead time in the frame at the expense of higher baud rate.

2) Adapt the EC only during the period when there is no far-end data signal, reducing the word length requirement for the adaptation.

3) Recover timing during the period when there is no near-end transmitted data signal, using the wave difference method with a sampling rate four times the baud rate.

4) Consider using one of the simpler line codes, such as biphase.

Looking into the future, implementation complexity will be less an issue, and it will be desirable to minimize the amount of analog circuitry. Many designers would favor this approach even today. Under this scenario we could consider making the following modifications to the design.

1) Use data converter techniques that minimize analog circuitry such as delta-sigma modulation. This may also eliminate the need for nonlinear EC.

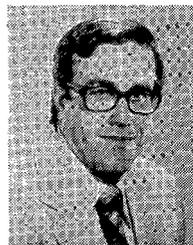
2) Abandon the automatic selection of hybrid termination in favor of increased EC accuracy.

3) Replace the \sqrt{f} switched capacitor equalizer with a baud rate transversal filter equalizer, retaining also the decision feedback equalizer.

With any of these design approaches it appears possible, with careful design, to accommodate the vast majority of subscriber loops the world over and at the same time minimize crosstalk and RFI problems.

REFERENCES

- [1] S. V. Ahamed, P. P. Bohn, and N. L. Gottfried, "A tutorial on two-wire digital transmission in the loop plant," *IEEE Trans. Commun.*, vol. COM-29 Nov. 1981.
- [2] D. Messerschmitt, "A transmission line modeling program written in C," *IEEE J. Select. Areas Commun.*, vol. SAC-2, p. 148, Jan. 1984.
- [3] —, "Echo cancellation for voice and data transmission," *IEEE Trans. Commun.*, vol. COM-32, Feb. 1984.
- [4] B. Aschrafi, P. Meschkat, and K. Szechenyi, "Field trial of a comparison of time separation, echo cancellation, and four-wire digital subscriber loops," in *Proc. Int. Symp. Subscriber Loops, Services*, Sept. 1982.
- [5] J.-O. Andersson, B. Carlquist, A. Bauer, and I. Dahlqvist, "An LSI implementation of an ISDN echo canceler: Design and network aspects," this issue.
- [6] K. Yoshida and N. Tamaki, "Subscriber loop noise considerations and the estimated TCM application range," in *Proc. IEEE Int. Conf. Commun.*, June 1985.
- [7] A. Brosio, C. Mogavero, and A. Tofanelli, "Echo canceler burst mode: A new technique for digital transmission on subscriber lines," in *Proc. IEEE Int. Conf. Commun.*, Chicago, IL, June 1985.
- [8] X. Moncalvo, G. Valbonesi, and X. Rizzotto, "ECBM: A low-cost echo canceler in a silicon boutique for terminals, PBXs, and CO access," this issue, pp. 1331-1336.
- [9] M. Fukuda, T. Tsuda, and K. Murano, "Digital subscriber loop transmission using echo canceler and balancing networks," in *Proc. IEEE Int. Conf. Commun.*, June 1985.
- [10] A. Kanemasa, A. Sugiyama, S. Koike, and J. Uoyama, "An ISDN subscriber loop transmission system based on echo cancellation," this issue, pp. 1359-1366.
- [11] A. Brosio, U. DeJulio, V. Lazzari, R. Ravaglia, and A. Tofanelli, "A comparison of digital subscriber line transmission systems employing different line codes," *IEEE Trans. Commun.*, COM-29, p. 1581, Nov. 1981.
- [12] H. Sailer, H. Schenk, and E. Schmid, "A VLSI transceiver for the ISDN customer access," in *Proc. IEEE Int. Conf. Commun.*, Chicago, IL, June 1985.
- [13] P. Adams, S. Cox, and P. Glen, "Long reach duplex transmission systems for ISDN access," *Brit. Telecom Technol. J.*, vol. 2, p. 35, Apr. 1984.
- [14] G. Brand, V. Madisetti, and D. G. Messerschmitt, "Multilevel signalling in the presence of near-end crosstalk in digital subscriber loops," *IEEE Trans. Commun.*, submitted for publication.
- [15] G. D. Forney, Jr., R. G. Gallager, G. R. Lang, F. M. Longstaff, and S. U. Qureshi, "Efficient modulation for band-limited channels," *IEEE J. Select. Areas Commun.*, vol. SAC-2, Sept. 1984.
- [16] J. W. Modestino, C. Massey, R. E. Bollen, and R. P. Prabhu, "Modeling and analysis of error probability performance for digital transmission over two-wire loop plant," this issue, pp. 1317-1330.
- [17] O. Agazzi, D. A. Hodges, and D. G. Messerschmitt, "Large-scale integration of hybrid-method digital subscriber loops," *IEEE Trans. Commun.*, vol. COM-30, p. 2059, Sept. 1982.
- [18] N. Holte and S. Stuefflotten, "A new digital echo canceler for two-wire subscriber lines," *IEEE Trans. Commun.*, vol. COM-29, pp. 1573-1581, Nov. 1981.
- [19] O. Agazzi, D. G. Messerschmitt, and D. A. Hodges, "Nonlinear echo cancellation of data signals," *IEEE Trans. Commun.*, vol. COM-30, p. 2421 Nov. 1982.
- [20] X. Maitre, M. Levy, and S. Surie, "The multi-memory echo canceler for two wire digital transmission," in *Proc. IEEE GLOBECOM*, Dec. 1984.
- [21] J. Tzeng, D. Hodges, and D. G. Messerschmitt, "Timing recovery in digital subscriber loops using baudrate sampling," this issue, pp. 1302-1311.
- [22] H. S. Lee, D. A. Hodges, and P. R. Gray, "A self-calibrating 15-bit CMOS A/D converter," *IEEE J. Solid-State Circuits*, vol. SC-19, p. 813, Dec. 1984.
- [23] P. Y. Li, J. M. Chin, P. R. Gray, and R. Castello, "A ratio-independent algorithmic A/D converter," *IEEE J. Solid-State Circuits*, vol. SC-19, p. 828, Dec. 1984.
- [24] L. Lerach, H. Sailer, and G. Schollmeier, "VLSI components for the ISDN customer access," in *Proc. IEEE GLOBECOM*, Nov. 1984.
- [25] H. Takatori and T. Suzuki, "Low-power line equalizer for digital subscriber loop," in *Proc. IEEE GLOBECOM*, Nov. 1984.
- [26] T. Suzuki, H. Takatori, M. Ogawa, and K. Tomooka, "Line equalizer for a digital subscriber loop employing switched capacitor technology," *IEEE Trans. Commun.*, vol. COM-30, p. 2074, Sept. 1982.
- [27] T. Chujo, N. Ueno, A. Takada, Y. Hino, and M. Fukuda, "A line termination circuit for burst-mode digital subscriber loop transmission," in *Proc. IEEE GLOBECOM*, Nov. 1984.
- [28] K. Szechenyi, F. Zapf, and D. Sallaerts, "Integrated, full-digital U-interface circuit for ISDN subscriber loops," this issue, pp. 1337-1349.
- [29] G. Ungerboeck, "Fractional tap-spacing equalizer and consequences for clock recovery in data modems," *IEEE Trans. Commun.*, vol. COM-24, p. 856, Aug. 1976.
- [30] K. H. Mueller and M. Muller, "Timing recovery in digital synchronous data receivers," *IEEE Trans. Commun.*, vol. COM-24, p. 516, May 1976.
- [31] O. Agazzi, C.-P. J. Tzeng, D. G. Messerschmitt, and D. A. Hodges, "Timing recovery in digital subscriber loops," *IEEE Trans. Commun.*, vol. COM-33, pp. 558-569, June 1985.
- [32] J. Tzeng, D. Hodges, and D. G. Messerschmitt, "Baud rate timing recovery in digital subscriber loops," in *Proc. IEEE Int. Conf. Commun.*, Chicago, IL, June 1985.
- [33] H. Sari, L. Desperben, and S. Moridi, "Optimum timing recovery for digital equalizers," in *Proc. IEEE GLOBECOM*, Dec. 1985.
- [34] R. Gitlin and J. Salz, "Timing recovery in PAM systems," *Bell Syst. Tech. J.*, vol. 50, p. 1645, May-June 1971.
- [35] D. D. Falconer, "Timing jitter effects on digital subscriber loop echo cancellers—Parts I and II," *IEEE Trans. Commun.*, vol. COM-33, pp. 826-838, Aug. 1985.
- [36] —, "Investigation of synchronization parameters in a digital subscriber loop transmission system," this issue, pp. 1312-1316.
- [37] D. G. Messerschmitt, "Asynchronous and timing-jitter insensitive data echo cancellation," *IEEE Trans. Commun.*, to be published.
- [38] S. Cox, "Clock sensitivity reduction in echo cancelers," *Electron. Lett.*, vol. 21, July 4, 1985.
- [39] R. Carpenter, S. Cox, and P. Adams, "Jitter compensation in echo cancelers," in *Proc. IASTED Int. Symp. Appl. Signal Processing, Digital Filtering*, June 1985.



David G. Messerschmitt (S'65-M'68-SM'78-F'83) received the B.S. degree from the University of Colorado, Boulder, in 1967, and the M.S. and Ph.D. degrees from the University of Michigan, Ann Arbor, in 1968 and 1971, respectively.

He is a Professor of Electrical Engineering and Computer Sciences at the University of California at Berkeley. From 1968 to 1977 he was a Member of the Technical Staff and later Supervisor at Bell Laboratories, Holmdel, NJ, where he did systems engineering, development, and research on digital transmission and digital signal processing (particularly relating to speech processing). His current research interests include analog and digital signal processing, adaptive filtering, digital communications (on the subscriber loop and fiber optics), architecture and software approaches to programmable digital signal processing, communication network design and protocols, and computer-aided design of communications and signal processing systems. He has published over 70 papers and has 10 patents issued or pending in these fields. Since 1977 he has also served as a consultant to a number of companies. He has organized and participated in a number of short courses and seminars devoted to continuing engineering education.

Dr. Messerschmitt is a member of Eta Kappa Nu, Tau Beta Pi, and Sigma Xi, and has several best paper awards. He is currently a Senior Editor of *IEEE Communications Magazine*, and is past Editor for *Transmission Systems of the IEEE TRANSACTIONS ON COMMUNICATIONS* and past member of the Board of Governors of the *IEEE Communications Society*.