Asynchronous and Timing Jitter Insensitive Data Echo Cancellation

DAVID G. MESSERSCHMITT, FELLOW, IEEE

Abstract—An approach to the implementation of asynchronous and timing jitter insensitive data echo cancellation is described. This approach introduces a small amount of jitter in the transmitted data signal, or alternatively in the received signal sampling, and uses a simple digital phase-locked loop together with the storage of two sets of echo canceler coefficients. The effect of derived timing jitter on the echo cancellation accuracy is completely eliminated for a loop timed transceiver (as in a digital subscriber loop network termination transceiver), and is easily reduced to negligible levels for a nonloop timed transceiver (as in a digital subscriber loop line card transceiver or a voiceband data modem). In the case of a voiceband data modem, this approach is one method to achieve asynchronous echo cancellation without the need to recover and resample a continuous-time far-end data signal.

I. INTRODUCTION

IN this paper we address the problem of designing an echo canceler (EC) for a full-duplex data transmission transceiver. Such an EC must achieve a high degree of accuracy, typically as high as 70 dB echo reduction in the worst case, and as a result, effects such as timing jitter and nonlinearities, which are not normally critical in data transmission systems, become very important. In this paper we specifically address the problem of timing jitter and its effect on the EC.

A full duplex data transmission system [1] is shown in Fig. 1. Specifically shown is the case of a digital subscriber loop, where the two directions of data transmission must be synchronized. The transceiver which is in the central office is called the line termination (LT) transceiver, and the transceiver on the subscriber premises is the network termination (NT) transceiver. The LT transceiver transmit clock is a master clock provided by the switch, or in some cases it may be a clock derived from a stable switch clock using a PLL. The timing of the data stream received from the switch is derived in the NT transceiver, and that timing is used to sample that same data stream at the output of the hybrid. Since the data stream sent back toward the LT must be synchronous with the switch clock, the NT transmitter uses the same receive timing (so-called "loop timing"). Back at the LT, this data stream arrives with an arbitrary phase relative to the transmit clock, although the average frequency is guaranteed to be the same, and hence, the timing must also be derived here and used to sample the data stream at the output of the hybrid.

Timing jitter causes a reduction in the degree of EC accuracy in the EC at both ends. At the LT, any changing phase in the receive timing will cause the effective impulse response of the echo path to change slightly with time, and unless this change is very slow the adaptation of the EC may

Fig. 1. Timing in an echo canceler (EC) digital subscriber loop.

not be able to follow it. In addition, any jitter caused by PLL derivation of a transmit clock from the switch may result in a reduction in EC accuracy. At the NT, any receive timing jitter causes the transmit pulses to be jittered in identical fashion. Since past transmitted pulses had phase jitter, the effective echo impulse response also changes, since the echo signal is sampled with a delayed version of the same timing signal. In this case, very slow jitter does not have to be tracked by the EC, since it affects the transmitted pulse and received sampling in the same way. However, higher frequency jitter which causes significant deviation of the transmitted pulse phase during a time period which is on the order of the maximum echo delay to be canceled will cause a reduced EC accuracy.

There are several solutions to the problem of EC accuracy. One solution is to keep the timing jitter extremely small [2]. This is undesirable since it requires an analog PLL, which is difficult to design and realize. A digital PLL (DPLL) is easier to realize, but results in discrete jumps in phase which cause a temporary unacceptable reduction in the EC accuracy. Alternatives which use a DPLL must somehow compensate for the effect of phase jumps. One approach is to use interpolation techniques in conjunction with choice of a particular line code [3]. Another approach is to design the frame format so that there are intervals of time where EC accuracy is less critical, and force the phase slips to occur at the beginning of those times [4]–[7].

The purpose of this paper is to show that the effect of timing jitter on the EC accuracy can be *completely eliminated* in the context of the simpler DPLL without any compromise in the choice of frame format or line code. The price that we pay is the storage of two sets of EC coefficients rather than one, and an increase in the complexity of the EC control mechanism. There is no increase in the EC computational requirements, nor in the rate of conversion of the data converters. Our technique can be used in both the LT and NT, and is compatible with both baud rate sampled EC's [8] and interleaved EC's [9]. Researchers at British Telecom have independently studied a technique similar to that proposed here [10]–[12]. Rather than perform a direct adaptive interpolation of the echo response, as we do here, they add an additional adaptation loop to estimate the derivative of that

Paper approved by the Editor for Digital Communications of the IEEE Communications Society. Manuscript received September 19, 1985; revised June 4, 1986. This work was supported by grants from Advanced Micro Devices, Fairchild Semiconductor, Harris Corporation, National Semiconductor, Racal-Vadic, and Bell Communications Research, with a matching grant from the University of California MICRO program.

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

IEEE Log Number 8611276.

response and then use that derivative to perform approximate interpolation.

Our proposed technique also results in a slight but controlled increase in the jitter in the received data stream. However, since this received data stream is much less sensitive to jitter than the EC (the signal-to-noise ratio which is acceptable is on the order of 20 dB rather than 70 dB), this is a very desirable tradeoff. Even this increased data jitter can be eliminated at the expense of doubling the EC computational requirements and D/A conversion rate, but this is unnecessary.

In brief, the technique proposed here should make it considerably easier to achieve the 70 dB EC accuracy objective, should greatly simplify the design of the timing recovery PLL, and on net should also result in an appreciable reduction in design complexity due to the elimination of the need for an analog PLL.

In Section II we describe the proposed technique. We also show how the technique is applicable to asynchronous EC, as in a voiceband data modem, enabling a simple and accurate interpolation to new phases as the relative transmit and receive phases slip with time.

II. ELIMINATING THE EFFECT OF TIMING JITTER

The proposed technique will be described in this section, first for the case where the two directions of full duplex transmission are synchronous, and then in Section II-F for the asynchronous case.

A. Digital PLL

The basic principle of a DPLL is shown in Fig. 2 for the case of the NT transceiver. An independent crystal oscillator provides a stable clock with nominal frequency Nf_b , where f_b is the baud (symbol) rate. We assume that baud rate sampling is used [8], although the technique is easily extended to higher sampling frequencies. For baud rate sampling, the oscillator frequency is divided by N to obtain the basic frequency used for sampling the hybrid output, and also to determine the times to transmit the data symbols. Since the oscillator will have some small frequency offset from the similar clock provided at the switch, a programmable divider allows a count which is occasionally $N - 1$ or $N + 1$, which provides a slip in phase by a time interval $1/Nf_b$. By controlling the interval between these slips, the frequency can be adjusted over the range

$$\frac{N}{N+1} f_b \leq f \leq \frac{N}{N-1} f_b.$$

A timing circuit estimates the direction in which the phase of the receive sampling clock should slip so as to maintain synchronism with the received data stream. When a slip in one direction or the other is indicated, the control circuit causes both counters to count by $N - 1$ or $N + 1$, depending on the desired direction of a slip. By forcing the same number and direction on the slips applied to the transmit and receive clocks, even if they are not precisely at the same time, the transmitted data stream is forced to be synchronous (albeit with a small phase jitter) to the receive data stream.

If we ignore the problem of the jitter effect on the EC, the nominal divide ratio can be chosen strictly on the basis of the effect of the phase jumps on the received data stream. The allowable phase jitter of the sampling clock depends on the horizontal eye opening, which in turn depends on the excess bandwidth. A jitter on the order of 3 percent of one baud interval should be quite tolerable, so that we might consider letting $N = 32$. If the maximum required frequency offset is one part in 10^5 , which should be easily obtainable with an inexpensive crystal oscillator, then in the worst case a slip will be required every 3125 baud intervals. The fact that these slips are very infrequent will be important in the sequel.

Fig. 2. A digital phase-locked loop (DPLL) in the network termination (NT) transceiver.

Unfortunately, the DPLL phase slips cause a temporary reduction in EC accuracy, as illustrated in Fig. 3, which shows the received sampling phase and transmit phase before and after one of these slips. Plotted on the vertical axis is the chosen phase out of N possibilities. The receive and transmit phase slip together, since synchrony between the two clocks is desired for loop timing. The effective impulse response of the echo depends on the relative transmit and received phase (not the absolute phases). If the EC has a memory of L baud intervals, then the EC accuracy is affected by the L past transmitted pulse phases and the current received sampling phase. Following the slip, the relative phases of the received sampling and the most recently transmitted pulses are the same as before the slip, and therefore the EC can accurately cancel the echo, due to these most recently transmitted pulses using the same coefficients as before the slip. The pulses transmitted earlier, however, enter the echo with a different phase and, hence, are less accurately canceled. The problem goes away after L baud intervals, however, since by then all the transmitted pulses within the memory of the EC have the correct phase relative to the received sampling.

Since the impact on EC accuracy is transient, one approach to jitter insensitive EC is therefore to force all slips to occur during a period of time where EC accuracy is not critical [4]–[6]. This requires a modification to the frame format, and an undesirable increase in the baud rate, which can be avoided using the technique proposed here.

B. Controlled Asynchronous Echo Cancellation

Given the configuration of Fig. 2, in principle the effect of jitter on the EC can be entirely eliminated. To see this, note that there are only N possible phases with which the data symbols are transmitted, and N phases for the receive sampler. If the transmit symbol modulates a pulse which is transmitted at time $(nN + k)/Nf_b$ for the n th symbol, where k is the transmitted phase out of N possibilities, then the echo pulse is of the form

$$r \left(t - \frac{nN + k}{Nf_b} \right)$$

for an echo response to a single transmitted pulse. If we sample this response with phase j out of N possibilities, then the sampling time is $(lN + j)/Nf_b$ for the l th baud interval. The echo response samples are therefore

$$r \left(\frac{(l-n)N + j - k}{Nf_b} \right).$$

Thus, if the EC knows or can adaptively learn the echo response to a single transmitted pulse sampled at a rate Nf_b , then the EC precisely cancels the echo for any transmit and

Fig. 3. Illustration of the relative phase error introduced by a phase jump in an NT DPLL.

receive phases k and j . This, of course, depends on the fact that the EC knows the values of k and j . The DPLL realization is fundamentally capable of achieving a higher accuracy of EC than the analog PLL.

Unfortunately, for a given maximum canceled echo delay, the scheme just described would increase the memory required for storage of the EC coefficients by a factor of N . This is because the echo response samples are required at a rate N times as great as normally required in a baud rate EC. We will show momentarily how this penalty of N can be reduced to only 2. First, however, note that the scheme does not increase the computational requirements of the EC. While the EC stores N sets of coefficients, one for each of the N possible relative transmit and receive phases, it need only calculate the echo replica once for each receive sample. It must keep track of the past transmitted phases, and use the appropriate coefficient for each delay. Also, since it can only update the coefficients that it uses in the calculation of the echo replica, the speed of convergence will be slowed. In fact, for a loop-timed transceiver there will be no opportunity to adapt the coefficients of the many relative phases which never occur, but for the same reason, neither must these phases be stored.

C. Limiting the Memory Penalty

Since the transmit and receive phases are controlled together in the NT, in fact it is not necessary to store and update N sets of EC coefficients. We will now show that two sets of coefficients will suffice.

The objectives that we must keep in mind are threefold.

1) We must store all the EC coefficients required to precisely cancel the echo for all possible relative transmit and receive phases.

2) We must have an opportunity to regularly adapt all the coefficients which are stored, so as to track accurately any time variation in the echo response due to changes in line temperature, etc.

3) We must be able to slip the phase in both directions.

The following approach will accomplish all these objectives. At the expense of a slight controlled jitter in the transmit data stream, or alternatively in the receive sampling, we can allow the EC to continuously adapt to two echo responses with adjacent phases. When a slip in one direction or the other is required, we specify a control policy for the sequence of transmitted and received phases which ensures that the two estimated echo responses are all that is required to precisely cancel the echo.

There are two variations on the technique we propose: one alternates the phase of the transmitted pulses, and the other alters the phase of the receive sampling. We explain the

former case first, since it is a little easier to understand, and then mention the latter.

Suppose that we alternate the transmit pulses between phase k and $k + 1$. More precisely, the transmitted pulses starting with the n th baud interval are transmitted at times

$$\frac{nN+k}{Nf_b}, \frac{(n+1)N+(k+1)}{Nf_b}, \frac{(n+2)N+k}{Nf_b}, \frac{(n+3)N+(k+1)}{Nf_b}, \dots$$

which requires that the divider counter for the transmit clock alternate between a count of $N + 1$ and $N - 1$. Further suppose that the receive sampler uses an arbitrary and constant phase j , i.e., samples at time $(nN + j)/Nf_b$ for the n th baud. Finally assume that the EC stores and adapts two sets of coefficients, one for each of the two transmitted phases. Let

$$g_0, g_1, g_2, \dots$$

be the set of coefficients of the EC for transmitted phase k and receive phase j , and let

$$h_0, h_1, h_2, \dots$$

be the set of coefficients for transmit phase $k + 1$ and receive phase j . These two sets of coefficients are illustrated in Fig. 4 for a portion of an echo pulse response. The echo replica calculation then uses a set of coefficients which alternates between the two sets of stored coefficients, since the transmit phases were alternated. Specifically, for the baud intervals where phase k was transmitted, the EC uses as its echo response for the purpose of calculating the echo replica

$$g_0, h_1, g_2, h_3, g_4, \dots$$

and for baud intervals where the transmitted phase was $k + 1$, the replica is calculated using the coefficients

$$h_0, g_1, h_2, g_3, h_4, \dots$$

In each case the replica is precise, independent of the transmitted phase jitter. Furthermore, the EC updates, using the cancellation error, whichever coefficients it uses in the calculation of the echo replica. In this fashion, each coefficient is updated precisely every second baud interval. This is why we alternated the transmit phases—to give an opportunity to update both sets of coefficients, thereby tracking any time variation of the echo response.

When the timing recovery circuit indicates a phase slip in the receive sampling is required, say in the direction from phase j to $j + 1$, then we must adjust the transmit phase in the same direction to maintain synchronism between the transmit and receive data streams. Thus, after the slip the transmit phase should be alternating between phases $k + 1$ and $k + 2$, in which the two sets of coefficients for receive phase $j + 1$ will be the same as before the slip, since the relative phases are unchanged. The phase of the transmit pulses and the receive sampling are plotted in Fig. 5 before the phase slip and after the slip. Note in both cases the alternating phase of transmitted pulses and the constant phase of receive sampling.

The transition from receive phase j to $j + 1$ requires some care to avoid the need for a new set of coefficients to maintain precise cancellation. The solution to this problem is also shown in Fig. 5. When the need for a slip is detected, we do not slip the receive sampling phase immediately. Before the receive phase is slipped, the transmit phase must be held constant at $k + 1$ for $L - 1$ baud intervals, where L is the maximum delay of cancellation. At the end of this interval, the receive phase is slipped, and the transmit phase resumes its alternation, this time between phase $k + 1$ and $k + 2$. The

Fig. 4. Sampled echo pulse response of three adjacent discrete phases for baud rate sampling.

Fig. 5. Sequence of transmit and receive phases corresponding to a phase jump in the NT which maintains exactly two relative phases.

period of constant phase ensures that the transmit and receive phases of k , $k+1$, and $k+2$ do not all appear within the memory of the EC at the same time, which would require three sets of coefficients instead of two.

Similarly, a slip in phase of the receive sampling from j to $j-1$ simply requires that the transmit phase be held constant at k for $L-1$ baud intervals, and then begin alternating between phases k and $k-1$ at the same time that the receive sampling phase is slipped.

A variation on the technique just described is to hold the transmit phase constant and continuously alternate the receive sampling between the two phases (just as in the lead-lag timing recovery technique [13]). The sequence of transmit and receive sampling phases corresponding to a phase slip is shown in Fig. 6. When the need for a slip is detected by the phase detector, the transmit phase is immediately slipped, but the receive sampling phase is temporarily held constant for the memory of the EC to avoid the introduction of a third relative phase. The receive sampling phase then resumes its alternation. The reader can verify that the relative phase between the transmit pulses and the current receive sampling assumes only two values throughout the phase slip.

Note that for both variations, each slip event actually takes approximately L baud intervals to achieve. This implies that the maximum rate of slip events is limited, with a limit that becomes more stringent as the length of the EC increases. Fortunately, as indicated earlier, for even modest clock accuracy requirements the worst case (shortest) interval between slips is very long. This is because the timing recovery circuit is not tracking actual phase jitter on the incoming data stream, but is tracking a small and nearly constant clock frequency offset.

Finally, it should be mentioned that the scheme just described requires a relatively complicated control to determine which set of coefficients to use during a phase slip. However, it is not necessary to alternate the phase between two phases on adjacent samples in the steady state. The alternation can occur less frequently, for example, once every L samples, which will simplify the control since the phases during a slip and in the steady state follow similar trajectories.

Fig. 6. Same as Fig. 5 except the receive sampling phase is alternated in the steady state rather than the transmitted pulse phase.

D. Canceler Convergence

Since each coefficient of the EC is only updated or adapted every second baud interval, the speed of convergence of the EC will be reduced by a factor of 2. This will present a problem only during the initial convergence at startup, since the steady-state variation of the echo response will be very slow.

During startup, the following modification is suggested. The difference between the g_j and h_j coefficients will be very slight, since they represent the samples of the echo response at slightly different phases. Therefore, during startup, simply require that the two sets of coefficients be equal to one another. When one coefficient is changed, then the other is changed by the same amount. (Equivalently, the deliberate jittering of the transmitted pulses or receive sampling can be suspended during timing acquisition.) The EC will then converge at the same rate as a conventional EC until the point at which the phase jitter becomes the limiting factor in the achievable EC accuracy. At this point, the EC can revert to the algorithm previously described.

E. Line Termination Transceiver

The scheme as described thus far applies only to the NT transceiver, which is loop timed to the received data stream. For this case the transmit phase is slaved to the receive phase, and therefore, the relative phase stays constant or almost constant. The approach described in Section II-C ensures that there are always just two relative phases between the transmitted symbols and received undesired interference or echo signal within the memory of the EC.

For the case of the LT transceiver, there are two cases to consider. The first case is where the crystal oscillator is replaced by a clock supplied by the switch, and this clock is jitter free or has only very low-frequency jitter components. The second case is where the LT transmit clock is derived from a switch clock using a PLL, in which case there is jitter on this clock (particularly if it is a DPLL). In either case the situation is still essentially as depicted in Fig. 2, but with two important distinctions.

- 1) The transmitted symbols must be synchronous with the switch clock, and therefore, while we are free to vary the countdown ratio from its nominal value of N , thereby introducing a slight amount of jitter on the transmit data stream, we can do so only if every count by $N+1$ is accompanied by another count by $N-1$, or vice versa, thereby keeping a constant average phase.

- 2) The incoming data stream is on the average synchronous with the transmitted data stream, since the remote NT transceiver is loop timed. The phase of the incoming data is unknown in advance, and depends on the precise propagation

delay from LT to NT and back. This propagation delay will vary slightly with changes in cable temperature, but these changes will be very slow, on the order of hours or days. There may also be a small jitter on the incoming data stream introduced by a phase-locked LT transmit clock or by the loop timing in the NT. From the perspective of received data detection, it should not be necessary to track this jitter.

One alternative we have here is to hold both transmitter and receiver countdown ratios constant, thereby allowing the sampling phase to be fixed over time but indeterminate relative to the phase of the received data stream. This approach will be suitable if the sampling rate is actually double the baud rate and fractionally spaced linear equalization [14] is used in the receiver, in which case the sampling phase can be arbitrary relative to the received data signal with little or no penalty in SNR after equalization. With this approach there will be no relative phase jitter between transmitter and receiver if the LT transmit clock is jitter free, and EC accuracy is not compromised. If the transmit clock is derived by a DPLL from the switch clock, then the situation is essentially the same as in the loop timed NT, and the same technique can be used.

For the remainder of this section, consider the alternative where either baud rate sampling is used in order to minimize the EC complexity, or a higher sampling rate without fractionally spaced equalization is used. This implies that we must adjust the received sampling phase to adjust for the arbitrary phase of the incoming data stream. Since the received data phase is nearly constant with time, the countdown ratio in the receiver can, in steady state, be held constant at N , except for an infrequent value of $N + 1$ or $N - 1$ to account for propagation delay changes. However, during the initial acquisition of receive timing, frequent adjustments of phase will be required before the steady-state phase is acquired.

These considerations imply that at the LT transceiver we must be prepared to change the relative phase of the transmitted and received data streams, albeit very infrequently in the steady state. In principle, then, we must be prepared to learn all N echo impulse responses, corresponding to the N received sampling phases, and not just two as in the NT transceiver. The transmit and receive sampling phases for this case are plotted in Fig. 7. The transmit phase can alternate as before, but when the receive sampling slips in phase, the transmit phase cannot also be slipped to hold the relative phase constant. As a result, each time the receive sampling slips, this introduces a new *relative* phase.

Let us see whether we can take advantage of the infrequent updates in phase to learn the new phases as we need them, rather than storing them perpetually. Refer back to Fig. 4, where the echo response $r(t)$ to a single transmitted pulse is shown. Using the technique described in Section II-C, two phases of this response separated by $1/Nf_b$ seconds can be learned, updated, and held in memory. As before, we denote these two phases by g_k and h_k . When it is required to slip the relative phase between transmitted data and receive sampling, inevitably a third phase of the echo response will enter the picture, which we have denoted by f_k . Depending on whether it is desired to advance or retard the phase, f_k may correspond to a sample just prior to or just after the g_k and h_k . For simplicity we treat only the latter case, as shown in Fig. 4.

We can attempt to interpolate f_k from the known samples as a method of avoiding storing all sampled phases of the echo response. In all cases of interest, the transmitted pulse will be band limited to f_b Hz, but (again for baud rate sampling) not band limited to $f_b/2$ Hz. For this case, we can in principle recover the f_k exactly since we have two samples per baud, even though they are not uniformly spaced [15]. Furthermore, the fact that the interpolated samples are close to the available samples should make our job easier.

Assume that the transmitted pulses are band limited to $(1 + \alpha)f_b/2$, where $0 \leq \alpha \leq 1$ is the fractional excess

Fig. 7. Same as Fig. 5 except that the relative transmit and receive phase is changing, as will be necessary in an LT transceiver in the absence of fractionally spaced equalization.

bandwidth of the data transmission system, and furthermore, that the sampling rate for the EC is nominally equal to the baud rate f_b . Then it is shown in the Appendix that a very accurate interpolation of f_k can be performed by the method of Fig. 8(a), where the g_k and h_k are applied to two filters $F_1(e^{j\lambda})$ and $F_2(e^{j\lambda})$ and the outputs are summed. The interpolated samples f_k can be exactly determined by this configuration and very accurate approximations can be obtained by low-order FIR filters. The required order of filter depends on the values of N and α , but three-tap FIR filters will suffice in most circumstances for $N = 64$, and five-tap filters may be required for $N = 32$. The adequacy of this simple filter is due to the fact that the interpolated sample and the two samples being used as the primary basis for the interpolation are so closely spaced in time. Any residual interpolation error will be corrected by the subsequent adaptation of the EC.

This interpolation introduces an additional computation every time a slip to a new relative transmit-receive phase is required. However, since in most circumstances the echo response is not expected to change over the period between slips, this computation need not be performed all at once, and need not appreciably affect the overall computation rate. Furthermore, if fractionally spaced equalization is used, it is not needed at all.

F. Voiceband Data Transmission

The conventional voiceband data modem requires that the two directions of transmission be asynchronous. Specifically, the transmitted data stream is synchronous with the local clock, and not locked to the received data stream, which was generated by an independent clock. In this respect, the voiceband data modem is similar to the line card digital subscriber loop transceiver with baud rate sampling, in that phase slips in the receive timing are required. A second difference is that the sampling rate is likely to be at least twice the baud rate, corresponding to the $\alpha = 0$ interpolation filter design described in the Appendix.

The technique which has been proposed for the subscriber loop application should be directly applicable to the voiceband modem. The differences are that a larger N can be used (since the baud rate is lower), simplifying the interpolation, but the worst-case clock offset is larger (resulting in more frequent slips). Due to the more frequent slips, it may not be necessary to alternate the transmit phase, but rather it can be held constant. For this case, the appropriate EC coefficient will depend only on the current (instantaneous) receive sampling phase, simplifying the control policy considerably. The EC coefficients for each new phase can be interpolated using the previous two phases as described in the Appendix. A remaining question is whether the EC can adapt quickly enough to overcome the accumulating interpolation error.

There are other approaches to the asynchronous EC problem for the voiceband data modem, including interpolation of the far-end signal by reconstruction of a continuous-

Fig. 8. Filters for interpolation of the echo pulse response. (a) Interpolation using two filters, one for each given phase. (b) Exact interpolation filter. (c) Exact interpolation filter for $F(e^{j\lambda})$ an all-pass filter.

time signal and resampling, interpolation in the sampled-data domain, and the adaptation of a fractionally spaced equalizer to the changing sampling phase using techniques similar to those described here.

III. CONCLUSIONS

A conceptually simple and easily implemented scheme for using a DPLL in timing recovery for a synchronous full duplex data transmission system has been demonstrated. In addition to allowing the use of the simpler DPLL technique, this scheme also completely eliminates the timing jitter as a limitation to the degree of EC enhancement which can be obtained. The practical significance of this is that the design of the timing recovery circuit is greatly simplified, and the design of the EC and data converters is also simplified, since a greater impairment from sources such as nonlinearities can be tolerated.

It is interesting that a similar scheme of alternating the sampling phase, called lead-lag sampling, has been proposed in the context of mean-square timing recovery [13], [16]. Lead-lag sampling may therefore serve the dual purpose of increasing EC accuracy and providing a robust method of baud rate timing recovery.

The price to be paid in this scheme is the extra memory required to store a second set of EC coefficients. This memory may be significant in some cases, and should be minimized by tricks such as storing the difference between the coefficients rather than the second set of coefficients, since this difference will have a much smaller dynamic range. An additional cost is the more complicated control circuitry for choosing the proper EC coefficients during a phase slip, although this cost can be minimized by choice of the control policy. Another issue which has not been addressed here is the combination of this technique with the several proposed approaches to nonlinear EC [12].

APPENDIX

If the transmitted pulses are band limited to $(1 + \alpha)f_b$ Hz, then the echo pulse response $r(t)$ is similarly band limited, and the Fourier transform of its samples at the baud rate is given by

$$\sum_k r \left(\frac{kN + n}{Nf_b} \right) e^{-j\lambda k} = f_b e^{j\lambda n/N} (R(\lambda f_b) + W_N^{-n} R^*(2\pi - \lambda) f_b) \quad (1)$$

for $0 \leq \lambda \leq \pi$ where

$$W_N = e^{j2\pi/N}.$$

This Fourier transform displays the aliasing which is inherent in sampling at less than the Nyquist rate.

Taking into account the Fourier transforms of the three signals shown in Fig. 8(a), two known and the third to be generated given by (1), we get for the region where $R((2\pi - \lambda)f_b) = 0$

$$e^{-j2\lambda/N} F_1(e^{j\lambda}) + e^{-j\lambda/N} F_2(e^{j\lambda}) = 1, \quad 0 \leq \lambda \leq (1 - \alpha)\pi \quad (2)$$

and for the region where there is aliasing distortion,

$$F_1(\lambda) = -e^{j2\lambda/N} W_N^{-1} \quad (3)$$

$$F_2(\lambda) = e^{j\lambda/N} (1 + W_N^{-1})$$

$$(1 - \alpha)\pi \leq \lambda \leq \pi.$$

We have some leeway in the low-frequency region since the two filters are not unique, and we can use that leeway to ensure that the filter has no phase discontinuities.

This establishes that the configuration of Fig. 8(a) can do an exact interpolation. The solution can be cast in the form of Fig. 8(b), where there are two delay filters, and the difference signal is filtered by $F(e^{j\lambda})$, which is an arbitrary filter for $0 \leq \lambda \leq (1 - \alpha)\pi$, and is

$$F(e^{j\lambda}) = e^{-j2\lambda/N}, \quad (1 - \alpha)\pi \leq \lambda \leq \pi \quad (4)$$

over the frequency band where there is aliasing. Since $F(e^{j\lambda})$ is an all-pass filter at aliasing frequencies, it is natural to assume that it is all-pass at all frequencies,

$$F(e^{j\lambda}) = e^{-j\theta(\lambda)} \quad (5)$$

where the phase $\theta(\lambda)$ is constrained by

$$\theta(\lambda) = \frac{2\pi}{N}, \quad (1 - \alpha)\pi \leq \lambda \leq \pi. \quad (6)$$

For this case, the two equivalent filters of Fig. 8(a) are easily manipulated into the form

$$F_1(e^{j\lambda}) = -e^{j\gamma(\lambda)} \quad (7)$$

$$F_2(e^{j\lambda}) = 2 \cos \left(\frac{\theta(\lambda)}{2} \right) e^{j\gamma(\lambda)/2}$$

where

$$\gamma(\lambda) = \frac{2\lambda}{N} - \theta(\lambda). \quad (8)$$

This form of the interpolation filter is shown in Fig. 8(c).

Note that for large N the angles $\theta(\lambda)$ and $\gamma(\lambda)$ are small, so that the two filters are given approximately by

$$F_1(e^{j\lambda}) \approx -1 \quad (9)$$

$$F_2(e^{j\lambda}) \approx 2$$

which is interestingly a linear interpolation from g_k and h_k to f_k . (In fact, for this case it is, strictly speaking, an extrapolation, but we call it an interpolation in anticipation of the more complicated filters to be derived which perform a true interpolation.) For sufficiently large N we would expect this linear interpolation to have sufficient accuracy. The desirable accuracy would be an interpolation error which is approximately 60–80 dB lower than the echo signal level (depending on the EC accuracy objective), for this would result in an interpolation error which is small relative to the normal error introduced by finite precision effects and the adaptation of the EC. For the digital subscriber loop with a baud rate of 160

Fig. 9. Interpolation error suppression versus frequency for an FIR filter approximation to $F_1(e^{j\lambda})$ and $F_2(e^{j\lambda})$ at frequencies below half the sampling rate (upper curve) and aliased frequencies (lower curve). (a) First-order linear interpolator, $N = 64$, $M = 0$. (b) Three-coefficient FIR interpolation filter, $N = 64$, $M = 1$, $\alpha = 0.5$. (c) Three-coefficient FIR interpolation filter, $N = 64$, $M = 1$, $\alpha = 0.25$. (d) Three-coefficient FIR interpolation filter, $N = 64$, $M = 1$, $\alpha = 0.0$. (e) Five-coefficient FIR interpolation filter, $N = 32$, $M = 2$, $\alpha = 0.5$.

kbits/s, the practical limit on N would be about 64, which corresponds to a 20 MHz clock (for voiceband data the practical value of N would be much larger, since the baud rate is more than an order of magnitude smaller). It is easy to calculate the interpolation error, which for any particular filters $F_1(e^{j\lambda})$ and $F_2(e^{j\lambda})$ has a Fourier transform

$$E(e^{j\lambda}) = E_1(e^{j\lambda})f_b R(\lambda f_b) + E_2(e^{j\lambda})f_b R^*((2\pi - \lambda)f_b) \quad (10)$$

where

$$E_1(e^{j\lambda}) = F_1(e^{j\lambda}) + e^{j\lambda/N} F_2(e^{j\lambda}) - e^{j2\lambda/N}$$

$$E_2(e^{j\lambda}) = F_1(e^{j\lambda}) + W_N^{-1} e^{j\lambda/N} F_2(e^{j\lambda}) - W_N^{-2} e^{j2\lambda/N}$$

have the interpretation as the transfer functions from the echo pulse response $r(t)$ to the interpolation error for, respectively, the low and aliased high frequencies in the echo pulse. The magnitudes of these two transfer functions in decibels are plotted in Fig. 9(a) for the simple linear interpolator and $N = 64$. The upper plot is the interpolation loss at input frequencies below half the sampling rate ($-20 \log_{10}(|E_1(e^{j\lambda})|)$), while the lower plot is the loss experienced by the aliased frequencies in the echo pulse ($-20 \log_{10}(|E_2(e^{j\lambda})|)$). One way to think of these plots is in terms of an input sinusoidal echo pulse, and as the frequency increases, the loss to interpolation error follows the upper curve, and then as the frequency exceeds the half sampling rate, it aliases back down below the

TABLE I
COEFFICIENTS OF FIR APPROXIMATION TO INTERPOLATION FILTERS

α	N	M	k	a_k	b_k
0.00	64	0	0	-1.0000000	2.0000000
1	64	0	-1	0.0112353	-0.0106553
			0	-0.9998723	1.9987316
			1	-0.0113630	0.0119237
2	64	0	-2	-0.0013486	0.0013235
			-1	0.0104850	-0.0098837
			0	-0.9998880	1.9987395
			1	-0.0105142	0.0111028
			2	0.0012658	-0.0012821
3	64	0	-3	0.0004168	-0.0004127
			-2	-0.0011414	0.0011127
			-1	0.0104242	-0.0098217
			0	-0.9998897	1.9987403
			1	-0.0104481	0.0110382
			2	0.0010420	-0.0010630
			3	-0.0004038	0.0004062

α	N	M	k	a_k	b_k
0.25	64	0	0	-1.0000000	2.0000000
1	64	0	-1	0.0178671	-0.0171508
			0	-0.9996749	1.9982800
			1	-0.0181923	0.0188708
2	64	0	-2	0.0002902	-0.0002418
			-1	0.0174372	-0.0168062
			0	-0.9996959	1.9983522
			1	-0.0174217	0.0180287
			2	-0.0006097	0.0006671
3	64	0	-3	-0.0005874	0.0005686
			-2	-0.0000822	0.0001504
			-1	0.0176546	-0.0170248
			0	-0.9996879	1.9983409
			1	-0.0176521	0.0182569
			2	-0.0002505	0.0003295
			3	0.0006055	-0.0006214

α	N	M	k	a_k	b_k
0.50	64	0	0	-1.0000000	2.0000000
1	64	0	-1	0.0187119	-0.0179780
			0	-0.9996431	1.9982157
			1	-0.0190688	0.0197623
2	64	0	-2	0.0047947	-0.0046118
			-1	0.0253054	-0.0247156
			0	-0.9993377	1.9978339
			1	-0.0250685	0.0256178
			2	-0.0056940	0.0058758
3	64	0	-3	0.0003923	-0.0003681
			-2	0.0049289	-0.0047902
			-1	0.0244605	-0.0238663
			0	-0.9993802	1.9978930
			1	-0.0242011	0.0247586
			2	-0.0055003	0.0056448
			3	-0.0007000	0.0007282

α	N	M	k	a_k	b_k
0.00	32	0	0	-1.0000000	2.0000000
1	32	0	-1	0.0223387	-0.0199814
			0	-0.9994893	1.9949293
			1	-0.0228494	0.0250521
2	32	0	-2	-0.0027802	0.0026628
			-1	0.0209377	-0.0185089
			0	-0.9995522	1.9949605
			1	-0.0210543	0.0233820
			2	0.0024490	-0.0024964
3	32	0	-3	0.0008478	-0.0008279
			-2	-0.0023818	0.0022518
			-1	0.0208213	-0.0183881
			0	-0.9995589	1.9949638
			1	-0.0209167	0.0232507
			2	0.0019839	-0.0020523
			3	-0.0007957	0.0008021

α	N	M	k	a_k	b_k
0.25	32	0	0	-1.0000000	2.0000000
1	32	0	-1	0.0353917	-0.0324539
			0	-0.9986997	1.9931258
			1	-0.0366919	0.0393281
2	32	0	-2	0.0002549	-0.0000795
			-1	0.0348736	-0.0323032
			0	-0.9987840	1.9934134
			1	-0.0348117	0.0371894
			2	-0.0015328	0.0017799
3	32	0	-3	-0.0011501	0.0010701
			-2	-0.0004964	0.0007469
			-1	0.0352947	-0.0327273
			0	-0.9987520	1.9933681
			1	-0.0352846	0.0376515
			2	-0.0008340	0.0011705
			3	0.0012225	-0.0012799

α	N	M	k	a_k	b_k
0.50	32	0	0	-1.0000000	2.0000000
1	32	0	-1	0.0370468	-0.0340334
			0	-0.9985726	1.9928692
			1	-0.0384742	0.0411642
2	32	0	-2	0.0086572	-0.0079258
			-1	0.0507852	-0.0483464
			0	-0.9973525	1.9913444
			1	-0.0498385	0.0519539
			2	-0.0122515	0.0129739
3	32	0	-3	0.0004788	-0.0003897
			-2	0.0092647	-0.0087230
			-1	0.0491331	-0.0466845
			0	-0.9975221	1.9915795
			1	-0.0480965	0.0502515
			2	-0.0115492	0.0121366
			3	-0.0017089	0.0018296

half sampling rate and follows the lower curve. The two curves meet at the half sampling rate because there is no discontinuity in the interpolation error as the input frequency passes through the half sampling rate.

The linear interpolation is clearly not adequate unless we are willing to allow a large decrease in the EC accuracy immediately following the slip in phase until the EC coefficients are allowed to readapt to the new phase. One approach that was tried to improve the interpolation accuracy was to use the Lagrange interpolation formula to derive higher order interpolators which include g_{k-1} , g_{k+1} , and so forth in the interpolation. This approach was not found to introduce an appreciable improvement because it does not directly take into account the aliasing distortion.

An alternate approach which worked beautifully was to design FIR filters to directly approximate the desired responses of (6) and (7). In particular, to ensure that the phase was smooth and had no discontinuities, for a given α the phase was chosen to be

$$\theta(\lambda) = \begin{cases} \sin\left(\frac{\lambda}{2(1-\alpha)}\right) & 0 \leq \lambda \leq (1-\alpha)\pi \\ \frac{2\pi}{N} & (1-\alpha)\pi \leq \lambda \leq \pi. \end{cases} \quad (11)$$

This choice ensures that the phase and the derivative of the phase are continuous at all frequencies, making the response easy to approximate accurately even with a low-order FIR filter. Note also from (6) and (7) that, fortuitously, the phase is continuous at $\lambda = \pi$.

For the filters of (7), (8), and (11), the actual FIR approximation to the transfer functions $F_1(e^{j\lambda})$ and $F_2(e^{j\lambda})$ were calculated using the frequency sampling method [17]. In each case a noncausal filter with $2M + 1$ coefficients of the form

$$\sum_{k=-M}^M a_k e^{-jk\lambda}$$

was used, where the transfer function matched the desired transfer function at the $2M + 1$ frequencies $e^{j2\pi m/(2M+1)}$, $0 \leq m \leq 2M$. Specifically, the approximation to $F_1(e^{j\lambda})$ is given by the inverse DFT

$$a_k = \frac{1}{2M+1} \sum_{m=0}^{2M} F_1(e^{j\lambda}) e^{j2\pi mk/(2M+1)}, \quad -M \leq k \leq M$$

with a similar approximation to $F_2(e^{j\lambda})$ with coefficients b_k . The resulting design reduces the interpolation error dramatically even for small values of M . For example, the interpolation error for the same value $N = 64$ is shown in Fig. 9(b)-(d) for $M = 1$ (a three-tap FIR filter). Three values of excess bandwidth α are shown, $\alpha = 0.5$ (50 percent excess bandwidth) in Fig. 9(b), $\alpha = 0.25$ in Fig. 9(c), and $\alpha = 0$ in Fig. 9(d). The latter case of no excess bandwidth would be appropriate if the sampling rate actually exceeded the Nyquist rate; for example, if the sampling rate were twice the baud rate. As α decreases, the minimum EC enhancement improves, as one might expect. The EC enhancement is inadequate for the aliased frequencies between 0 and $(1-\alpha)\pi$, but recall that the transmitted data signal and, hence, the echo pulse does not include these frequency components.

While the interpolation error is probably adequate for $M = 1$, the interpolation error decreases dramatically for larger values of M . This can be exploited to reduce the value of N , which is related to the clock frequency, at the expense of additional computational complexity for the interpolation. For example, the case $N = 32$, $M = 2$, and $\alpha = 0.5$ is shown in Fig. 9(e). The coefficients of all the filters plotted in Fig. 9 are given in Table I.

The actual total interpolation error depends on the spectrum of the echo pulse. As an example, the interpolation error was calculated in the time domain, using some of the filter designs given in Table I, for a raised cosine received echo pulse with the same excess bandwidth as assumed in the FIR filter design. The results are shown in Table II. The tradeoff between N and M is clearly displayed in this table.

TABLE II
INTERPOLATION ERROR FOR FIR INTERPOLATION FILTERS FOR A RAISED
COSINE ECHO PULSE RESPONSE

α	N	M	Interpolation error (dB)
0.5	64	0	58
		1	75
		2	86
0.25	64	0	59
		1	77
		2	93
0.0	64	0	59
		1	91
		2	100
0.5	32	0	46
		1	63
		2	73
		3	94

REFERENCES

- [1] D. G. Messerschmitt, "Echo cancellation in speech and data transmission," *IEEE J. Select. Areas Commun.*, vol. SAC-2, pp. 283-297, Mar. 1984.
- [2] D. D. Falconer, "Timing jitter effects on digital subscriber loop echo cancellers," *IEEE Trans. Commun.*, vol. COM-33, pp. 826-838, Aug. 1985.
- [3] N.-S. Lin, D. Hodges, and D. G. Messerschmitt, "Partial response in digital subscriber loop," in *Proc. IEEE GLOBECOM*, Dec. 1985.
- [4] 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.
- [5] L. Lerach, H. Sailer, and G. Schollmeier, "VLSI components for the ISDN customer access," in *Proc. IEEE GLOBECOM*, Nov. 1984.
- [6] 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.
- [7] D. G. Messerschmitt, "Design issues for the ISDN U-interface transceiver," *IEEE J. Select. Areas Commun.*, vol. SAC-4, pp. 1281-1293, Nov. 1986.
- [8] 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.
- [9] 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.
- [10] R. Carpenter, S. Cox, and P. Adams, "Jitter compensation in echo cancelers," in *Proc. IASTED Int. Symp. Appl. Signal Processing, Digital Filtering*, June 1985.
- [11] S. Cox, "Clock sensitivity reduction in echo cancelers," *Electron. Lett.*, vol. 21, July 4, 1985.
- [12] P. Adams, R. Carpenter, and S. Cox, "Adaptive cancellation of non-linear timing-varying echos in digital data transmission," presented at IEEE Colloq. Adaptive Filters, Oct. 10, 1985.
- [13] H. Sari, L. Desperben, and S. Moridi, "Optimum timing recovery for digital equalizers," in *Proc. IEEE GLOBECOM*, Dec. 1985.
- [14] G. Ungerboeck, "Fractional tap-spacing equalizer and consequences for clock recovery in data modems," *IEEE Trans. Commun.*, vol. COM-24, p. 856, Aug. 1976.
- [15] D. G. Messerschmitt, "Approaches to multiple 8 kHz nonuniform sampling in PCM channel banks," *IEEE Trans. Circuit Theory*, Jan. 1975.
- [16] R. Gitlin and J. Salz, "Timing recovery in PAM systems," *Bell Syst. Tech. J.*, vol. 50, p. 1645, May-June 1971.
- [17] L. R. Rabiner and B. Gold, *Theory and Application of Digital Signal Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1975.

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 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 also 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 the IEEE COMMUNICATIONS MAGAZINE, and is past Editor for Transmission of the IEEE TRANSACTIONS ON COMMUNICATIONS and past member of the Board of Governors of the IEEE Communications Society.