

Equalizer Training Algorithms for Multicarrier Modulation Systems

Jacky S. Chow, John M. Cioffi*, and John A.C. Bingham
Amati Communications Corporation, Palo Alto, CA 94301

Abstract

In a multicarrier data transmission system the length of a symbol is often limited by the maximum permissible delay of data from input to output. If, as is usual, the length of the impulse response of the channel is not negligible compared to the permitted symbol length, the impulse response can be shortened by passing the received signal through a time-domain equalizer before demodulation. The best performance for a given computational complexity can then be achieved by appending to each block of samples of the transmit signal a cyclic prefix that is the same length as the shortened impulse response.

This paper describes several algorithms for designing the time-domain equalizer, which use adaptation in the frequency domain and windowing in the time domain in order to minimize the mean squared error of the equalized response.

1 Introduction

Equalization of a multicarrier signal that has passed through a distorting channel was originally discussed by Hirosaki [1], but the equalization of the particularly efficient (in terms of both bandwidth utilization and amount of computation) form of multicarrier that is based on fast Fourier transforms was first considered in [2]. This paper describes algorithms based on [2] that have been implemented in commercially available digital signal processors at data rates up to 6.3 Mbit/s.

The second section defines the problem, and explains the basic approach in a way that is different from, but complementary to, that in [2]. The third section describes in detail the basic steps from which four algorithms can be constructed; Section 4 discusses the relative merits of these algorithms; Section 5 reports on results obtained with high-speed digital subscriber line (HDSL) and asymmetric digital subscriber line (ADSL) systems; Section 6 extrapolates from the results to suggest some future work.

*John M. Cioffi is also with Stanford University

2 Background and Definition of the Problem

Although a multicarrier signal is most conveniently generated using an inverse FFT (IFFT), there are several conceptual advantages to considering it as the sum of a set of Quadrature Amplitude Shift Keyed (QASK) signals¹, for which the baseband pulse shape is a rectangle of duration T . Similarly, it is useful to consider the FFT in the receiver as performing conventional demodulation of each of the carriers followed by a baseband matched filter, which is implemented as an integrate-and-dump. The frequency-separation, $\Delta f (= 1/T)$, of the carriers is such that if the rectangular pulses are not distorted the baseband signals are orthogonal, and there is no inter-carrier interference (ICI). The problem, therefore, is to maintain—or restore—this orthogonality if the channel distorts the pulses.

A theoretically ideal multicarrier system would use an infinitesimally small frequency spacing between carriers and, consequently, an infinite-length symbol. It would be immune to distortion because the (finite) length of the impulse response of the channel (the transient at the beginning of each set of data symbols) would be negligible compared to the length of the symbol over which the integration is performed. Most systems, however, have a maximum tolerable delay from data input to output (latency), and since the end-to-end delay through a multicarrier transceiver is typically three symbols, this defines a maximum symbol length.

If, as is usual, the duration of the impulse response of the channel, T_{imp} , is not negligible compared to the permissible symbol duration, T_{max} , there are three fairly straightforward ways of avoiding ICI and ISI (interference between successive symbols modulated onto the same carrier)²:

- A. Fully equalize the channel with a conventional adaptive tapped delay line; this has the disadvantage that the amount of computation may be too great for programmable DSP implementation at the data rates required for HDSL or ADSL.
- B. Use a symbol (baseband pulse) length, T' , that is greater than T , and perform the integration (i.e., collect samples for the FFT) over only the latter T of each

¹This approach to multicarrier modulation is described in more detail in [3].

²There are two other methods (see [4], [5] and [6], and [7]), which require considerably more computation, but they have not been used for HDSL systems, and are beyond the scope of the present paper.

symbol, when the transient response of the channel is assumed to have subsided to a negligible level; that is, $T' = (1 + \nu/N)T$, where N = the number of samples used for the IFFT and FFT.

Because every one of the carriers has an integer number of cycles in the integration period, T , this extending of the baseband pulses can be easily implemented by cyclically prefixing the block of N samples of the output of the IFFT by a repeat of the last ν samples. Since these added samples are redundant this reduces the throughput efficiency to $\alpha = N/(N + \nu) = T_{max}/(T_{max} + T_{imp})$.

For many systems—and particularly for subscriber line applications—the latency requirements and the distortion of the channel are such that this efficiency would be intolerably low.

- C. For any given system and computational capability, [2] showed that the greatest throughput can be achieved by a combination of methods A and B: that is, by using a fairly short (i.e., computationally feasible) equalizer to constrain the impulse response of the channel to just $\nu + 1$ samples, and using a cyclic prefix of ν samples. A transceiver system that uses this method was described in [8].

A partial equalizer, which shortens the impulse response of a channel to some pre-defined length, is also needed as the pre-filter for Maximum Likelihood Sequence Detection (MLSD) receivers, in which a Viterbi decoder is applied to M^ν states, where M is the number of levels for a PAM system. Design methods for these equalizers were described in [9], [10], and [11]. These methods all used Least Mean Squares (LMS) adaptation in the time domain for learning the parameters of both the equalizer and the Shortened Impulse Response (SIR), and suffered from slow and uncertain convergence.

3 Equalizer Training Algorithms

The training of an equalizer for data transmission usually consists of two parts: the major part during start-up using a carefully chosen data sequence that is known by the receiver, and then minor adjustments during data transmission to correct for slow, small changes of the channel. This paper discusses only the initial training; a future paper will discuss adaptation during data transmission.

The discrete-time channel impulse response is given by $h_i, i = 0, 1, \dots, K$. The goal of these training algorithms is to design a finite-impulse response (FIR) equalizer $w(D)$ with $L + 1$ taps so that the equalized channel response, $w(D) * h(D)$, approximates any arbitrary SIR ($b(D)$), which has only $\nu + 1$ contiguous non-zero samples.

Figure 1 shows the basic training algorithm. A pseudo-random (binary) sequence is used to generate a training vector X of length N . An inverse fast Fourier transform (IFFT) of length N is performed to obtain the time-domain training vector $x(D)$. This training vector $x(D)$ is then sent repeatedly over the channel to form a periodic signal of period N .




Figure 1: A basic block diagram of the training algorithms

At the receiver, the received vector $y(D)$ of length N should strictly be defined by

$$y(D) = x(D) * h(D) + n(D), \quad (1)$$

where $n(D)$ is the received noise vector. For most communication channels with additive noise of zero mean, however, the effect of noise becomes negligible when we average the observation of $y(D)$ over a long period of time. Then $y(D)$ can be considered as the cyclic convolution of $x(D)$ and $h(D)$. A local copy of $x(D)$ at the receiver, along with $y(D)$, is used to train the SIR, $b(D)$, and the FIR equalizer $w(D)$ as follows:

- i. B , which is the discrete Fourier transform (DFT) of $b(D)$ (of length N), is updated in the frequency domain using the current values of X , Y , and W (DFTs of $x(D)$, $y(D)$, and $w(D)$, respectively).
- ii. A windowing is performed on $b(D)$ to limit the SIR to have $\nu + 1$ or fewer non-zero samples.
- iii. W , the frequency domain version of the FIR equalizer $w(D)$, is updated using the current values of X , Y , and B .
- iv. A windowing is performed on $w(D)$ to limit the FIR equalizer to have $L + 1$ or fewer non-zero taps.

The above steps are preceded by setting $b(D)$ and $w(D)$ to some initial values, and the steps are repeated until suitable

criteria have been met. Section 3.2 describes one specific criterion for monitoring the convergence of these algorithms.

Each of the above steps can be realized by more than one method. The following subsections describe some of the methods for each step.

3.1 Update B

We have used two methods for updating B : (1) frequency domain LMS and (2) frequency domain division.

(1) Frequency domain LMS




Figure 2: Update B : Frequency domain LMS

Figure 2 shows the block diagram of the frequency domain LMS update. We included a subscript for both $b(D)$ and $w(D)$ to clarify whether the samples/taps have been updated (u) or updated and windowed (w). We define an error signal E by

$$E = B_w X - W_w Y, \quad (2)$$

where B_w is the FFT of $b_w(D)$, the windowed SIR, and W_w is the FFT of $w_w(D)$, the windowed FIR equalizer.

The LMS update is then

$$B_{u,k+1} = B_{w,k} + 2\mu_b E X^*, \quad (3)$$

where k is the time index in symbols (blocks), μ_b is the stepsize for updating B , X^* is the complex conjugate of X , and the multiplication $E X^*$ is performed component by component.

(2) Frequency domain division

There are two equivalent ways to perform the frequency domain division. We either form $z(D) = w_w(D) * y(D)$ in the time-domain, as illustrated in figure 3, and then perform an FFT on $z(D)$ to obtain Z , or we perform the FFT of $y(D)$ and $w_w(D)$ to obtain $Z = W_w Y$. In either case,

$$B_{u,k+1} = \frac{W_{w,k} Y}{X}, \quad (4)$$

where all multiplication and division are done component by component.




Figure 3: Update B : Frequency division

3.2 Window $b(D)$

To limit the number of non-zero samples of the SIR $b(D)$, we perform a windowing as shown in Figure 4:

- A time-domain updated response $b_u(D)$ is formed by taking the inverse FFT of the updated B .
- The $\nu + 1$ consecutive taps with the highest total energy are found by a cyclic search through the N samples of $b(D)$.
- A time-domain windowing is performed. The simplest windowing function is a rectangular window, which simply zeros all the other $N - (\nu + 1)$ samples. Other windowing functions may be used, and the relative merits of various functions are under current investigation. The ratio of the energy outside the window (which is discarded) to that in the $\nu + 1$ samples is used as the measure of convergence of the adaptation.
- To prevent the SIR and the FIR equalizer from converging to the trivial solution, $E = B = W = 0$, we restrict the total energy of $b_w(D)$ to be either above a certain threshold limit or normalized to some preset value.
- An FFT is then performed on $b_w(D)$ to convert it back to the frequency domain, B_w .




Figure 4: Window B: time-domain windowing

3.3 Update W

Methods of updating W are similar to those for B . With frequency-domain LMS,

$$E = B_w X - W_w Y, \quad (5)$$

and

$$W_{u,k+1} = W_{w,k} + 2\mu_w E Y^*. \quad (6)$$

Similarly, with frequency-domain division,

$$W_{u,k+1} = \frac{B_{w,k} X}{Y}. \quad (7)$$

There is a subtle difference between the frequency-domain division on W and on B , when actual implementation issues are considered. We will explore some of these issues in Section 4.

3.4 Window $w(D)$

Parallel to Section 3.2, windowing $w(D)$ can be performed as follows:

- A time-domain updated response $w_u(D)$ is formed by taking the inverse FFT of the updated W .
- The $L+1$ consecutive taps with the highest total energy are found by a cyclic search through $w_u(D)$.
- A time-domain windowing is performed. The simplest windowing function is a rectangular window, which simply zeros all the other $N-(L+1)$ taps. Other windowing function may also be used.

- If we monitor the total energy on $b_w(D)$, then there is no need also to monitor the total energy on $w_w(D)$ since they track each other.
- An FFT is then performed on $w_w(D)$ to convert it back to the frequency domain, W_w .

4 Discussion

There are four possible combinations for updating W and B using the two algorithms described in Section 3:

1. frequency-domain division for W and B
2. frequency-domain division for W and frequency-domain LMS for B
3. frequency-domain division for B and frequency-domain LMS for W
4. frequency-domain LMS for W and B

In terms of convergence, method 1 is the most unstable since there is no direct control of the rate of adaptation. Both W and B are trying to converge in one step (division) and they could end up conflicting with each other. Methods 2 - 4 offer various degrees of control for the rate of adaptation by means of stepsizes in the LMS updates. With a careful selection of stepsizes, all these three methods should converge to the same result. However, the values of the stepsizes are the most critical with method 4 since the relative stepsizes between updating W and B could result in similar conflicts as in the case of method 1.

Methods 2 and 3 are very similar but there are at least two important differences. In method 2, B is the driving force and W is the follower with the division (unless the stepsize on B is too large, which then could be unstable), whereas the reverse is true in method 3. For the case when we desire a much shorter SIR than the FIR equalizer ($\nu \ll L$), method 2 would appear to be superior because the windowing on the SIR is likely to be more drastic than the windowing on the FIR equalizer. With method 2, W can then be adapted to the new, windowed B quickly³. The reverse is true when $L \ll \nu$.

Another difference between methods 2 and 3 is a practical consideration. In most cases, a division requires much more computational power than a multiplication. We could rewrite equations (4) and (7) as follows:

$$B_{u,k+1} = \frac{W_{w,k} Y X^*}{|X|^2}, \quad (8)$$

$$W_{u,k+1} = \frac{B_{w,k} X Y^*}{|Y|^2}, \quad (9)$$

where X^* is the complex conjugate of X , and $|X|^2$ is the square magnitude of X . Again, all multiplications are done component by component.

³The limiting case of $\nu = 0$ corresponds to a conventional equalizer for symbol-by-symbol detection of single-carrier signals. For this the components of $B_{w,k}$ would all be equal, and the updating of W would be equivalent to inverting the frequency response of the channel.

Equation (8) would appear to require a division, but since X is a repetitive and known training sequence, the division for method 3 can be implemented by multiplication by the stored inverse of $|X_i|^2$. On the other hand, Y also has the effect of the channel H and noise on it, and its inverse cannot be stored; equation (9) does require a division. Therefore, method 3 is usually easier to implement than method 2.

5 Results

We have used Method 3 (division for B and LMS for W) in receivers for both HDSL (a dual-simplex system at 2.048 Mbit/s in each direction) and ADSL (a frequency-division-multiplexed system combining conventional telephone service—known colloquially as POTS—with duplex transmission of up to 6 Mbit/s from central office to remote terminal and 28 kbit/s from remote to central). The criterion of performance of an equalizer was defined as the decrease in equivalent SNR resulting from remanent distortion; it was calculated by simulation and measured on actual hardware/software for all systems.

For the HDSL system the equalizers have to deal only with the distortion of the loop. The decrease in SNR was less than 0.25 dB. For the ADSL system the performance was very different in the two directions. To receive the wide-band downstream ADSL signal the equalizer must compensate for the loop plus the sharp low-end cut-off of the filters used to remove the transmitted upstream signal. The equalizer was somewhat less successful at this; the degradation was about 0.5 dB. On the other hand, to receive the narrower upstream signal the equalizer must compensate for very sharp cut-offs at both the low end (to remove the POTS signal) and the high end (to remove the transmitted downstream signal). It was much less successful at this; we judged the degradation to be about 3 dB; several of the carriers near the edges of the band were limited in their capacity by remanent distortion rather than by noise.

6 Future Work

From observation of the convergence of the upstream ADSL equalizer we hypothesized that the attempt to simultaneously learn both w and b may result in many local minima of the error. When there is severe line distortion, the minima may be very different, and the algorithm rarely finds the smallest.

Since the only important criterion for $b(D)$ is that it be limited to $\nu + 1$ samples, we are investigating a modification of the method of Figure 1 in which the update of W tries to drive the rest of the samples to zero. The steps of the method are:

1. Calculate B^4 by frequency-domain division as shown in equation (4).

⁴Since we are not trying to learn B , “update” is no longer the appropriate word.

2. Wall $b(D)$.⁵ Perform the first two steps of Section 3.2, and then calculate the error for the LMS update of W as the transform of the wall part of $b(D)$.

3. Update and window W as in method 3 of Section 4. Preliminary simulations of this method have been very promising; details will be given in a later paper.

References

- [1] B. Hirosaki. An Analysis of Automatic Equalizers for Orthogonally Multiplexed QAM Systems. *IEEE trans on Communications*, 28(1):73–83, January 1980.
- [2] J. S. Chow and J.M. Cioffi. A Cost-Effective Maximum Likelihood Receiver for Multicarrier Systems. In *1992 International Conference on Communications*, pages 948–952, Chicago, June 1992.
- [3] J.A.C. Bingham. Multicarrier Modulation for Data Transmission : An Idea Whose Time Has Come. *IEEE Communications Magazine*, 28(5):5–14, May 1990.
- [4] J.L. Holsinger. Digital Communication over Fixed Time-continuous Channels with Memory—with Special Application to Telephone Channels. *Lincoln Lab. Technical Report*, 366, October 1964.
- [5] S. Kasturia, J. Aslanis, and J.M. Cioffi. Vector Coding for Partial-Response Channels. *IEEE Trans. on Information Theory*, 36(4):741–762, July 1990.
- [6] J.W. Lechleider. The Optimum Combination of Block Codes and Receivers for Arbitrary Channels. *IEEE Trans. on Communications*, 38(5):615–621, May 1990.
- [7] D. Chafee. A Multicarrier Approach for V.fast. In *Report of CCITT Special Rapporteur’s meeting*, Paper L.111, Fort Lauderdale, February 1991.
- [8] J.S. Chow, J.C. Tu, and J.M. Cioffi. A Discrete Multitone Transceiver System for HDSL Applications. *IEEE Journal on Selected Areas in Communications*, 9(6):895–908, August 1991.
- [9] D.D. Falconer and Jr. F.R. Magee. Adaptive Channel Memory Truncation for Maximum Likelihood Sequence Estimation. *Bell Syst. Tech. J.*, 52:1541–1562, November 1973.
- [10] D.G. Messerschmitt. Design of a Finite Impulse Response for the Viterbi Algorithm and Decision Feedback Equalizer. In *1974 International Conference on Communications*, pages 37D.1–37D.5, June 1974.
- [11] J.A.C. Bingham. *The Theory and Practice of Modem Design*. John Wiley & Sons, New York, NY, 1988.

⁵A “walling” operation is the complement of windowing: the retained samples are those that are outside the chosen window (i.e., they would splatter on the wall).